From Perception to Protection: A Developer-Centered Study of Security and Privacy Threats in Extended Reality (XR)

Kunlin Cai[†], Jinghuai Zhang[†], Ying Li[†], Zhiyuan Wang[¶], Xun Chen[§], Tianshi Li[‡], and Yuan Tian[†]
[†]University of California, Los Angeles. Email: {kunlin96, jinghuai1998, yinglee, yuant}@g.ucla.edu
[‡]Northeastern University. Email: tia.li@northeastern.edu
[§]Independent Researcher. Email: xunchen@outlook.com
[¶]University of Virginia. Email: vmf9pr@virginia.edu

***Abstract*—The immersive nature of XR introduces a fundamentally different set of security and privacy (S&P) challenges due to the unprecedented user interactions and data collection that traditional paradigms struggle to mitigate. As the primary architects of XR applications, developers play a critical role in addressing novel threats. However, to effectively support developers, we must first understand how they perceive and respond to different threats. Despite the growing importance of this issue, there is a lack of in-depth, threat-aware studies that examine XR S&P from the developers' perspective. To fill this gap, we interviewed 23 professional XR developers with a focus on emerging threats in XR. Our study addresses two research questions aiming to uncover existing problems in XR development and identify actionable paths forward.**

By examining developers' perceptions of S&P threats, we found that: (1) XR development decisions (e.g., rich sensor data collection, user-generated content interfaces) are closely tied to and can amplify S&P threats, yet developers are often unaware of these risks, resulting in cognitive biases in threat perception; and (2) limitations in existing mitigation methods, combined with insufficient strategic, technical, and communication support, undermine developers' motivation, awareness, and ability to effectively address these threats. Based on these findings, we propose actionable and stakeholder-aware recommendations to improve XR S&P throughout the XR development process. This work represents the first effort to undertake a threat-aware, developer-centered study in the XR domain—an area where the immersive, data-rich nature of the XR technology introduces distinctive challenges. [1]

I. Introduction

With technological advancements, extended reality (XR) has become increasingly accessible in various aspects of daily life, offering immersive experiences that blur the boundaries between physical and digital worlds. XR developers have been developing numerous XR applications which revolutionizing how people learn, work, and interact. By 2025, the global XR market is projected to reach USD 87.3 billion, with the applications segment accounting for the largest share at 65%, signaling significant growth in XR applications [1]. However, the rapid expansion of the XR applications is a double-edged sword, offering substantial benefits while introducing notable security and privacy (S&P) risks [2].

The concerns related to S&P in XR applications—encompassing augmented reality (AR), virtual reality (VR), and mixed reality (MR)—require a dedicated investigation due to their unique characteristics: (1) XR provides users with unparalleled *immersive* experiences [3] by enabling innovative interaction designs unique to spatial computing, such as life-like avatar embodiment [4] and gaze-based or emotion-driven interactions [5], [6]. These advancements introduced new security concerns, including but not limited to XR side-channel attack [7], [8], [9], immersive digital manipulation [10], [11], [12], identity threats [13], and intellectual property threats from blending real and virtual content [14], [15]. (2) XR inherently requires an extensive collection of *multimodal user and environment data*, such as gestures, gaze, voice, physiological signals, location, and movement. The combination of these data streams provides a granular, real-time representation of users' state, where prior studies [16], [17], [18], [19] have shown that such fine-grained data opens up new avenues for attackers to exploit XR systems and compromise user privacy.

The interaction designs and data collection channels in XR are primarily determined by developers at the application level. Since developers play a key role in creating these designs [20], their design choices directly shape the S&P posture of applications. However, current XR research [21], [22] primarily focuses on user-centered studies, aiming to understand the S&P issues that users care about. Although existing studies provide valuable insights from the user perspectives, they often fall short in connecting S&P issues with the actual development process and uncovering the key factors hindering S&P development. Therefore, there's an urgent need for dedicated developer-centered studies to gain a deeper understanding of the causes behind emerging threats in XR applications.

A portion of this work was conducted during Kunlin Cai's internship with Xun Chen at Samsung Research America.

[1]The full version can be found at https://arxiv.org/abs/2509.06368.

Network and Distributed System Security (NDSS) Symposium 2026
23-27 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.230807
www.ndss-symposium.org

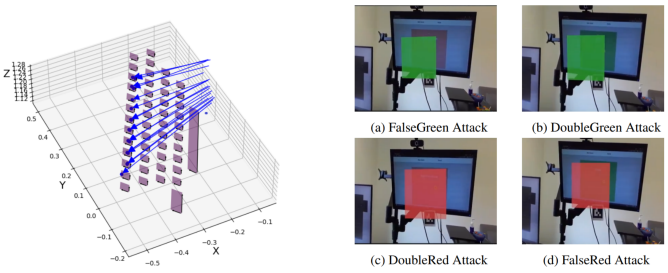

(a) VR Keylogging Attack [8] (b) AR Perception Attack [10]

Fig. 1: Attack examples shown to participants (developers).

Specifically, developers' awareness, misconceptions, attitudes, and capabilities [20], [23], [24] in S&P can directly affect their development decision. However, only a few studies have examined S&P in XR development at a high level, often focusing on general concerns rather than specific threats. For example, Adams et al. [25] interviewed developers about their general concerns and attitudes about XR S&P. These studies only rely on developers' subjective perceptions and concerns, which limits the identification and exploration of *unknown unknowns*—that is, situations where developers are unaware of, or cannot systematically recall XR-specific threats (i.e., threats introduced by XR interaction design or data collection), and reflect on them within the context of their app development. Given the emergence of diverse threats in XR, dedicated studies are needed to investigate these threats, examine how developers perceive them, and assess how these threats affect XR development practices. The absence of a threat-focused, developer-centered study makes it challenging to understand: (1) the relationship between XR development and the presence of different threats in XR applications, (2) the challenges developers face in implementing effective mitigation strategies, and (3) the ways to provide actionable solutions to support developers in addressing these threats.

To address this research gap, our work takes an initial step toward providing an in-depth understanding of how developers perceive emerging threats and existing mitigation strategies in XR, aiming to fill the missing but necessary knowledge to handle both current and emerging threats. More formally, we have two research questions:

RQ1: What are developers' perceptions of emerging S&P threats in XR?

RQ2: What are the developers' perceptions of current mitigation practices and the support from the XR community?

To address the research questions, we conducted semi-structured interviews with 23 professional XR developers. To investigate developers' perspectives on threats unique to or amplified by XR, we curated and presented a diverse set of sensitive data sources and attack scenarios (e.g., the attacks illustrated in Figure 1). These examples, sourced from top-tier conference publications, represent a broad and emerging threat landscape in XR and enabled us to elicit developer feedback. We then asked developers to reflect on the necessary mitigations for these threats and the responsibilities various stakeholders in the XR community should undertake. Following the interviews, we conducted a qualitative analysis using a bottom-up, open-coding approach to identify developers' perceptions of different S&P threats and mitigation.

Our interviews reveal that developers recognize the importance of XR S&P but are hindered by awareness gaps and unmet supports from the XR community. From developers' perspectives on the presented threats (Section IV, V), we identified factors that make it difficult for XR developers to recognize potential S&P threats. They include suboptimal practices caused by awareness gaps, sensitive use cases (e.g., the collection of rich sensor data from the user's viewpoint), and the tension between S&P and the immersive nature of XR environments (e.g., increased potential for user misuse). Furthermore, in Section VI, we found gaps in developers' awareness of existing tools for threat mitigation, suggesting that developers may still fail to take effective actions to protect their users from known threats. We also observed that many community solutions either reduce utility (e.g., blocking raw camera feeds), rely on outdated documentation (e.g., Meta's API docs), or require external S&P experts—impractical for XR developers, especially those at small companies with limited budgets. These external issues could exacerbate developers' challenges for mitigation in XR development, awareness gaps and suboptimal practices.

Our findings suggest that despite the growing number of S&P threats revealed in academic research and threats occur in real-world applications, XR developers face multiple challenges to keep up-to-date knowledge and promptly address them in their actual development process. In Section VII, we synthesize two core problems, *awareness gaps* and *diffusion of responsibility*, and discuss how the emerging nature of XR interactions and the user-experience-driven nature of XR threats contribute to these problems. Following these insights, we propose stakeholder-aware recommendations to enhance S&P in XR.

We summarize our contributions as follows:

- To the best of our knowledge, we are the first to conduct a developer-centered study exploring developers' perceptions (e.g., awareness, misconceptions, attitudes, and capabilities) regarding diverse emerging threats in XR.
- We conducted 90-minute interviews with 23 professional XR developers to identify contexts in which S&P threats are considered important or amplified during XR development (e.g., the amplification of sensitive data through immersive designs, the collection of rich sensor data, unmoderated XR social experiences, etc.).
- We further identified gaps in developers' perceptions and knowledge about existing mitigations and support, and then proposed improvement directions from both developers' and the broader community's perspectives. These findings highlight the urgent need for enhanced support in XR S&P strategies, technologies, and communication.

II. Background and Related Work

A. Security and Privacy in Extended Reality (XR)

XR encompasses emerging technologies that deliver immersive experiences, including Augmented Reality (AR), Virtual Reality (VR), and Mixed Reality (MR). Recent research highlights that the unique characteristics of XR development make it particularly vulnerable to S&P threats [26]. Specifically, XR applications collect an extensive range of fine-grained human data (e.g., motion, voice, and virtual behavior) which often contains more granular and sensitive user information than traditional applications. Breaches in XR also pose more severe risks to user identity [27], [18], [16], surrounding environments [28], [29], and typing information [30], [8]. XR environments also introduce new attack surfaces that developers must consider when developing interactive designs. Researchers have identified XR-specific problems with overlooked design choices and specialized functionalities [8], [7], virtual social experience [31], [32], [33], [34], [35], content security [36], [37], [38], user perception [10], [11], [12], side-channels [39], [9], and insecure implementations [40], [41]. Moreover, there are only a few standards, policies, and guidelines [42], [43] that address S&P in the context of XR development. Although these policies provide suggestions for avoiding basic security issues (e.g., identity, intellectual property) and privacy leakage (e.g., location, financial information, and personal preferences), they remain insufficient. They fail to cover many XR-specific threats [14] and types of sensitive data [26], and they lack an understanding of the relationship between XR development and existing threats. Our work seeks to address this gap by providing a deeper understanding of XR developers, identifying S&P issues in XR development, and offering actionable directions for S&P-aware XR development.

B. Developer-Centered Qualitative Study on S&P

While user-centered studies offer valuable insights into the current situation, they lack an in-depth understanding of how development processes contribute to emerging S&P threats. To this end, research focused on developers' perceptions—such as their awareness, misconceptions, attitudes, and capabilities—is essential, given developers' central role in the software lifecycle. Such developer-centered studies can better reveal the causes of S&P issues and inform the creation of usable, adoptable solutions. For instance, Acar et al. [20] argue that researchers should consider developers in the S&P landscape. Building on this, researchers have examined how developers contribute to S&P issues, including their role in vulnerable implementations [44], [45], opinions on privacy concerns [46], [47], and perceptions of S&P documents [23], [48].

However, there is a lack of developer-centered studies focusing specifically on XR S&P. A better understanding of how threats in XR S&P affect development is needed to identify existing problems and build stronger support around the development process [49], [50]. Such support is currently limited in XR, largely due to emerging threats associated with new interaction designs in XR. Most recent XR S&P-related studies primarily focus on users [21], [22], [51], [52] and experts with notable experience and knowledge in XR [53]. While there is a general developer involved study [25] on VR developers' perspectives on S&P, it does not address XR-specific threats, such as those related to XR interactions and extensive data collection. Abhinaya et al. [51] examined both users' and developers' perceptions of harassment in XR, uncovering several key concerns. However, by focusing exclusively on harassment, it does not compare or analyze other S&P threats, which limits its ability to surface broader development challenges, or offer actionable recommendations for emerging risks from the developers' perspective. A significant gap remains, as none of these studies examine developers' perceptions or responses to emerging XR threats, and the challenges of designing S&P-conscious XR applications. As previously discussed, this gap needs to be addressed to better understand the causes of unique S&P issues in XR. To this end, we summarized XR-specific threats from existing literature and integrated domain expertise to conduct the first developer-centered and threat-aware study, aiming to understand developers' perceptions, the reasons behind those perceptions, and to provide usable solutions.

