A Multi-Layered Privacy Permission Framework for Extended Reality

Shady Mansour
shady.mansour@campus.lmu.de
LMU Munich
Munich, Germany

Verena Winterhalter
verena.winterhalter@ifi.lmu.de
LMU Munich
Munich, Germany

Florian Alt
florian.alt@ifi.lmu.de
LMU Munich
Munich, Germany
University of the Bundeswehr
Munich, Germany

Viktorija Paneva
viktorija.paneva@ifi.lmu.de
LMU Munich
Munich, Germany

Abstract

Extended Reality (XR) systems bring arrays of sensors closer to the user's body, enabling the collection of extensive user and contextual data, from motion and biometrics to behavioral analytics, that users might not be aware they are sharing. This poses significant risks to users' privacy. Yet, despite the immersive and dynamic nature of XR, most platforms still rely on static, text-based privacy mechanisms inherited from traditional 2D interfaces. We propose a new paradigm of *continuous consent in XR*, where privacy decisions unfold as a relational, context-aware, and renegotiable process embedded in the experience – not a single consent event. To this end, we propose a *Multi-Layered Privacy Framework* spanning five interdependent layers: regulatory compliance, technical implementation, permission models, user experience, and user perception and cognition. We then introduce the *User Privacy Journey Model*, which operationalizes the framework as a sequential user pathway: from onboarding and contextual prompts to in-experience control and post-session review, along with the *XR Privacy Checklist* to support practical adoption. By rethinking consent as a continuous journey, we present a new paradigm for XR privacy, one that opens a new research perspective on what "informed" consent means in immersive environments where the boundaries between self, system, and space are increasingly blurred.

CCS Concepts

• **Security and privacy → Usability in security and privacy**; • **Human-centered computing → HCI theory, concepts and models**; *Mixed / augmented reality*; *Virtual reality*.

ACM Reference Format:
Shady Mansour, Verena Winterhalter, Florian Alt, and Viktorija Paneva. 2025. A Multi-Layered Privacy Permission Framework for Extended Reality. In *New Security Paradigms Workshop (NSPW '25), August 24–27, 2025, Aerzen, Germany.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3774761.3774916

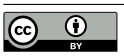

NSPW '25, Aerzen, Germany

ACM ISBN 979-8-4007-1875-5/2025/08
https://doi.org/10.1145/3774761.3774916

1 Introduction

Extended Reality (XR) systems are sensor-rich environments that can continuously sense, infer, and adapt – tracking user motion, behavior, emotion, intention, and even identity [27, 58, 59, 67]. What does privacy mean in such a context? We argue that existing privacy models fall short in XR as they treat consent as an isolated event, a one-time transaction, rather than a part of a dynamic, multifaceted, and evolving user journey.

XR gives us the opportunity — and the obligation — to rethink how we design for user agency, trust, and control in immersive environments, where data collection is pervasive, but not always transparent. Traditional privacy mechanisms, shaped by web and mobile interfaces, rely heavily on static, text-based policies and preemptive user agreements. Yet, prior research has highlighted many limitations of these approaches, even in non-immersive settings. Users are often unaware of their current privacy permission states [68], important information is inaccessible, buried in pages of legal text [61], and many users have only a limited understanding of how their privacy-related decisions impact their subsequent experience with using an application [80]. These limitations are amplified in XR, where interactions are spatial, embodied, adaptive, and deeply contextual. For example, recent work such as *GAZEploit*, demonstrates how leaked gaze information can be exploited for remote keystroke inference. By capturing and analyzing the virtual avatar's viewpoint, an attacker can estimate a user's gaze trajectory and reconstruct sensitive inputs such as passwords, emails, and messages [85]. As a result, the stakes for unintentional disclosures are higher, the complexity of privacy management increases, and the cognitive burden on users becomes significantly greater.

To address this, we propose a *Multi-Layered Privacy Framework* comprising five interdependent layers spanning from governance to user cognition (see Figure 1). These layers are: (0) Ethics & Governance, (1) System & Data Foundations, (2) Permission Models & Control Granularity, (3) User Interaction & Experience, and (4) User Perception & Cognitive Load. Our framework provides a structured way to integrate high-level principles with low-level implementation, ensuring privacy considerations are embedded throughout XR development – not retrofitted afterwards. It offers developers, researchers, and policymakers a structured lens through which to

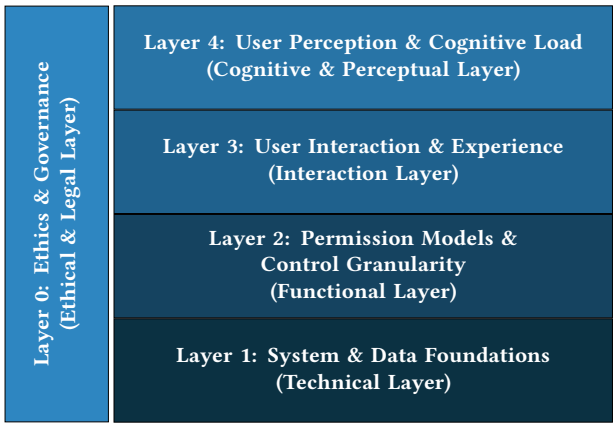

Figure 1: Multi-Layered Privacy Framework

design and evaluate privacy mechanisms in XR. Importantly, we present the framework not as a fixed or final solution, but as a living artifact intended to evolve alongside the XR landscape.

To operationalize this layered perspective, we introduce the *User Privacy Journey Model*, which maps the sequential stages at which users encounter and respond to privacy cues, from onboarding and contextual permission prompts to in-experience feedback and post-session reflection. This model reframes XR privacy not as a static contract, but as an experiential process that unfolds over time, supporting secure, sustained, and contextualized interaction design. Lastly, to support practical adoption, we offer the *XR Privacy Checklist*, a concise, layer-by-layer tool to evaluate and implement privacy-by-design practices throughout the life cycle of the XR app.

2 Paradigm

In line with NSPW's mission to rethink foundational assumptions in security and privacy, we propose a paradigm shift in how privacy is conceptualized and operationalized in XR environments, away from a one-time static checkpoint towards **continuous, in-context consent**, echoing Morton and Sasse's argument that "privacy is a process, not a PET" [57]. In XR, where data collection is continuous and interaction is embodied, privacy cannot be reduced to a single decision upfront. Instead, our *Multi-Layered Privacy Framework* and *User Privacy Journey Model* remodel privacy as a renegotiable lived experience that unfolds over time and evolves with user needs and understanding. We provide actionable tools to facilitate this transition from reactive compliance to proactive design: rather than patching privacy into XR systems after the fact, we embed it throughout the lifecycle of the user experience. In doing so, our work offers a forward-looking agenda for privacy-by-design in XR, opening new technical, ethical, and UX/UI research perspectives.

3 Background and Research Problem

The research community has acknowledged the privacy challenges arising from the unprecedented scale and sensitivity of data collection in immersive environments, including biometric signals, behavioural patterns, and contextual cues [58].

In response, prior work has explored concrete technical solutions [51], interface designs [65], and proposed ethical guidelines and frameworks [2, 29]. For example, the XRSI framework [29] provides valuable taxonomies and compliance guidance for developers and organizations, outlining what data should be protected and delineating responsibility across various stakeholders. However, these approaches fall short in addressing how users engage with privacy decisions over time.

Morton and Sasse [57] introduced the *Privacy–Security–Trust (PST) Framework*, conceptualizing privacy as an ongoing organizational process grounded in ethics, culture, and information management. Our work extends this perspective to the experiential layer of immersive technologies. While PST primarily examined how institutions operationalize privacy as part of corporate practice, our framework reinterprets privacy-as-process through the lens of the user's lived, embodied experience in XR.

In research ethics, Klykken [36] similarly argues that consent should be treated not as a static, pre-fieldwork agreement, but as an evolving process responsive to participant understanding over time. We extend this critique to XR: users are often asked to consent to pre-use terms before engaging with actual data practices in the experience. In such contexts, meaningful consent cannot be fully realized upfront.

Despite a recognition of these issues, a lack of a structured, end-to-end model that systematically integrates the full spectrum of factors shaping privacy in XR remains, from system-level data flows to user cognition and experience. The following questions emerge:

- How can privacy in XR be designed as a continuous, user-centered experience rather than a one-off consent event?
- What frameworks can guide the integration of legal, technical, experiential, and cognitive considerations into XR privacy design?
- How can systems support user understanding, trust, and agency across the entire lifecycle of XR interactions?