III. Methods

A. Recruitment

We recruited 23 professional XR developers with experience in developing and distributing XR applications. Participants were recruited by searching on LinkedIn using the keywords of *XR/VR/AR developers*. We assessed their qualifications by reviewing their LinkedIn profiles and personal CVs, and directly contacted those who met our criteria. Participants were required to have published at least one complete XR-specific project on a public distribution platform (e.g., Meta Quest) and to be currently employed as XR developers. Additionally, we confirmed their XR experience by asking about their past projects through an invitation email or message. Once their experience were verified, we proceeded by inviting them for a Zoom interview. Of the 412 qualified developers we contacted (from June 2024 to June 2025), 34 participants signed up, and 23 ultimately participated in the study. The decision to finalize the number at 23 was based on reaching saturation of themes, where no new themes emerged from additional interviews. All 23 participants completed the main study (approximately 90 minutes) and were compensated $70 USD.

Participant Demographics: Among our participants, eight identified as female, fourteen as male, and one preferred not to disclose their gender. The group was relatively young, with eight participants aged 18-25, eleven aged 25-35, three aged 35-45, and one aged 45-55. Despite randomly inviting all qualified participants we found online, most were young males, reflecting the gender and age makeup of the XR developer community, as observed in similar studies [54], [25]. Table I shows the information on our diverse participant pool. Specifically, our participants have experience developing 15 types of XR applications across different domains. Moreover,

our participants have published applications on 12 different platforms. Among all of our participants, only two participants have less than 1 year of experience, six have 1-2 years of experience, eight have 2-5 years of experience, six have 5-10 years of experience, and one has more than 10 years of experience. Notably, two participants have contributed to XR applications with over one million downloads.

B. Study Design

After recruitment, the study included a pre-study survey and an interview phase. The pre-study survey collected participants' demographic information, development experience, and application publishing channels. The survey was completed within three days prior to each scheduled interview.

During the main study session, we conducted semi-structured interviews with each participant to gather information on their development practices, perceptions of the threats we demonstrated, and opinions on the XR industry's security and privacy (S&P). We assured participants of anonymity and emphasized that our study aimed to address issues affecting the entire developer community, not to test individuals.

This study was approved by the Institutional Review Board. At the beginning of each study session, the interviewer briefed the participants on the study goals and procedures and then asked them to sign a consent form. The semi-structured interview contains the following four modules:

Background Questions: The interviewer asked participants about their XR development background, including training, involvement in S&P design, experiences with S&P issues, and perspectives on community practices. To elicit unique insights, the interviewer tailored prompts using participants' pre-study survey responses (e.g., application details). Both prepared and impromptu follow-up questions were used to explore topics such as their role in S&P-related decisions.

Threats in XR: The second part of the interview aimed to understand participants' perceptions of S&P-related threats in XR, including sensitive data, which we extended based on [26] (Table II), data leakage channels (Table V), and potential attacks. As detailed in Table III, we reviewed notable and representative attacks documented in existing XR S&P literature from top venues, classifying them as XR-related (e.g., perception manipulation) or XR-amplified (e.g., social attacks). Building on the definitions proposed in [55] and [14], we organized these threats into categories for clarity and demonstration convenience. Initially, we categorized them based on primary attack purposes outlined in [55], then refined the categorization by incorporating variations in attack surfaces related to XR design choices and use cases.

We then presented these threats (e.g., perception, social attacks in XR) to participants in two different orders (i.e., ascending and descending) to gather their feedback. Due to time constraints, the demonstration included all threat categories but did not exhaustively cover every XR S&P threat identified in the literature. However, as our aim was to explore developers' perceptions, we believe the selected examples were sufficient and broadly representative for XR threats. After the demonstration, the interviewer asked participants to reflect on the threats regarding their practicality and importance. The interviewer followed up with questions to assess their comprehension of the content and ask for potential mitigation.

Mitigation and Best Practices: The third part of the interview focused on understanding participants' awareness of existing mitigation strategies in XR and their perspectives on the responsibilities of various stakeholders (e.g., policymakers, users, platform providers) in the XR community. We compiled tools and practices from existing literature into a slide deck as provided in Table IV, asked developers about their awareness of these mitigations, and presented explanations and examples. We then asked participants to use the tools to address the threats mentioned previously and assess their understanding.

C. Qualitative Analysis

In alignment with our two research questions, we conducted a qualitative analysis of interview transcripts and screen recordings using a bottom-up open coding approach facilitated by MAXQDA. Following the best practices by Saldana [75], our analysis involved two rounds of coding to ensure a thorough examination and identification of key themes.

Three researchers independently analyzed four interviews in the first round of coding to develop an initial codebook through repeated transcript readings and iterative discussions. Daily meetings were held to confirm, refine, and merge codes, resulting in a preliminary codebook comprising 124 codes. Building on this foundation, the researchers engaged in axial coding analysis. This phase involved merging similar codes and organizing them into overarching themes, providing a structured framework to address the research questions.

Following the initial coding process, the task of coding the remaining data was undertaken by three researchers using the agreed-upon codebook. Each data sample was then coded by a primary coder and verified by another researcher. Any new codes emerging from this coding process were discussed among all three researchers and incorporated into the codebook upon reaching consensus. Following the recommendations of McDonald et al. [76], inter-rater reliability was not calculated for the coding process as the primary objective was to identify emergent themes rather than to seek coder agreement. This approach allowed the researchers to remain open to novel insights and patterns within the data.

The final codebook consists of 81 codes grouped into 18 themes. The complete codebook is provided in Appendix C.

Labeling Criteria: During our study, we analyzed the ratings of user responses to understand the relationship between developers' awareness and their threat perceptions. The primary labeling was performed by an author with expertise in XR S&P. To mitigate potential bias and ensure consistency, each response was independently coded by a second researcher, selected from a pool of three co-authors on the paper. The second coder then discussed the labeling decisions with the XR expert to reach an agreement. This approach reduced the

TABLE I: **Participant Overview** – Years of experience in XR development, number of XR applications developed, the categories and styles of developed applications, and the platforms where these applications were published and accessed.

ID	XR Development YOE	# of Apps	App Categories	App Style	Publish Channel
P1	0-1 Years	1	Single player	Education	Meta
P2	2-5 Years	2	Single player	Design, Action, Puzzle, Fitness	Meta
P3	1-2 Years	4	Single & Multi-player	Design, Education	itch.io
P4	2-5 Years	12	Single & Multi-player	Social, Action, Role-playing	Apple, Meta
P5	1-2 Years	4	Single & Multi-player	Design, Action, Role-playing, Puzzle	Steam, itch.io
P6	1-2 Years	3	Single player	Retail, Sport	Snap
P7	5-10 Years	2	Single & Multi-player	Social, Music	Steam, Meta
P8	5-10 Years	5	Single & Multi-player	Social, Event	WebXR
P9	5-10 Years	7	Single & Multi-player	Education, Healthcare, Action, Sport	Meta, Steam, Pico, WebXR
P10	>10 Years	15	Single & Multi-player	Social, Education, Retail, Design, Action, Sport	Meta, Apple
p11	2-5 Years	4	Single & Multi-player	Design, Action, Puzzle	Spark AR
P12	2-5 Years	5	Single player	Design, Action, Puzzle	Meta, Apple, Google Play
P13	5-10 Years	>40	Single & Multi-player	Social, Education, Healthcare, Retail	Meta, Directly to business
P14	1-2 Years	3	Single & Multi-player	Social, Design, Action	Meta
P15	2-5 Years	8	Single & Multi-player	Education, Healthcare, Retail, Action	Steam, itch.io
P16	2-5 Years	2	Single player	Education, Healthcare, Relax	Meta, Steam, Pico, itch.io
P17	0-1 Years	8	Single &Multi-player	Social, Museum, Other	itch.io
P18	5-10 Years	29	Single & Multi-player	Social, Education, Healthcare, Design, Puzzle	Steam, Apple, WebXR
P19	5-10 Years	2	Single player	Education, Design, Social	Meta, Apple
P20	1-2 Years	3	Single player	Retail, Design	Meta, Apple
P21	2-5 Years	6	Single player	Education, Social AR Filter, Other	Snap, YouTube, Instagram, Meta
P22	1-2 Years	3	Single & Multi-player	Education, Social	WebXR
P23	2-5 Years	20	Single player	Education, Design, Action	Meta

TABLE II: Sensitive data in XR we collected and summarized from the literature [26], [56].

No.	Data Type	Sensitive Information Examples
1	Motion-related data (e.g., motion tracking/gesture) [16], [18], [57]	Identity, age, gender, physical state, biometric information, etc.
2	Voice data (e.g., microphone data) [30], [28], [58]	Speech content, gender, biometric, etc.
3	Camera data (e.g., surrounding environment) [26], [59]	Environment, personal information, identity, other person identity, etc.
4	Device related data (e.g., network/HMD) [39], [26], [40]	Wealth, network, IP address, etc.
5	Spatial data (e.g., location/play area) [26]	Wealth, identity, address, etc.
6	Health related data (e.g., heart rate) [56], [60], [61]	Cognition, emotion, health, etc.
7	Biometric (e.g., eye tracking/gait) [62], [63], [64]	Attention, identity, authentication related information, etc.
8	Observation (e.g., avatar/behavior) [8], [7]	Age, gender, preference, health, emotion, etc.
9	Feedback (e.g., haptics) [65], [56]	Virtual behavior, users state, etc.

risk of prejudice toward individual participants and improved the reliability of the labeling process. The interrater reliability scores are as follows: (1) awareness of attacks, we obtained a Cohen's Kappa of 0.812; (2) for awareness of mitigation, the Cohen's Kappa is 0.920; (3) and for mitigation quality, the quadratic weighted Kappa is 0.827. These scores indicate a high level of consistency among coders [77].

- Awareness of attacks: We classified participants as "Aware" of an attack if they mentioned it without our prompt, stated that they had considered such an attack before, or provided additional examples of similar attacks.
- Awareness of mitigations: Participants were marked "Unaware" of a defense mechanism if they mentioned being unaware of it or misused it for solving demonstrated attacks.
- Quality of developer-proposed mitigation: The researcher rated the mitigation strategies proposed by developers on a *scale of three*, where *one* indicates strategies that do not make much sense, *two* indicates strategies that make some sense but are missing important details, and *three* indicates strategies that are well-thought-out or align with realistic solutions adopted in the industry.

Threat Model: Our threat model focuses on (1) external adversaries (e.g., network/system intruders, malicious XR users, bystanders, or third-party API providers capable of generating harmful content) and (2) careless developers who may inadvertently introduce attack vectors or harmful content to XR applications. We do not consider malicious developers as we assume study participants do not have malicious intent.

During the presentation of each attack, we further discuss the specific attack channels and attacker capabilities. For example: (1) Many XR platforms and applications (e.g., VRChat) allow user-generated content (e.g., custom rooms or game objects), which—if not properly validated—can enable content, perception, or physiology attacks. (2) Malicious third-party API providers can introduce input, content, or software side-channel attacks when their services are integrated without proper vetting. (3) External adversaries may exploit hardware vulnerabilities to manipulate sensor inputs or engage in ma-

TABLE III: Security and privacy attacks in XR we summarized from the literature.

No.	Attack Category	Attack Targets	Explanation
1	Shoulder Surfing Attacks	Spy XR user [66], Spy on bystanders [67], etc.	Recording without other people's consent
2	Software Side-channel Attacks	Infer keystroke [8], [7], re-identify user [18], etc.	Exploiting functionality data as side-channel
3	Input Attacks	Compromise input integrity [68], [69], [14], [70], DoS [71], [72], etc.	Manipulating inputs to trigger dangerous operation
4	Social Attacks	Harass another user [31], [32], [34], [35], social engineering [33], etc.	Utilizing virtual social experience as attack channels
5	Content Attacks	Overlay malicious content [36], [37], impersonate others [38], etc.	Introducing malicious virtual content in XR
6	Perception Attacks	Manipulate perception [11], [10], [12], manipulate memory [73] etc.	Manipulating perception and memory
7	Physiology Attacks	Introduce disorientation [11], introduce content dizziness [74], etc.	Introducing motion sickness and discomfort

licious activities within social XR applications, leading to shoulder-surfing, input-based, or social attacks.

IV. Developers' Perceptions on Privacy in XR

XR technologies typically involve a wide range of user data, which introduces novel privacy threats. This section explores developers' perceptions of these threats, particularly those arising from the unique characteristics of XR user data. In our pre-study survey, developers rated privacy in XR as highly critical, with a median Likert score of 6/7; only three developers rated it 4 (moderately critical) or lower. With this in mind, we further investigate our research question on developers' perceptions of privacy in XR from three perspectives, following the demonstration described in Section III-B:

- Developers' awareness of sensitive data in XR
- Developers' perceptions of sensitive data in XR on:
  - What makes data overall more sensitive?
  - What makes certain data perceived as less sensitive?
- Developers' perceptions of leakage channels in XR on:
  - What makes leakage channels overall more realistic?
  - What makes leakage channels perceived as less realistic?

Additionally, we also analyze potential developer misconceptions at the end of this section. Understanding these perceptions is essential, as their awareness, misunderstanding, and capability directly influence their practices, application designs, and protection mechanisms in XR [20], [23].

A. Developers have limited awareness of sensitive data in XR

When asked about the sensitive data types used or collected in their applications——prior to being introduced to the definitions of sensitive data types and sensitive information in XR (Table II)——developers, on average, identified only *2.1* types of XR-sensitive data type (e.g., motion, biometric) and *0.5* types of non-XR data type (e.g., email, password) as being collected in their applications. This initial finding raises concerns about whether developers' applications indeed collect only this limited set of data types or if developers fail to recognize certain types of data as sensitive.

Developers come to realize that their applications collect more types of sensitive data than they had previously thought. After being presented with XR-sensitive data definition, developers identified significantly more XR-sensitive data types, recognizing an average of *4.8* types collected in their applications and acknowledging protection needs. P3 reflected: *"It makes you realize that some data that seems pretty harmless is actually a lot more potent,"* highlighting potential risks from developers' lack of awareness.

B. What makes data overall more sensitive in XR

After being presented with various types of sensitive data (Table II) and their implications, developers were asked to rate the sensitivity of each data type. From their ratings, we found a median Likert score of 5.0/7.0. We analyzed their reasoning and identified the key factors as follows.

Immersive interaction reduces the awareness of data leakage. P4, P12 and P23 mentioned that the immersive interaction design in XR makes voice data more prone to leakage and increases its sensitivity: *"When you talk to people in VR, it brings back the same feeling that you're talking to people in real life, so often people don't have the expectation that they are being recorded."* Additionally, P7 and P8 raised concerns about unintentional self-disclosure in XR applications due to non-intuitive interaction design. Informed by their prior experiences in XR, P8 noted that this could increase the sensitivity of voice and expressed an intention to incorporate clearer notification mechanisms in his future applications as mitigation: *"I did not realize that they could hear me. There was no point at which I was told or made aware of, or had visual feedback on the fact that I had an open microphone."*

Advanced tracking sensors produce sensitive data with greater details. P7, P8, P13-16, and P19-22 highlighted that data collected in XR contain rich sensitive information, which allows adversaries to infer user information. P13, a healthcare XR developer, emphasized that XR technology amplifies the sensitivity of data leakage in healthcare applications: *"Because this [XR] device involves a lot of activity and emotion expressing capability. Using the controller like with the phone, you just use the vibration or some haptic features to identify things, but XR can track my hand motion, and in future it can comes with some sensor for your legs, or even if not, the sensors available in the excess can still detect your leg movement and the energy and the pace you move."* Similarly, P13 highlighted that XR camera data is more sensitive than data from other devices as it provides more detailed information and is physically attached to the users' direct view. Considering the sensitivity of the camera data, P8 mentioned that they would

only collect this data for valid reasons and properly protect (e.g., via encryption) in development.

The value of XR data incentivizes extensive collection and surveillance. P7, P13, P14, and P19-21 further noted that the economic value of the sensitive information inferred from the data types collected in XR can drive XR headsets, applications, and service providers to collect user data (e.g., camera and spatial data.), significantly increasing the privacy sensitivity of these data. P13 noted that scanning a room for guardian system with an XR headset can leak data to the platform:*"Where is my room? Where is my couch? Where is my table? Where is my TV? So this data definitely has a business value. Because, let's say most of the ADs now understand what we talk and what I'm interested in, and based on that, they will publish those ADs, right? ... at the same time, certain people are not comfortable with sharing that information."* These concerns align with findings in [40], [78], which show that data is already being collected for profiling, analytics, and personal advertisements in XR. More critically, P7, P8, P16, and P21 expressed concerns that headset providers or governments could exploit camera and biometric data for surveillance purposes since these data provide detailed information about XR users and their private space (e.g., home).

Opaque XR infrastructures impede effective mitigation. P4, P9, and P13 highlighted that XR devices are increasingly being used for medical purposes, but the black-box XR systems present challenges related to data protection and compliance, thereby increasing the sensitivity of medical data. P9 mentioned the challenges he faced when shipping his applications with devices as a medical service, particularly in meeting compliance requirements from government agencies (e.g., VA) which require full visualization on collected data in devices [79]: *"The problems we have with the headsets are that it is bytedance, or it's Meta, HTC. There's no open source. But on the same day, you know, it's like we're following the API callback to mainland China, ... We're part of the VA, the US government doesn't want that."*

C. What makes certain data less sensitive

While developers generally perceive data as more sensitive in XR, they still rate some data types as less sensitive (score $<$ 4). we analyzed their reasoning and identified the key factors:

The perceived benefits tend to overshadow the risks of privacy leakage. P4, P7, and P17 described a trade-off, classifying functionally essential XR data types as less sensitive and prioritizing technical performance over privacy concerns. P4 stated: *"I think it's sensitive, but it's required for the application to work. It's hard to make a gesture-based game if you can't get the gesture."* Similarly, the importance of motion data for enabling immersive experiences contributes to P7's perception of it as less sensitive: *"If you want to exist in the same place as somebody else. You're gonna need to know what their hands are doing with their bodies, and people voluntarily add extra trackers to their bodies to have the avatar be more realistic."* This contradicts the finding in [16], where the connection between motion data and crucial XR operations increases the potential to infer important information.

The insufficient knowledge of developers tends to obscure the consequences of privacy leakage. P1, P2, P10, P13, P15, P16, P18, P20, and P21 stated that they do not consider certain data types (e.g., motion, feedback) to be highly sensitive even after our demonstration, as they could not identify plausible risks of sensitive information leakage based on their experience. For example, P7 mentioned: *"I don't think it's super sensitive. I think most kinds of data that you're gonna get from that, it's like, Oh, this person whacked their controller into a wall. And so there's a wall there."* Moreover, P4, P10, P13, and P16 admitted that they were not aware and had not previously considered the potential consequences of using or collecting feedback data (e.g., haptics). Therefore, they perceive feedback data as less sensitive as P13 said: *"So this is a very niche and new thing, which we don't know, because the applicant is completely new, and people are still figuring out how to use. It's hard for me to say how much threat from my experience."*

D. What makes leakage channels more realistic in XR

After presenting potential sensitive data types and possible data leakage channels, we collected developer feedback on which channels are more likely in XR. We identified high realism ratings with a median Likert score of 6.0/7.0. Analysis of their reasoning identified the key factors as follows:

Prevalence of development misoperations due to technological immaturity. P4, P6, P7, P9, P10, P12, P13, P16, and P18-22 reported, based on their own experiences or those of XR developers they worked with, developers tend to blindly rely on third-party packages (e.g., Photon [80]) when developing XR applications. For example, P12 noted: *"It's a fact that people use APIs all the time without really knowing what's under it ... If you're using Unity to build anything, you're gonna use a bunch of packages."* P7, P9, P10, P12, P13, P15, and P16 also noted that insecure operations, as a data leakage channel, are likely to occur in XR. P13 believes the immaturity of XR technology and market makes this trend especially concerning: *"When it comes to the AR/VR set of things. It's still a maturing technology where a lot of freelancers, small agencies, people from all sizes of teams, and backgrounds are coming in. So, as I mentioned, no matter what the protocols we bring in, at the end of the day. It's a responsibility of the developing team."* These threats stem from the current landscape of the XR industry, which is largely composed of individual developers and small enterprises [81]. This context suggests that developers should be careful when using third-party packages to avoid potential data breach.

Ease of self-disclosure by users in an immersive environment P1, P4, P5, P7, P9, P13, P14, P17 and P21 expressed concerns that beyond intentional information sharing, XR's immersive nature can more easily lead to the unintentional disclosure of private information (e.g., behavior patterns, voice). For example, P17 noted that users may accidentally leak private data and emphasized that developers should take

responsibility for helping to protect them: *"I am a firm believer that users are stupid. So if you, if you have something, chances are they will screw it up. Doesn't mean it's intended. But there's always something right?"* P14 emphasized that such misuse is more likely to occur in XR due to its immersive interaction, which tends to encourage self-disclosure, aligning with the findings in [82], [83]: *"In VRChat. I think a lot of times we are interacting with friends in a very natural way. I wouldn't pay extra attention to how I behave. So that data could be very suggestive of who I really am."*

E. What makes certain leakage channels less realistic in XR

While developers generally agree with the presented leakage channels as realistic, we identified cases where they give a lower score (score < 4). To better understand the reasoning behind these ratings, we classified developers' perception of leakage channels being less realistic due to:

The responsibility is not acknowledged in XR development. P3, P6, P7, P9, P10, P13, P15, P16, and P21-23 mentioned that they were unfamiliar with hardware-side channels for data leakage. For example, P3 stated: *"I am not too knowledgeable about this one, so I'm giving a lower score."* Additionally, P16 believed that developers should not be responsible for it: *"I think that is possible. But I'm not really concerned about that. And also as a developer, I can't really do much about that."* Similarly, P7, P15 and P17 suggested that this issue should be addressed by XR headset providers and governments.

The belief that local XR applications are inherently secure. When asked about the necessity of implementing protections (e.g., encryption or secure data storage) to prevent XR data leakage in their applications for the types of data demonstrated in the study, P1-3, P6, P7, P10-12, P16 and P17 believed that data handled by local XR applications (i.e., those without network functionality or limited to single-player use) would be inherently secure, and therefore would not require additional security measures. For example, P11 commented: *"I think it's just fine because it's a local APP. So it only processes those data and uses it, but doesn't actually send it to anything."*

F. Analysis: Potential misconceptions on privacy in XR

Based on the developers' perceptions above, we analyze and summarize the following misunderstandings about XR privacy:

Misconceptions regarding data sensitivity. From Section IV-C, we observed two common misconceptions: (1) the belief that the data with higher utility is less sensitive, and (2) the assumption that unfamiliar data is also less sensitive. However, according to the GDPR [84] and widely accepted research guidelines [85], [86], [87], any data that can be used to reveal an individual's identity or personal information (e.g., financial or health data) should be classified as sensitive and protected using appropriate mechanisms (e.g., [27]).

Misconception on the data leakage channels. FromSection IV-E, we observed the following misconceptions: (1) Misconceptions of assuming that mitigating data-leakage channels (e.g., hardware side channels) is outside their scope. For example, developers may believe this responsibility lies with hardware teams or other stakeholders. However, prior work emphasizes that software developers should play an active role in preventing leakage, even when it originates in hardware. Developers are encouraged to participate in co-design efforts and implement software-level protections such as isolation [88], [89], [90]. (2) The misconception that local XR applications are inherently secure. While local applications may offer more security than online apps, they remain vulnerable to leakages. For example, through unsafe Android API calls, side channels (e.g., rendering and motion), and insecure local storage [41].

V. Developers' Perceptions on Security in XR

In addition to privacy threats, In this section, we explore developers' perceptions and understanding of threats happening in XR. Our pre-study survey reveals that developers generally perceive security in XR as a critical issue, with a median Likert score of 5/7. Among the 23 developers surveyed, only four rated XR security $<= 4$ (Moderately Critical). Building on these findings, we further investigate our research question on developers' perceptions of security in XR from three aspects, following the demonstration in Section III-B:

- Developers' awareness of security attacks in XR
- Developers' perceptions of attack importance regarding:
  - What makes attacks in XR overall more important?
  - What makes certain attacks perceived as less important?
- Developers' perceptions of attack practicability:
  - What makes attacks in XR overall more practical?
  - What makes certain attacks perceived as less practical?