Our work addresses this gap by offering a complementary, user-centric perspective grounded in human-computer interaction. We use literature synthesis, comparative system analysis, and scenario-based design to structure key aspects of XR privacy into a *Multi-Layered Privacy Framework*, and operationalize it through the *User Privacy Journey Model* and the *XR Privacy Checklist*.

4 The Multi-Layered Privacy Framework

The *Multi-Layered Privacy Framework* structures privacy permissions in XR into interdependent layers, where each layer represents a progressively abstracted level of decision-making, interaction, and perception in the XR privacy ecosystem. Each layer builds upon the previous one, ensuring a clear conceptual separation between system-level mechanics (how data is accessed) and user experience (how permissions are perceived and granted).

4.1 Layer 0: Ethics (Ethical & Governance Layer)

Regulatory Frameworks (GDPR, CCPA) Legal regulations for data privacy, such as the EU's GDPR and California's CCPA, do in general extend to XR platforms, but they introduce unique implementation challenges. GDPR, for example, mandates principles like data minimization, purpose limitation, informed consent, and the

right to access or delete personal data [18]. XR systems, with their rich data streams, must interpret these requirements for new data types. One concrete issue is that XR tracking data can be classified as *biometric data* under GDPR (as it can uniquely identify a person or reveal health information). Jarin et al. [32] note that both GDPR and CCPA define certain behavioral and biometric data (e.g., motion patterns or gaze) as sensitive personal information, thus requiring explicit consent and special handling. This means that an XR app collecting head and hand trajectories might legally need to obtain explicit informed consent and perhaps even offer an opt-out, much as a website needs special consent to collect sensitive categories. However, studies have found a compliance gap: for example, many VR apps have privacy policies that *do not adequately disclose* the collection of sensor data or the inferences that can be drawn [32]. This could put them at risk of non-compliance. Kim [35] argues that the current mode of obtaining consent, usually via long terms of service agreed to on app install or device setup, is not truly "informed" in VR contexts. They call for moving beyond just text-based consent to more dynamic, ongoing consent mechanisms that align with the interactivity of XR environments. Additionally, regulations give users rights like the right to deletion, which in XR entails not just deleting an account but possibly all the telemetry tied to a user. Implementing such rights is technically challenging if data is stored in complex ways (e.g., distributed across a user's device, company servers, cloud services).

Research in the law and tech community is actively discussing these issues. They highlight that compliance is not just a legal formality but also an ethical design task: how to ensure users truly understand and control their data in line with GDPR/CCPA [29, 84]. For example, GDPR's emphasis on transparency could inspire more in-app explanations of what data is being used for, and its consent requirement might push developers to implement that just-in-time, granular permission system we will discuss in Layer 2, rather than blanket consents. We also see emerging industry standards from groups like the XR Safety Initiative (XRSI), which has published an "XR Privacy Framework" that interprets regulations for XR and provides guidelines to developers [29]. Adhering to these frameworks can help build user trust because they signal that a product is following recognized best practices for privacy.

Building and Enforcing Privacy Regulations For recommendations and best practices to end up in final user-facing products, collaboration across multiple actors is needed. Industry (e.g., manufacturers and developers) could lead by example, following a self-regulation approach. While user privacy is mainly for the benefit of the user, it can also serve as a competitive differentiator, allowing companies to position themselves and their products as user-friendly and privacy-conscious. Not only could it make them and their products more attractive for privacy-aware users, clear privacy policies and data handling practices could also lead to users sharing more data (as they can trust that it is handled carefully). However, prior research indicates that self-regulation alone might not be sufficient or even fail [16, 22], particularly when there are economic incentives for companies to collect user data [33]. Therefore, laws and regulations, along with compliance oversight and checks, remain critical for broad adoption. A key challenge already observed in mobile and web environments is a disconnect between formal compliance and usability, where companies technically meet legal requirements on paper but fail to provide a user-friendly experience that actually achieves the intended goal of providing users with meaningful choice (such as cookie banners, which technically comply with GDPR but have poor usability in practice [25]). There is a risk of these issues carrying over into XR, where immersive interactions could further obscure consent mechanisms or embed them within poorly designed interfaces (we discuss deceptive patterns in XR in subsection 4.5). Putting the regulations in place will take time, and crafting them so that they stay applicable over time with new developments in XR will take careful consideration to avoid future gaps between users, technological developments, and legal frameworks. Anticipating new developments, from technical advances to changes in how the technology is used, is crucial for having proactive, forward-looking regulations.