Similar to the previous section, we analyze these responses and highlight potential misconceptions held by developers. We focus on these questions because developers' awareness, misconceptions, and capabilities directly influence the security and mitigation they adopt in XR applications [20], [23].

A. Developers exhibited limited awareness of XR attacks.

Without our demonstration of the categorized XR attacks, developers recalled an average of only *0.9/7.0* XR-specific attack types listed in Table III. Furthermore, 6/23 developers reported being unaware of any XR-specific attacks. These findings suggest that, without reminders, developers may struggle to adequately consider potential XR-specific threats during the design and development of their applications. To better understand this gap, we further investigate whether the lack of awareness stems from developers never having encountered these threats or simply needing a prompt to recall them.

Developers had never heard of many attacks prior to our demonstration. Based on developers' feedback and our analysis (as described in Section III-C), developers had not heard of *3.7/7.0* attack categories prior to our demonstration. Moreover, only 10/23 developers were introduced to more than half of our demonstrated attacks. Nevertheless, after the demonstration, developers expressed that they gained new insights into important attacks they did not consider but would be helpful for their future development. As P9 remarked

regarding the software side-channel attacks used to infer keystrokes [8], *"I've built keyboards in VR, and I will think about the possibility of these attacks now."*

Developers' attack awareness undermines their perceptions of attack importance and practicability. Developers assigned lower scores to attacks that they did not know before. To investigate how awareness influences perceptions, we categorized ratings by attack awareness and applied Mann-Whitney U-tests to assess significant differences. Our analysis revealed a significant difference ($U = 4245.5$, $z = -3.458$, $p < 0.001$) between developers' ratings and awareness of attacks. Developers rated known attacks higher median importance score of 7.0/7.0 versus 6.0/7.0 for attacks they were unaware of before. Similar to attack importance, for attack perceived practicability, developers gave significantly higher ratings to the attacks they were aware of ($U = 5032.5$, $z = -6.125$, $p < .001$), with a median score of 7.0/7.0 compared to 5.0/7.0 for those they were unaware of before.

B. What makes attacks overall more important in XR

After being presented with various attack types (Table III) and their implications, developers rated attack importance with a median score of 6.0/7.0. We analyze their reasoning and identify the key factors as follows:

Immersive experiences are considered highly personal. P2, P8-10, P12, P14, and P16-22 noted that attacks like social attacks are particularly important because XR experiences are highly personal. P12 stated: *"It's a completely different feeling when people are in your personal space, you're using voice and you can hear them, you can talk to them, you can see their hands. It's very different than a social media app that you can ignore on a screen ... I think the damage can definitely be a lot more than other media, like I had a kid that came up to me and just sneezed in my face, and I felt so disgusted, I know, like it's virtual, and it was like: dude. Get away from me."* P12, P14, and P22 emphasized that social attacks, from these previous negative experiences, can be highly disruptive and should be mitigated in both their application and other multiplayer XR applications. As P22 noted: *"Especially social XR platforms like VRChat, at least consider it as part of the game's design."*

XR functionalities may negatively impact the large user population. P6, P10, P19, P22 and P23 identified attacks as important due to potential large-scale sensitive information leakage from the community's perspective. As P6 stated about the shoulder-surfing attack using XR devices: *"This sounds like a very big issue. Because you know AR glasses are coming soon right? By a bunch of companies, or you know what, a lot of them have already been around. I think mass adoption is going to be coming within the next couple of years."* Similarly, without hesitation, P4, P6, P7, P9-11, P16, and P21 believed social attacks are the most important attacks in XR because they affect a large number of users. To avoid similar threats happening in their applications, developers stated that they would actively consider more of these attacks during implementation. For example, P20 mentioned their efforts to reduce the risk of social, physiological, and perceptual attacks:*"That's why we're always trying to consider that in our design. And we're gonna do a lot like testings internally before we bring it to test down users."* Moreover, P8, P15, P16 and P18 mentioned that developers are worried that negative experiences caused by these attacks could drive users away.

XR applications are beginning to carry out higher-risk activities. P3, P4, P7, P8, P13, P14, P16, and P21-23 indicated that attacks like side-channel attacks could be used to infer sensitive user information during higher-risk activities in XR, such as password entry, private behaviors, conversations, or other personal details, making these attacks especially important to consider in XR development. P16 commented: *"Those can track a lot more physical data, and also your health data, and even the like. Your motions, signatures, or even capture, if you have certain injuries."* Similarly, P13 mentioned that this threat will become more severe in the near future as higher-risk operations, such as banking, are introduced in XR. Developers and the XR community should prioritize their focus on protecting data used in these higher-risk applications.

The immersive nature of XR makes mitigation more challenging. P6, P16, P19, and P22 believe that attacks are more important in XR as they are harder to mitigate for developers, which aligns with [91]. P6 used social attacks as an example: *"You probably aren't able to do real-time voice detection."* Similarly, P16 commented: *"I think it is very important ... I also think it's difficult to moderate, most things like Rec Room, they're unmoderated experiences."*

C. What makes certain attacks less important in XR

While developers generally perceive attacks as important in XR, they consider some attacks less important (score < 4). We analyzed their reasoning and the key factors as follows:

Users are capable of performing mitigation. As highlighted by P3, P7-9, and P14, physiology attacks can be perceived as less important since users can simply remove headsets or close their eyes for mitigation instead of implementing mitigation during the development process. For example, P8 mentioned:*" I think it's too easy to just take off the headset. It's not like I'm holding you hostage or anything."* Similarly, P7 mentioned that closing eyes can also prevent these attacks. However, this shifts the burden of mitigation onto users, which contradicts recommendations for preventing physiology attacks [71], [92].

Certain attacks are perceived as applicable only within specific XR environments. P4, P6, P7, P10, P16, P17, P20, and P21 mentioned that they consider certain attacks less important when they believe those attacks are only effective under certain circumstances. For example, P4 commented on physiology attacks: *"I think the damage is more or less minimal because it's impacting an individual user ... in most cases it's more so harmless and just a Disney effect."*

D. What makes attacks overall more practical in XR

Similar to the process of evaluating attack importance, we asked developers how practical the attacks were and identified

a median Likert score of 6.0/7.0. We analyzed their reasoning and identified the key factors as follows:

XR provides stealthiness and flexibility through hardware and immersive technology. P6, P8-10, P13, P14, P16, P19, P21 and P22 perceived attacks (e.g., shoulder-surfing attacks) as practical because of their stealthiness to perform attacks when comparing XR to other mobile devices. P13 mentioned: *"If you're using a mobile or a camera. People at least will get an awareness, right? Somebody is trying to steal something from me. Different cameras again, it's hard. But this one is like an easy option for a threat, and hard to identify, for the person who is being watched."* Moreover, P6, P8-10 and P16 mentioned that social attacks are more flexible in XR, making them more practical from a technical perspective. Based on these experiences, P16 and P19 emphasized that they want to contribute to mitigation for better user education and improved notification systems by developers and platforms.

Encouraging user-generated content in XR may enable adversaries. P4, P9, P16 and P22 noted that since XR environments encourage users to create their own content (e.g., game rooms), the likelihood of physiology attacks increases. P4 noted: *"I think in VRChat, you can throw a flash bomb to introduce seizure...We rely on the users to like deal with this. Platform allows these kinds of behaviors."* P4 was also concerned with content attack introduced by user-created content: *"That's actually one of the things we tried to prevent when we're doing an event on Horizon, because the user can just jump on stage or show models or images that's not great."* Considering this issue, P4, P16, and P22 have mentioned the need in their applications for better access control and moderation tools to mitigate malicious user-generated content.

E. What makes certain attacks less practical in XR

While developers mostly agreed attacks were practical, some gave lower scores (< 4). We analyzed their reasoning and identified the key factors as follows:

Execution of attacks requires significant technical expertise. P10, P12, P15, P17, and P23 mentioned that attacks, such as perception and input attacks, are technically difficult in XR, making them less practical. For example, P15 highlights perception attacks: *"You need a lot of technical know-how to be able to manipulate boundaries like that."* Similarly, P4, P7, P9, P15, and P16 noted that conducting input attacks requires technical expertise in XR or an understanding of users' states.

User-experience oriented attacks are perceived as less motivating. P6, P8, P13, P14, P16, P17, P21, and P22 mentioned that attacks (e.g., perception attacks), seem less motivating for attackers, as they primarily target user experiences rather than direct benefits, making them impractical. As P6 commented: *"You can't really gain anything from this. Aside from harming people, basically like you can't earn money from it right?"*

F. Analysis: Potential misconceptions on security in XR

Based on the developers' perceptions above, we analyze and summarize the following misconceptions about XR security:

Misconceptions on attack importance. From Section V-C, we identified the following misconceptions: (1) considered that attacks that can be mitigate by users as less important in development. While users can sometimes mitigate attacks (e.g., removing a headset during a physiology attack), developers should not shift the responsibility to users. Instead, they should implement robust defenses, especially when users may be unaware of potential threats. This view is supported by [93], [94]. (2) Misconception that the attacks only affect specific groups (e.g., physiology attacks on individuals with seizures) are less important. Developers should not overlook these risks. All users should be considered equally in security design, especially when user safety is at stake [95].

Misconceptions on attack practicability. From Section V-E, we identified the following misconceptions: (1) Attacks with technical complexity are perceived as less practical. A more rigorous approach to assessing attack feasibility involves using a threat model that assumes attackers possess the necessary skills and have reasonable access to relevant information [96]. (2) Developers may underestimate attacks due to a lack of appreciation for attacker motivations. Specifically, they may overlook user experience–oriented attacks, assuming they are less likely. However, as XR technology is increasingly deployed in critical domains such as the military [97] and healthcare [98], these types of attacks remain both feasible and attractive to adversaries. Moreover, all user deserves full protection, as emphasized in ethical development guidelines [99].

VI. Developers' perceptions on current mitigation practices and supports on XR S&P

After examining developers' perceptions of S&P threats in XR, we aim to understand whether they are capable of addressing these threats with appropriate mitigation strategies. Based on developers' responses to threats and our questions, we address our second research question by summarizing the findings that answer: (1) What are the limitations of the current mitigation practices for XR S&P threats? (2) What are the missing supports in current mitigation practices for XR S&P threats? Specifically, we present insights on these questions from the perspective of both the developers and the broader XR community (i.e., other stakeholders). This is important because we aim to identify the limitations of current mitigation practices and address existing problems. Moreover, we seek to understand how other stakeholders in community can better support developers in ensuring S&P.

A. Limitations of current practices

As a first step, we identified limitations in proposing mitigations, awareness of existing mitigation practices, and perceptions of those practices provided in XR communities.

Developers lack the ability to propose effective mitigation strategies. After demonstrating each attack, we asked developers to propose a mitigation strategy and evaluated their responses using the criteria described in Section III-C. Developers sometimes provided impractical or unclear mitigation

strategies. Additionally, developers' awareness of attacks influenced their ability to propose effective mitigation strategies. We conducted a Mann-Whitney U-test to compare mitigation scores between the two groups. Developers provided significantly better mitigation for attacks they were aware of, rating familiar attacks higher ($U = 3956.5$, $z = -2.48$, $p < .01$).

Many developers are unaware of existing policies and standards that inform best practices for privacy in XR. 12/23 developers were unaware of the GDPR [84], and 14/23 were unaware of the CCPA [100], both of which are relevant to S&P in XR applications. Moreover, none were aware of XR-specific standards such as Rosenberg et al. [43].

Developers are unaware of existing mitigation practices for XR. When presented with mitigation tools for S&P in XR development, developers noted that they were unaware of some of these tools but believed these tools could be helpful in addressing S&P challenges in XR. For example, 14/23 developers reported unfamiliarity with program analysis tools (e.g., Ghidra [101] and Roslyn [102]) used for detecting vulnerabilities, and 14/23 were unfamiliar with existing system- or game-engine-level protection methods (e.g., Arya [103]).

Existing mitigation sacrifices utility in exchange for S&P. P3-6, P12-14, and P21 expressed concerns that existing privacy mitigation practices (e.g., restricting access to camera feeds) constrain developers' creativity and limit the potential of XR applications. P6 noted: *"They don't provide any access to the camera and information at all, It really limits what you can develop as a developer."*

Platforms need to improve the S&P checking process for application publication. Developers highlighted that all current platforms publishing S&P policies need improvement. In our study, 14/23 participants had published applications on the Meta Quest platforms. Among them, four reported that their submissions did not undergo a detailed review of S&P, and five recommended that Meta enhance its review process by implementing more rigorous content checks and providing clearer feedback, guidelines, and documentation. Compared to Meta, fewer developers in our study (6/23) had published on the Apple platform. However, the satisfaction rate was slightly higher: only two of these six developers believed that Apple could benefit from stricter reviews or more developer guidance, while the remaining four considered Apple's policies to be adequate. Additionally, P11, P15, P19, and P22 observed that S&P requirements are largely absent on other distribution platforms (e.g., SideQuest, itch.io, and AR platforms).

Disclosure and feedback on S&P problems is not sufficient. P1, P4-10, P13-16, P19, and P22 emphasized the importance of timely disclosure of S&P issues to ensure effective mitigation, but noted that such disclosure is largely lacking at present. To address this problem, developers called for the establishment of more testing and disclosure channels—both formal (e.g., beta testing, bug bounty programs) and informal (e.g., app reviews, user comments)—to help identify implementation and design flaws related to S&P issues. For example, P15 noted that collaboration and disclosure among developers was invaluable in his previous XR development projects at larger companies. *"Have developers that are all working on like VR or XR in general, when you publish it, and they all test it, and tell you what's good, what's not, and then you fix that... We need dedicated developer-focused forums, similar to Reddit, to gather feedback from experienced developers. This feedback can help you refine your application and enhance its privacy compliance."* Moreover, P13 called for additional tests from users for S&P issues in their applications and help uncover previously unknown problems.*"I would like users to be open about the issues they are facing. However, I haven't gotten any feedback provided by my end users regarding the protection mechanism used in the application.*"

Regulations and best practices for XR are impractical or outdated. P1-11, P13-16 and P21, P22 expressed frustration with existing guidelines and documentation in XR by describing them as consistently difficult to follow. They characterized these resources as overly lengthy, frequently updated, and occasionally inaccurate which is largely attributed to the current stage of the XR industry. P16 also noted that this issue stems from the suboptimal practices of industry-leading platforms: *"Meta is the biggest vendor out there, and they have so much old documentation mixed with new documentation. Oftentimes it is impossible to know what is current, and sometimes even the current stuff is not well documented."*

B. Absence of supports from the XR community

In addition to the limitations of existing solutions, developers' suggestions for mitigating threats and other suboptimal practices in XR (e.g., outdated documentations) revealed the lack of support within the XR industry. We categorized these needs into three groups demonstrated below:

1) Strategic Support: From the developers' response, we identify a need for strategic support (i.e., the provision of guidance or resources) as follows.

Better support for developers to address economic and utility concerns in S&P-aware XR development. P2, P6-10, P12, P13, P15-19, and P21 highlighted that the tension between utility and security is fundamentally rooted in the economic constraints of XR development. P7 remarked:"*Security is definitely something we want to consider sooner rather than later, but it's essential to have more cost-effective solutions.*" Moreover, as previously discussed, developers called for strategies different from existing ones (e.g., blocking the camera feed) that do not compromise utility. P2 and P11 also emphasized the need for strategic support to balance the economic needs of smaller companies with S&P in XR:*"[Companies/Policymakers] should reward developers that demonstrate greater responsibility toward the technology, to encourage them to build a compliant application."* P5, P7-9, P16, and P22 also emphasized the need for community to provide external S&P experts, which developers can not afford.

Better guidance for best S&P practices in XR. After demonstrating the threats, P1-3, P5-13, P15, P16, and P22 emphasized the need for guidance and professional resources

to address emerging attack surfaces in XR. For example, P5 and P16 highlighted the responsibility of industry leaders (e.g., Meta, Apple) to provide accessible resources such as checklists, templates, tutorials, and consulting services to support developers, particularly junior developers, in S&P practices such as privacy policy drafting. Moreover, developers found that general S&P policies (such as GDPR and CCPA) were often tedious, costly to implement. For example, P8 mentioned:*"The biggest problem I have with the GDPR is that it's so annoying that I would beg to have it go away."* Moreover, P20 mentioned that many requirements might not be suitable for XR:*"I don't made any experience that's telling the people ahead that we're tracking your motion data because that's kind of ruins the magic of the [XR] experience."* Based on this, P2, P4, P6, P8, P11, and P18-22 called for more XR-specific regulations to address the unique use cases (e.g., notification, immersive social experiences) in XR applications.

Better consideration for underrepresented groups. All female developers (P1, P11, P18, and P19), raised concerns about insufficient support in XR for underrepresented groups. P11 pointed out that physiological attacks disproportionately affect women, since many XR testing focus primarily on male participants, and called on developers to address this imbalance. This mirrors our finding here that only female participants flagged these issues. P19 described having experienced social attacks: *"I got attacked a couple of times, especially when your avatar is a woman or a minor"* and urged stronger anti-discrimination measures and requirements in social XR experiences. Besides the worry of female developers, many developers (P6-13, P16, P17, P19, P21, and P23) emphasized that current XR platforms lack adequate protections for children, underscoring the urgent need for age-appropriate security safeguards in the XR community.

2) Technical Support: Developers also called for stronger technical support as follows:

Moderation tools: P2, P5, P8, P9, P11, P13, P14, P16, P17 and P19-22 requested moderation tools (e.g., vision-based) for content and user behavior management in XR applications. P13 specifically highlighted the need for integrated moderation tools to block malicious content.

I/O processing tools: P2, P6, P7, P8, P10, P11, P14, P16, and P19 emphasized the need for input/output processing tools. For instance, P6 suggested differential privacy-based automatic blurring as an output processing method for mitigating shoulder surfing attacks, while P8 highlighted that input checking, such as packet inspection can help prevent input attacks.

User education and notification systems: P1, P4, P6-10, P13, P16, P17, and P19-22 emphasized the need for improved user education and notification systems as a technique to mitigate content, social, shoulder-surfing, input, and physiology attacks. For example, P16 stressed the importance of implementing an interaction design to notify users of content attacks.

3) Communication Support: Developers also identified gaps in communication support as follows:

Communication for developers to understand XR S&P knowledges. P1, P2, P4, P5, P7, P11-13, P16, and P18-22 stressed the importance of providing more threats knowledge with clear guidance on S&P in XR development. P2 highlighted: *"When I see, like XR related posts it's always about. Oh, what's the new tech like AI and stuff like that? But it's never about like security. So I think, having the community participate and educating everyone about it, and then giving suggestions."* Moreover, all developers hope researchers could share findings promptly and educate the broader community through accessible channels (e.g., online tutorials and developer summits). P1, P2, P8, P15, and P19 also mentioned the need for more communication channels dedicated to S&P. The absence of such channels may also contribute to the existing communication problems (e.g., disclosure on S&P issues).

Communication for technical transparency (P2, P4, P5, P6, P7, P8, P9, P11, P12, P13, P14, P16 and P19). Developers expressed concerns about the lack of communication on technical transparency in the XR ecosystem, particularly regarding the black-box nature of third-party APIs (e.g., Avatar SDK, XR Tracker), which they feared could introduce S&P risks. Participants emphasized the need for clearer communication, reliability, and openness from third-party providers.

C. Analysis: Misconceptions on the supports for S&P in XR

Although most perceptions of unmet needs are valid, our analysis also identifies the following misconceptions.

Misconceptions on policies and standards. Developers mistakenly believe that there are no XR-specific S&P policies or standards. While we agree that developing more XR-specific S&P policies is beneficial, as discussed in Section VI-A, there are already some initial efforts in this area. Addressing developers' lack of awareness of existing policies is also a crucial part of the problem [84], [100], [43].

Misconceptions on technical support. Developers also mistakenly believe that there are no usable tools for mitigating the proposed threats (e.g., content attacks). While we agree that additional tools are needed to address XR S&P threats, there have already been initial efforts and tools made available to developers [27], [101], [102]. However, a lack of awareness among developers about these existing mitigation tools significantly limits their effectiveness.

Misconceptions on communication. Besides, developers hold the misconception that there are no usable communication channels for XR S&P. While we agree that more communication channels and feedback mechanisms would be beneficial, there are existing platforms—such as forums [104], [105], [106] and conferences [107], [108]—that already serve this purpose. Developers should proactively engage in S&P discussions through these channels, participate in bug bounty programs, and follow up with users to close the feedback loop.

VII. Discussion

In this work, we investigate how developers perceive and reason about real-world XR S&P threats. By examining

their perceptions of the sensitivity, realism, importance, and practicality of threats, we uncover gaps in development (see Section IV and V) and highlight key factors (see Section VI). Moreover, our comprehensive study reveals two common misconceptions: (1) awareness gaps regarding XR threats, and (2) diffusion of responsibility when discussing mitigation practices for XR threats. In this section, we further examine the consequences of these misconceptions, analyze their underlying causes, and propose corresponding solutions. We also discussed the limitation and future directions in Appendix A.

A. Awareness gaps due to the rapidly evolving nature of XR interaction technologies.

In Sections IV-A, V-A, and VI-A, we identified significant gaps in developers' awareness of emerging XR threats, existing mitigation tools, and support resources (e.g., reward programs [109], guidance [110]). These awareness gaps lead to normalcy bias—developers tend to assume the threats as safe because they are unaware of potential risks (e.g., perception attacks are widely underestimated, especially in high-stakes contexts like military or medical XR) [111], [112]. This underestimation reduces developers' motivation and capacity to implement mitigation, increasing the likelihood of suboptimal practices such as insecure interaction designs (e.g., poorly considered voice control in immersive contexts) and missing safeguards (e.g., unmoderated social XR). As noted earlier, developers identify these factors as amplifying XR threats.

We believe that this awareness gap exists due to the emerging nature of XR interaction technologies, where threats are closely tied to new interaction designs, especially the new modalities of data (e.g., motion, voice) and the heavy reliance on them to fully enable basic interactions in XR, rendering an inherent and difficult privacy-utility tradeoff. Knowledge about these emerging XR threats and mitigation support is difficult for developers to access because it evolves rapidly and is scattered across various sources. Unlike mature platforms such as mobile or PC, where established guidelines are readily available (e.g., [113], [114]), XR S&P findings are typically shared through academic publications, industry reports, or informal developer-to-developer communication, requiring developers' active engagement.

Hence, we advocate for the development of standardized communication channels for XR developers [115], [116], especially those integrated into tools that support the normal development workflow (e.g., IDE), as prior research in the mobile context suggests that these just-in-time reminders contextualized in the app implementation help raise developers' awareness and encourage the adoption of secure and privacy-preserving designs [117]. We argue that it is important to empower developers to proactively follow emerging XR trends, engage in community discussions, and utilize existing channels (e.g., forums [118], feedback systems, and app reviews) to stay informed of new threats and mitigation supports.

B. Diffusion of responsibility due to the user-experience oriented nature of XR threats.

From developers' perceptions on threats and mitigation (Section IV,V,VI), we identified a trend of diffusion of responsibilities [119] (i.e., the tendency to assume that addressing XR S&P issues is someone else's responsibility). For example, developers suggested that users can mitigate physiology attacks by removing the headset, a problematic stance that disrupts the user experience and shifts both responsibility and risk onto users. Although users and platform providers also share responsibility, developers play a crucial role in collaboratively addressing these risks [120], [121].

We believe this diffusion of responsibility is especially pronounced in XR because threats (e.g., physiology attacks that induce motion sickness) span both hardware and software domains and directly impact the users' body and mind [122], [14]. Unlike traditional exploits (e.g., remote code execution) that map cleanly onto traditional vulnerability taxonomies, XR threats often look like UX, hardware, or platform issues, causing a misalignment in developers' mental models for threat identification and mitigation. Notably, although major platform providers develop built-in S&P protection, emphasizing considerations such as data minimization, transparency, control, and on-device data processing [123], they are inevitably limited to relatively coarse-grained mitigations. For instance, although Apple Vision Pro confines camera and sensor inputs (e.g., eye-tracking) to on-device processing, an application-level gaze-driven UI still reveals where users are looking through button selections—effectively leaking the eye-tracking data [124]. Failing to effectively engage developers to mitigate threats represents a missed opportunity to appropriately balance S&P considerations with the nuanced demands of functionality and the complexities of social contexts.

We advocate for the development of a practical industry framework for XR that clearly assigns ownership of emerging threats to designated stakeholders and provides S&P checklists for each threat. This approach has proven effective in encouraging responsible development in other domains, such as the Web [125], mobile [126], and IoT [127]. While there have been initial efforts in this space (e.g., XRSI's recommendations [110]), our findings highlight a pressing need for: (1) more detailed practice guidelines that address real-world threat scenarios encountered in XR (e.g., social harassment [128], motion as a side channel [8]), similar to the OWASP Top Ten [125] in web development; and (2) accompanying enforcement mechanisms to motivate secure development practices (e.g., incentives or penalties from policymakers for S&P, additional testing support from users, and auditing tool development from researchers and organizations). Such an approach is critical for fostering greater awareness among developers about their S&P responsibilities and for motivating proactive ownership of XR threats.

VIII. Conclusion

In this work, we conducted an interview study with 23 professional XR developers to investigate their perceptions and

responses to threats and mitigation in XR development. This study represents an initial effort to address XR S&P issues through a developer-centered, threat-aware perspective. Our findings reveal that developers recognize XR S&P risks as closely tied to, and often amplified by, development decisions. However, there is limited awareness of these issues among developers, compounded by a lack of strategic, technical, and communication support from key stakeholders within the XR community. To address these issues, we analyzed developers' reasoning around threats and mitigation strategies to uncover existing problems in XR development and propose actionable, stakeholder-aware solutions. Our proposed future directions clarify responsibilities for key stakeholders and encourage stronger collaboration across the XR ecosystem. The insights from this study offer valuable directions for future research and practical advancements in XR S&P.

Acknowledgment

We sincerely thank the reviewers for their valuable feedback on the paper. This work is supported in part by the National Science Foundation (NSF) Awards 2323105, 2317184. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of sponsors.

References

[1] Y. Shinde, "Extended reality market to hit usd 519.5 bn by 2032," Jul 2024. [Online]. Available: https://scoop.market.us/extended-reality-market-news/

[2] J. Soroushian, "Thinking ahead about XR: Privacy and security in an immersive world," Bipartisan Policy Center, Jul 2021. [Online]. Available: https://bipartisanpolicy.org/blog/thinking-ahead-about-xr-privacy-and-security-in-an-immersive-world

[3] D. Weinstein, "What is extended reality?" Dec 2024. [Online]. Available: https://blogs.nvidia.com/blog/what-is-extended-reality/

[4] H. J. Smith and M. Neff, "Communication behavior in embodied virtual reality," in CHI, 2018.

[5] A. Plopski, T. Hirzle, N. Norouzi, L. Qian, G. Bruder, and T. Langlotz, "The eye in extended reality: A survey on gaze interaction and eye tracking in head-worn extended reality," CSUR, 2022.

[6] M. Speicher, B. D. Hall, and M. Nebeling, "What is mixed reality?" in CHI, 2019.

[7] Z. Yang, Z. Sarwar, I. Hwang, R. Bhaskar, B. Y. Zhao, and H. Zheng, "Can virtual reality protect users from keystroke inference attacks?" arXiv preprint arXiv:2310.16191, 2023.

[8] Z. Su, K. Cai, R. Beeler, L. Dresel, A. Garcia, I. Grishchenko, Y. Tian, C. Kruegel, and G. Vigna, "Remote keylogging attacks in multi-user vr applications," arXiv preprint arXiv:2405.14036, 2024.

[9] Y. Zhang, C. Slocum, J. Chen, and N. Abu-Ghazaleh, "It's all in your head (set): Side-channel attacks on {AR/VR} systems," in USENIX Security, 2023.

[10] K. Cheng, J. F. Tian, T. Kohno, and F. Roesner, "Exploring user reactions and mental models towards perceptual manipulation attacks in mixed reality," in USENIX Security, 2023.

[11] P. Casey, I. Baggili, and A. Yarramreddy, "Immersive virtual reality attacks and the human joystick," IEEE Transactions on Dependable and Secure Computing, 2019.

[12] W.-J. Tseng, E. Bonnail, M. McGill, M. Khamis, E. Lecolinet, S. Huron, and J. Gugenheimer, "The dark side of perceptual manipulations in virtual reality," in CHI, 2022.

[13] C. Shi, X. Xu, T. Zhang, P. Walker, Y. Wu, J. Liu, N. Saxena, Y. Chen, and J. Yu, "Face-mic: inferring live speech and speaker identity via subtle facial dynamics captured by ar/vr motion sensors," in Mobicom, 2021.

[14] A. Giaretta, "Security and privacy in virtual reality: a literature survey," Virtual Reality, 2024.

[15] J. A. De Guzman, K. Thilakarathna, and A. Seneviratne, "Security and privacy approaches in mixed reality: A literature survey," CSUR, 2019.

[16] V. Nair, L. Rosenberg, J. F. O'Brien, and D. Song, "Truth in motion: The unprecedented risks and opportunities of extended reality motion data," 2023.

[17] M. Wierzbowski, G. Pochwatko, P. Borkiewicz, D. Cnotkowski, M. Pabiś-Orzeszyna, and P. Kobyliński, "Behavioural biometrics in virtual reality: To what extent can we identify a person based solely on how they watch 360-degree videos?" in ISMAR-Adjunct. IEEE, 2022.

[18] V. Nair, W. Guo, J. Mattern, R. Wang, J. F. O'Brien, L. Rosenberg, and D. Song, "Unique identification of 50,000+ virtual reality users from head & hand motion data," in USENIX Security, 2023.

[19] C. Slocum, Y. Zhang, N. Abu-Ghazaleh, and J. Chen, "Going through the motions:{AR/VR} keylogging from user head motions," in USENIX Security, 2023.

[20] Y. Acar, S. Fahl, and M. L. Mazurek, "You are not your developer, either: A research agenda for usable security and privacy research beyond end users," SecDev, 2016.

[21] C. Katins, P. W. Woźniak, A. Chen, I. Tumay, L. V. T. Le, J. Uschold, and T. Kosch, "Assessing user apprehensions about mixed reality artifacts and applications: The mixed reality concerns (mrc) questionnaire," in CHI, 2024.

[22] A. SB, A. Agrawal, Y. Yao, Y. Zou, and A. Das, ""what are they gonna do with my data?": Privacy expectations, concerns, and behaviors in virtual reality," PETS, 2025.

[23] S. A. Horstmann, S. Domiks, M. Gutfleisch, M. Tran, Y. Acar, V. Moonsamy, and A. Naiakshina, ""those things are written by lawyers, and programmers are reading that." mapping the communication gap between software developers and privacy experts," PETS, 2024.

[24] M. Gutfleisch, J. H. Klemmer, N. Busch, Y. Acar, M. A. Sasse, and S. Fahl, "How does usable security (not) end up in software products? results from a qualitative interview study," in IEEE S&P, 2022.

[25] D. Adams, A. Bah, C. Barwulor, N. Musaby, K. Pitkin, and E. M. Redmiles, "Ethics emerging: the story of privacy and security perceptions in virtual reality," in SOUPS, 2018.

[26] G. M. Garrido, V. Nair, and D. Song, "Sok: Data privacy in virtual reality," arXiv preprint arXiv:2301.05940, 2023.

[27] V. Nair, G. M. Garrido, and D. Song, "Going incognito in the metaverse," arXiv preprint arXiv:2208.05604, 2022.

[28] H. Farrukh, R. Mohamed, A. Nare, A. Bianchi, and Z. B. Celik, "{LocIn}: Inferring semantic location from spatial maps in mixed reality," in USENIX Security, 2023.

[29] J. A. d. Guzman, A. Seneviratne, and K. Thilakarathna, "Unravelling spatial privacy risks of mobile mixed reality data," IMWUT, 2021.

[30] S. Luo, A. Nguyen, H. Farooq, K. Sun, and Z. Yan, "Eavesdropping on controller acoustic emanation for keystroke inference attack in virtual reality," in NDSS, 2024.

[31] L. Blackwell, N. Ellison, N. Elliott-Deflo, and R. Schwartz, "Harassment in social virtual reality: Challenges for platform governance," CSCS, 2019.

[32] G. Freeman, S. Zamanifard, D. Maloney, and D. Acena, "Disturbing the peace: Experiencing and mitigating emerging harassment in social virtual reality," CSCW, 2022.

[33] P. Jansen and F. Fischbach, "The social engineer: An immersive virtual reality educational game to raise social engineering awareness," in Extended Abstracts of the 2020 Annual Symposium on Computer-Human Interaction in Play, 2020.

[34] K. Schulenberg, G. Freeman, L. Li, and C. Barwulor, "" creepy towards my avatar body, creepy towards my body": How women experience and manage harassment risks in social virtual reality," CSCW, 2023.

[35] Q. Zheng, S. Xu, L. Wang, Y. Tang, R. C. Salvi, G. Freeman, and Y. Huang, "Understanding safety risks and safety design in social vr environments," CSCW, 2023.

[36] K. Cheng, A. Bhattacharya, M. Lin, J. Lee, A. Kumar, J. F. Tian, T. Kohno, and F. Roesner, "When the user is inside the user interface: An empirical study of ui security properties in augmented reality," in USENIX Security, 2024.

[37] C. Slocum, Y. Zhang, E. Shayegani, P. Zaree, N. Abu-Ghazaleh, and J. Chen, "That doesn't go there: Attacks on shared state in {Multi-User} augmented reality applications," in USENIX Security, 2024.

[38] S. Tariq, A. Abuadbba, and K. Moore, "Deepfake in the metaverse: Security implications for virtual gaming, meetings, and offices,"

in Proceedings of the 2nd Workshop on Security Implications of Deepfakes and Cheapfakes, 2023.
[39] A. Al Arafat, Z. Guo, and A. Awad, "Vr-spy: A side-channel attack on virtual key-logging in vr headsets," in IEEE VR, 2021.
[40] R. Trimananda, H. Le, H. Cui, J. T. Ho, A. Shuba, and A. Markopoulou, "{OVRseen}: Auditing network traffic and privacy policies in oculus {VR}," in USENIX Security, 2022.
[41] H. Guo, H.-N. Dai, X. Luo, Z. Zheng, G. Xu, and F. He, "An empirical study on oculus virtual reality applications: Security and privacy perspectives," in ICSE, 2024.
[42] Cyber XR Coalition, "Immersive technology standards for privacy, safety, and security," Cyber XR Coalition, Tech. Rep., 2021, accessed July 22, 2025. [Online]. Available: https://cyberxr.org/wp-content/uploads/2021/05/Immersive_Technology_Standards.pdf
[43] L. Rosenberg, C. Wallace, K. Pearlman, and B. Choudhary, "The metaverse and standards," 05 2023.
[44] D. Votipka, K. R. Fulton, J. Parker, M. Hou, M. L. Mazurek, and M. Hicks, "Understanding security mistakes developers make: Qualitative analysis from build it, break it, fix it," in USENIX Security, 2020.
[45] H. Assal and S. Chiasson, "Security in the software development lifecycle," in SOUPS, 2018.
[46] T. Li, E. Louie, L. Dabbish, and J. I. Hong, "How developers talk about personal data and what it means for user privacy: A case study of a developer forum on reddit," CSCW, 2021.
[47] M. Tahaei, A. Frik, and K. Vaniea, "Privacy champions in software teams: Understanding their motivations, strategies, and challenges," in CHI, 2021.
[48] T. Li, K. Reiman, Y. Agarwal, L. F. Cranor, and J. I. Hong, "Understanding challenges for developers to create accurate privacy nutrition labels," in CHI, 2022.
[49] E. Hine, I. N. Rezende, H. Roberts, D. Wong, M. Taddeo, and L. Floridi, "Safety and privacy in immersive extended reality: An analysis and policy recommendations," Digital Society, 2024.
[50] S. Das, C. Faklaris, J. I. Hong, and L. A. Dabbish. Now Foundations and Trends, 2023.
[51] S. Abhinaya, A. Sabir, and A. Das, "Enabling developers, protecting users: Investigating harassment and safety in vr," arXiv e-prints, pp. arXiv–2403, 2024.
[52] J. Cao, A. Das, P. Emami-Naeini et al., "Understanding parents' perceptions and practices toward children's security and privacy in virtual reality," arXiv preprint arXiv:2403.06172, 2024.
[53] M. Abraham, P. Saeghe, M. Mcgill, and M. Khamis, "Implications of xr on privacy, security and behaviour: Insights from experts," in Nordic CHI, 2022.
[54] S. A. Karre, N. Mathur, and Y. R. Reddy, "Is virtual reality product development different? an empirical study on vr product development practices," in ISEC, 2019.
[55] DARPA, "Intrinsic cognitive security (ics)," Oct 2023. [Online]. Available: https://sam.gov/opp/cfaf7a3e51fc4f62ae4120f88d52f418/view
[56] S. Pahi and C. Schroeder, "Extended privacy for extended reality: XR technology has 99 problems and privacy is several of them," Notre Dame Journal on Emerging Technologies, apr 2023. [Online]. Available: https://ndlsjet.com/extended-privacy-for-extended-reality-xr-technology-has-99-problems-and-privacy-is-several-of-them/
[57] Y. Wu, C. Shi, T. Zhang, P. Walker, J. Liu, N. Saxena, and Y. Chen, "Privacy leakage via unrestricted motion-position sensors in the age of virtual reality: A study of snooping typed input on virtual keyboards," in IEEE Security and Privacy, 2023.
[58] D. Cayir, R. Mohamed, R. Lazzeretti, M. Angelini, A. Acar, M. Conti, Z. B. Celik, and S. Uluagac, "Speak up, i'm listening: Extracting speech from zero-permission vr sensors," in NDSS, 2025.
[59] Y. Chandio, N. Bashir, and F. M. Anwar, "Stealthy and practical multimodal attacks on mixed reality tracking," in AIxVR. IEEE, 2024.
[60] T. Zhang, Z. Ye, A. T. Mahdad, M. M. R. R. Akanda, C. Shi, Y. Wang, N. Saxena, and Y. Chen, "Facereader: Unobtrusively mining vital signs and vital sign embedded sensitive info via ar/vr motion sensors," in ACM CCS, 2023.
[61] Z. Ye, A. T. Mahdad, Y. Wang, C. Shi, Y. Chen, and N. Saxena, "Bpsniff: Continuously surveilling private blood pressure information in the metaverse via unrestricted inbuilt motion sensors," in IEEE Security and Privacy, 2025.
[62] N. Noah and S. Das, "Privacy and security in extended reality: Exploring the risks of external biometric data collection," Available at SSRN 4780358, 2024.
[63] I. Jarin, Y. Duan, R. Trimananda, H. Cui, S. Elmalaki, and A. Markopoulou, "Behavr: User identification based on vr sensor data," arXiv preprint arXiv:2308.07304, 2023.
[64] H. Wang, Z. Zhan, H. Shan, S. Dai, M. Panoff, and S. Wang, "Gazeploit: Remote keystroke inference attack by gaze estimation from avatar views in vr/mr devices," in ACM CCS, 2024.
[65] J. Gugenheimer, W.-J. Tseng, A. H. Mhaidli, J. O. Rixen, M. McGill, M. Nebeling, M. Khamis, F. Schaub, and S. Das, "Novel challenges of safety, security and privacy in extended reality," in CHI Extended Abstracts, 2022.
[66] S. R. K. Gopal, D. Shukla, J. D. Wheelock, and N. Saxena, "Hidden reality: Caution, your hand gesture inputs in the immersive virtual world are visible to all!" in USENIX Security, 2023.
[67] T. Denning, Z. Dehlawi, and T. Kohno, "In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies," in CHI, 2014.
[68] K. Lebeck, T. Kohno, and F. Roesner, "How to safely augment reality: Challenges and directions," in Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications, 2016.
[69] S. Dastgerdy, "Virtual reality and augmented reality security: A reconnaissance and vulnerability assessment approach," arXiv preprint arXiv:2407.15984, 2024.
[70] J. Shang and J. Wu, "Secure voice input on augmented reality headsets," IEEE Transactions on Mobile Computing, 2020.
[71] D. Cayir, A. Acar, R. Lazzeretti, M. Angelini, M. Conti, and S. Uluagac, "Augmenting security and privacy in the virtual realm: An analysis of extended reality devices," IEEE Security & Privacy, 2023.
[72] M. El-Hajj, "Cybersecurity and privacy challenges in extended reality: Threats, solutions, and risk mitigation strategies," in Virtual Worlds. MDPI, 2024.
[73] E. Bonnail, W.-J. Tseng, M. Mcgill, E. Lecolinet, S. Huron, and J. Gugenheimer, "Memory manipulations in extended reality," in CHI, 2023.
[74] D. Freeman, P. Haselton, J. Freeman, B. Spanlang, S. Kishore, E. Albery, M. Denne, P. Brown, M. Slater, and A. Nickless, "Automated psychological therapy using immersive virtual reality for treatment of fear of heights: a single-blind, parallel-group, randomised controlled trial," The Lancet Psychiatry, 2018.
[75] J. Saldaña, "The coding manual for qualitative researchers," 2021.
[76] N. McDonald, S. Schoenebeck, and A. Forte, "Reliability and inter-rater reliability in qualitative research: Norms and guidelines for cscw and hci practice," CSCW, 2019.
[77] M. L. McHugh, "Interrater reliability: the kappa statistic," Biochemia medica, 2012.
[78] GDPR Advisor, "Gdpr and augmented reality advertising: Ensuring consumer privacy," https://www.gdpr-advisor.com/gdpr-and-augmented-reality-advertising-ensuring-consumer-privacy/, 2025.
[79] U.S. Department of Veterans Affairs, "Va publications," https://www.va.gov/vapubs/, n.d.
[80] E. Games, "Photon: The networking library for unity." [Online]. Available: https://www.photonengine.com
[81] XR4Europe, "European xr industry report 2025," XR4Europe, Tech. Report, May 2025. [Online]. Available: https://xr4europe.eu/wp-content/uploads/European-XR-Industry-Report-2025.pdf
[82] P. Sykownik, D. Maloney, G. Freeman, and M. Masuch, "Something personal from the metaverse: goals, topics, and contextual factors of self-disclosure in commercial social vr," in CHI, 2022.
[83] D. Maloney, S. Zamanifard, and G. Freeman, "Anonymity vs. familiarity: Self-disclosure and privacy in social virtual reality," in VRST, 2020.
[84] European Parliament and Council of the European Union, "Regulation (EU) no 2016/679: General data protection regulation (GDPR)," Official Journal of the European Union, May 2016. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj
[85] Office of the Vice Provost for Research, Lehigh University, "Guidance: Data security and privacy," Online, 2025. [Online]. Available: https://research.lehigh.edu/policies-guidance-forms/guidance-data-security-and-privacy
[86] Office of Research Cyberinfrastructure, University of Central Florida, "Research data privacy," Online, 2025. [Online]. Available: https://rci.research.ucf.edu/resource/research-data-privacy/

[87] Staff, Faculty of Archaeology, Leiden University, "Personal and sensitive data," Online, 2025. [Online]. Available: https://www.staff.universiteitleiden.nl/vr/archaeology/research-data-management-in-archaeology/personal-and-sensitive-data

[88] D. Dangwal, M. Cowan, A. Alaghi, V. T. Lee, B. Reagen, and C. Trippel, "Sok: Opportunities for software-hardware-security codesign for next generation secure computing," in Proceedings of the 9th International Workshop on Hardware and Architectural Support for Security and Privacy, 2020.

[89] M. Yang, T. Ahmed, S. Inagaki, K. Sakiyama, Y. Li, and Y. Hara-Azumi, "Hardware/software cooperative design against power side-channel attacks on iot devices," IEEE Internet of Things Journal, 2024.

[90] A. Dubey, R. Cammarota, A. Varna, R. Kumar, and A. Aysu, "Hardware-software co-design for side-channel protected neural network inference," in HOST. IEEE, 2023.

[91] D. Castro, "AR/VR Poses New Content Moderation Challenges That Policymakers Should Address," ITIF, Tech. Rep., Feb. 2022. [Online]. Available: https://itif.org/publications/2022/02/28/arvr-poses-new-content-moderation-challenges-policymakers-should-address

[92] T. Porcino, D. Reilly, E. Clua, and D. Trevisan, "A guideline proposal for minimizing cybersickness in vr-based serious games and applications," in SeGAH. IEEE, 2022.

[93] C. Wijayarathna and N. A. G. Arachchilage, "Am i responsible for end-user's security? a programmer's perspective," arXiv preprint arXiv:1808.01481, 2018.

[94] P. L. Gorski, L. L. Iacono, and M. Smith, "Eight lightweight usable security principles for developers," IEEE Security & Privacy, 2022.

[95] Y. Wang, "Inclusive security and privacy," IEEE Security & Privacy, 2018.

[96] A. Shostack, Threat modeling: Designing for security. John wiley & sons, 2014.

[97] A. Industries, "Anduril and meta team up to transform xr for the american military," 2025, accessed: 2025-07-14. [Online]. Available: https://www.anduril.com/article/anduril-and-meta-team-up-to-transform-xr-for-the-american-military/

[98] L. Nurture, "How augmented reality (ar) is transforming healthcare in 2025: Benefits and applications," 2025, accessed: 2025-07-14. [Online]. Available: https://www.letsnurture.com/blog/how-augmented-reality-ar-is-transforming-healthcare-in-2025-benefits-and-applications.html

[99] J. Gogoll, N. Zuber, S. Kacianka, T. Greger, A. Pretschner, and J. Nida-Rümelin, "Ethics in the software development process: from codes of conduct to ethical deliberation," Philosophy & Technology, 2021.

[100] California State Legislature, "California consumer privacy act," 2018. [Online]. Available: https://oag.ca.gov/privacy/ccpa

[101] NSA's Research Directorate, "Ghidra," https://ghidra-sre.org/, 2024.

[102] Mikadumont, "Code analysis using .net compiler platform (roslyn) analyzers - visual studio (windows)," 2024. [Online]. Available: https://learn.microsoft.com/en-us/visualstudio/code-quality/roslyn-analyzers-overview?view=vs-2022

[103] K. Lebeck, K. Ruth, T. Kohno, and F. Roesner, "Arya: Operating system support for securely augmenting reality," IEEE S&P, 2018.

[104] "Developer category - meta community forums," https://communityforums.atmeta.com/category/developer, 2025, accessed: 2025-07-17.

[105] "Xr & spatial computing – nvidia developer forums," https://forums.developer.nvidia.com/c/omniverse/xr-spatial-computing/707, 2025, accessed: 2025-07-17.

[106] "Xr community – immersiveunity," https://xrcommunity.immersiveunity.com/, 2025.

[107] "Awe usa 2025: The world's #1 xr event," in Augmented World Expo (AWE), 2025, immersive spatial computing expo, June 2025. [Online]. Available: https://www.awexr.com/usa-2025

[108] "IEEE conference on virtual reality and 3d user interfaces," in IEEE VR, 2025, 32nd annual conference held in Saint-Malo, France. [Online]. Available: https://ieeevr.org/2025/

[109] Meta Platforms, Inc., "Building the metaverse responsibly," https://about.fb.com/news/2021/09/building-the-metaverse-responsibly/, 2021, accessed: 2025-04-20.

[110] XRSI, "The metaverse and standards," Jul 2023. [Online]. Available: https://xrsi.org/publication/the-metaverse-and-standards

[111] S. Chattopadhyay, N. Nelson, A. Au, N. Morales, C. Sanchez, R. Pandita, and A. Sarma, "A tale from the trenches: cognitive biases and software development," in ICSE, 2020.

[112] S. Magazine. (2021) The next frontier in cybersecurity: Mitigating normalcy bias. Accessed: 2025-04-16. [Online]. Available: https://www.securitymagazine.com/articles/96934-the-next-frontier-in-cybersecurity-mitigating-normalcy-bias

[113] Google, "Android security best practices," https://source.android.com/docs/security/best-practices, 2025, accessed: 2025-04-20.

[114] K. A. Scarfone, W. Jansen, and M. Tracy, "Sp 800-123. guide to general server security," 2008.

[115] V. Krauß, A. Boden, L. Oppermann, and R. Reiners, "Current practices, challenges, and design implications for collaborative ar/vr application development," in CHI, 2021.

[116] XR Association, "XR Association Releases State of the Industry Report," 2023. [Online]. Available: https://xra.org/xr-association-releases-state-of-the-industry-report/

[117] T. Li, Y. Agarwal, and J. I. Hong, "Coconut: An ide plugin for developing privacy-friendly apps," IMWUT, 2018.

[118] "Meta community forums," https://communityforums.atmeta.com/, accessed: 2025-04-20.

[119] J. M. Darley and B. Latané, "Bystander intervention in emergencies: diffusion of responsibility." Journal of personality and social psychology, 1968.

[120] A. Naiakshina, A. Danilova, C. Tiefenau, M. Herzog, S. Dechand, and M. Smith, "Why do developers get password storage wrong? a qualitative usability study," in ACM CCS, 2017.

[121] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner, "Android permissions: User attention, comprehension, and behavior," in SOUPS, 2012.

[122] F. Roesner, T. Kohno, and D. Molnar, "Security and privacy for augmented reality systems," Communications of the ACM, 2014.

[123] Apple Inc., "Apple Vision Pro Privacy Overview," February 2024. [Online]. Available: https://www.apple.com/privacy/docs/Apple_Vision_Pro_Privacy_Overview.pdf

[124] M. S. Smith, "Privacy is just no longer a thing in augmented reality?" IEEE Spectrum, 2024, accessed: 2025-04-23. [Online]. Available: https://spectrum.ieee.org/apple-vision-pro-privacy

[125] Open Worldwide Application Security Project (OWASP), "Owasp top ten," https://owasp.org/www-project-top-ten/, 2021, accessed: 2025-04-22.

[126] Z. R. Alkindi, M. Sarrab, and N. Alzeidi, "User privacy and data flow control for android apps: Systematic literature review," Journal of Cyber Security and Mobility, 2021.

[127] I. Brass, L. Tanczer, M. Carr, M. Elsden, and J. Blackstock, "Standardising a moving target: The development and evolution of iot security standards," in Living in the Internet of Things: Cybersecurity of the IoT. IET, 2018.

[128] A. Alhakamy, "Extended reality (xr) toward building immersive solutions: the key to unlocking industry 4.0," ACM Computing Surveys, 2024.

[129] E. M. Redmiles, N. Warford, A. Jayanti, A. Koneru, S. Kross, M. Morales, R. Stevens, and M. L. Mazurek, "A comprehensive quality evaluation of security and privacy advice on the web," in USENIX Security, 2020.

[130] R. Cummings, G. Kaptchuk, and E. M. Redmiles, "" i need a better description": An investigation into user expectations for differential privacy," in ACM CCS, 2021.

[131] A. Zenner, M. Speicher, S. Klingner, D. Degraen, F. Daiber, and A. Krüger, "Immersive notification framework: Adaptive & plausible notifications in virtual reality," in Extended abstracts of CHI, 2018.

[132] N. Sabri, B. Chen, A. Teoh, S. P. Dow, K. Vaccaro, and M. Elsherief, "Challenges of moderating social virtual reality," in CHI, 2023.

[133] M. Lyu, R. D. Tripathi, and V. Sivaraman, "Metavradar: Measuring metaverse virtual reality network activity," POMACS, 2023.

[134] Meta, "Hand and body privacy notice," 2024. [Online]. Available: https://www.meta.com/help/quest/articles/accounts/privacy-information-and-settings/hand-tracking-privacy-notice/

[135] Meta, "Meta quest virtual reality check (vrc) guidelines," https://developer.oculus.com/resources/publish-quest-req/, 2024.

[136] F. Roesner and T. Kohno, "Security and privacy for augmented reality: Our 10-year retrospective," in VR4Sec: 1st International Workshop on Security for XR and XR for Security, 2021.

APPENDIX

A. Limitations and Future Work

Since our study involve demonstrations of XR risks, there is a potential for bias in developers' perceptions of S&P. However, this is unavoidable as our aim was to capture their views on broader XR S&P threats, despite possible gaps in their knowledge. It is also a common practice for research in threat-aware developer studies on other platforms [129], [130].

Our work takes the first step towards understanding developers' perceptions and challenges for building S&P-friendly XR apps. There would be many valuable future directions building on top of our work, such as vetting developers' actual implementations for potential S&P issues, and development frameworks that support S&P-by-design XR developments.

B. Labeling Criteria and Examples

We provide additional details to elaborate on our labeling criteria from Section III-C.

Awareness of attacks: Participants were labeled "aware" if they met any of the following criteria; otherwise, "unaware".
(1) Proactively mentioned attacks, e.g., discussing social attacks or experiences before the demonstration of this attack.
(2) Proactively provided additional or similar examples during demonstration, e.g., mentioning user-generated light bombs triggering seizures (physiological attack).
(3) Clearly stated familiarity during or after demonstration, e.g., sharing shoulder-surfing experiences or knowing developers with mitigations for physiological/social attacks.

Awareness of mitigation: Participants were labeled "unaware" if they met either criterion below; otherwise, "aware":
(1) Explicitly stated unawareness or requested details during/after presentation, e.g., first-time hearing about mitigations like program analysis, or asking for resources.
(2) Misused tools in the matching task, e.g., suggesting Ghidra [101] (static analysis) for dynamic social attacks, indicating misunderstanding.

Quality of developer-proposed mitigation: We rated proposed mitigations on a three-point scale:
(1) Score 1 - Ineffective or inapplicable: e.g., no solution provided (e.g., stating they do not know), vague or high-level (e.g., *"Yeah, just have good security measures."*), or irrelevant (e.g., for input attacks: *"Discomfort is like, okay, just stop using the Headset for a minute. Let the situation fix itself"*).
(2) Score 2 - On track but unclear/missing details: e.g., for content attack: *"This is honestly on the developer side, like, be mindful on the way you create apps to ensure that nobody can piggyback off of what you've created, and get access, and actually show their content in your content."* (correct direction but lacks specifics like API checks or program analysis). For input attacks: *"Maybe if there was a way to disable the hand overlay, and so you can see your true hands had passed through, or something. I'm not too certain about this one this time."* (suggests moderation but unsure/lacks details).
(3) Score 3 - Valid, clear, aligns with prior work or XR applications: e.g., for physiology attacks: *"From my personal experience, it was like efficient use of data to ensure that, like on the developer side, like, you're ensuring that your data is being handled correctly, so that there's no overhead and not a lot of overload on the system already, because by itself is very intensive, especially since you're rendering every scene twice. So you have to be mindful of like, what kind of assets you use, and if they're like. compliant or not. As for the other side, if somebody else is trying to do that, like DDoS attacks and stuff like that, honestly, that's security, and like networking."* (multiple valid solutions with examples).

C. Finalized Codebook

- RQ1:
  - Awareness of Threats
    - Mentioned without prompt
    - Provided additional examples
    - Clearly mentioned awareness of attacks
  - Awareness of Mitigation
    - Specifically mentioned unaware of a mitigation.
    - Misuse mitigation tools in follow-up questions
  - What makes data more sensitive
    - Immersive interaction design
    - Advanced tracking sensors
    - Economic value of user data
    - Potential surveillance by XR platform
    - Opaque XR infrastructure
    - Other (e.g., data abuse)
  - What makes certain data less sensitive
    - Significant benefits overshadow risks
    - Insufficient knowledge of developers
    - Don't believe data can leak sensitive information
    - Other (e.g., also collected in other platforms)
  - Why data leakage channels more realistic
    - Development misoperations
    - Relied on 3rd party APIs
    - Misused by user
  - Why certain data leakage channel less realistic
    - Developers unfamiliar with the channel
    - Unrelated to software development
    - Local apps are safe
  - What makes attacks in XR overall more important
    - Immersive experiences increase attack severity
    - Affect the growth of the XR industry
    - Broader negative impact in XR
    - Affect critical operations (e.g., banking) in XR
    - Hard-to-Mitigate characteristics in immersive social environment
    - Other (e.g., perceivable benefit for attackers)
  - What makes certain attacks in XR less important
    - User can mitigate
    - Only for certain scenarios
    - Only target certain user groups
    - Other (e.g., less important consequence)

TABLE IV: Mitigation tools and best practices for XR security and privacy

No.	Approach	Examples	Applicable to
1	User notification&education	Privacy policy [41], [40], in-app notification [131], etc.	Content attacks, Perception attacks, Data leakage, etc.
2	Static analysis	Ghidra [101], [41], Roslyn [102], etc.	Software side-channel, Content attacks, Data leakage, etc.
3	Runtime analysis	Moderation [132], packet analysis [133], [40], etc.	Social attacks, Physiology attacks, Data leakage, etc.
4	OS and Game Engine enforcement	Arya [103], OS level data control [134], etc.	Content attacks, Input attacks, Perception attacks, etc.
5	Legislation and policymaking	Laws [100], [84], best practices in XR [110], [135], etc.	Shoulder surfing attacks, Social attacks, Content attacks, etc.
6	Input and output processing	Input or output checking [136], [68], differential privacy [16], etc.	Software side-channel attacks, Shoulder surfing attacks, etc.

TABLE V: Data Leakage Channels in XR we Collected From Literatures and Discussions Like [8], [9], [14].

No.	Leakage Channel	Example
1	App designed functions	Motion data for rendering avatar
2	Hardware side-channels	Rendering side-channels
3	Third-party services	Analysis APIs
4	Insecure operations	Unsafe storage, unencrypted data

  - What makes attacks in XR overall more practical
    - More Stealthy attacks enabled by XR hardware
    - Flexibility of attacks in immersive space
    - Already seen similar attacks happening because of implementation issues
    - User-generated content introduces attacks
  - What makes certain attacks in XR less practical
    - Require technical knowledge from attackers
    - Require knowledge about the usecases/environment of the user
    - Less motivating to attackers
- RQ2:
  - Quality of mitigation:
    - Score1: The proposed mitigation is ineffective.
    - Score2: On the right track. Not very clear, mistaken, or missing important details.
    - Score3: Valid, clear results that align with solutions in prior research works or adopted in applications.
  - Limitations XR threats mitigations
    - Developer: Awareness of existing mitigation
    - Community: S&P design sacrificed utility
    - Community: high requirement of external resources
    - Community: Impractical or outdated resources
  - Developers' General Unmet Needs from The Entire XR Community - Communication
    - More Community Communication Channels: emerging attack channels
    - More Community Communication Channels: XR components that are vulnerable to attacks
    - Attack Mitigation
    - How to implement S&P practices in XR application
    - Lack of open platforms for communication
    - Potential solutions to threats
  - Developers' General Unmet Needs from The Entire XR Community - Strategy
    - Strategic Support : S&P hurting utility
    - Strategic Support : utility compromise S&P
    - Strategic Support : standard and best practices
  - Developers' General Unmet Needs from The Entire XR Community - Technical
    - Better Technical Solutions : moderation
    - Better Technical Solutions : I/O
    - Better Technical Solutions : hardware
    - Better Technical Solutions : authentication
    - Better Technical Solutions : notification
    - Better Technical Solutions : Other
  - Specific Unmet Needs from Each Key Stakeholder in XR Development
    - More Responsible API Provider : Third-Party API as leakage channels
    - More Responsible API Provider : Malicious Third-Party API
    - More Support from API Providers To Handle Updates
    - More Transparent Industry Infrastructures from XR Platform Providers
    - More Instructions from Policymakers
    - More Active and Responsible Publishing Platforms
    - More Engaged and Responsible Users
    - More Guidance from Researchers
- Other Information
  - Future Considerations: Challenge and Direction
    - AI
    - Hardware
    - Children
    - More Complicated Attacks
    - Bias and Ethics
    - Discuss all potential problems now, even if they're not happening yet
  - Developer Proposed Mitigation
    - Shoulder Surfing Mitigation
    - Software Sidechannel Mitigation
    - Input Mitigation
    - Social Mitigation
    - Content Mitigation
    - Perception Mitigation
    - Physiology Mitigation