Summary Regulatory considerations are driving the XR industry to rethink consent and data governance from the ground up. While academic and industry literature increasingly calls for privacy-by-design in XR [1, 29, 35, 63, 64, 84], existing implementations often treat legal requirements as static checklists. Here, we reframe compliance as a dynamic design challenge. Rather than treating regulatory norms as endpoints, we extend the view of consent to a *relational model*, acknowledging that consent cannot be preemptively secured but must be *renegotiated* in context, aligned with user understanding and situational use.

4.2 Layer 1: System & Data Foundations (Technical Layer)

Data Collection and Inferences Modern XR systems collect extensive sensor data (headset orientation, hand/controller tracking, eye gaze, etc.), which can reveal sensitive information beyond the immediate usage context. Prior work has shown that seemingly innocuous motion data can be leveraged to uniquely identify users and infer personal attributes. For example, Nair et al. [60] demonstrated that using just head and hand tracking logs, a machine learning model could identify individual VR users with over 94% accuracy among thousands of candidates, highlighting the biometric uniqueness of motion patterns. Similarly, research on combined eye and body tracking data found that even with privacy protections on one data stream, the other unprotected stream could be exploited for *unauthorized user identification* [7]. These findings underscore that raw sensor outputs (e.g., IMU data, eye images) often contain latent personal information (physical traits, behaviors, even emotional or cognitive states) that can be *inferred* without the user's knowledge [49, 55, 56, 67, 69, 70]. Users may not realize, for instance, that subtle head movements can disclose height or gait, or that gaze patterns can reveal interests and emotional responses. This duality between raw and inferred data raises serious privacy concerns: *XR platforms must treat sensor telemetry not just as technical data but as potentially sensitive personal data.*

OS-Managed vs. App-Managed Permissions and Privacy Enforcement An essential question in XR privacy is determining who manages permission UX and enforces privacy; should it be handled at the OS level or delegated to individual apps? Conventional computing environments typically utilize OS-managed permissions (like Android or iOS dialogs), ensuring consistency and

centralized control. In contrast, app-managed permissions often appear as custom prompts or implicit user actions within an app.

Current XR platforms exhibit a mixture of both practices. Standalone XR systems (e.g., Oculus/Meta Quest) tend to implement standardized OS-level permission dialogs for accessing sensors like cameras, microphones, or eye tracking, requiring explicit user consent per app [53]. Additionally, the Meta Quest OS enforces privacy by physically cutting power to external cameras and microphones under certain conditions, effectively providing a reliable hardware-based kill switch [53]. This approach guarantees uniformity, lets users revoke permissions through a centralized settings interface, and supports the Principle of Least Privilege, in which "every program and every user of the system should operate using the least set of privileges necessary to complete the job" [74], by sandboxing apps and restricting default access.

In contrast, PC-based VR platforms, such as SteamVR/OpenVR, historically provided minimal OS-level mediation. For instance, an analysis of SteamVR (a PC-based VR platform) found that *SteamVR apps do not have any runtime permission constraints* preventing them from reading sensors; any app can access tracking data by default [32]. Consequently, users must rely heavily on app developers to be transparent and responsible or manually adjust available settings, which leads to potential inconsistencies and privacy vulnerabilities. Furthermore, the lack of OS-level mediation on some platforms has led to inconsistencies and potential privacy gaps. Jarin et al. [32] note that most VR apps' privacy policies do not clearly disclose VR sensor data usage, and few apps provide optional permissions for this data, implying that, without OS enforcement, users are often in the dark about data practices.

Researchers argue for stronger OS-level privacy frameworks to close these gaps, advocating for a global mediation approach that can manage privacy effectively across multiple applications and interactions, ensuring robust, consistent protection [38]. Nonetheless, even with comprehensive OS-level permissions, VR applications often require additional, finer-grained controls specific to their features, for instance, enabling users to toggle voice chat within the app despite the OS already having granted microphone access, which can then be managed on the application level.

Evidence from current research and industry practices suggests the most effective approach combines both OS-managed core permissions and app-managed, context-specific privacy controls [31, 38, 72]. The operating system establishes a baseline of privacy protection and user awareness by enforcing uniform rules and consent requirements, while individual apps offer more granular, context-specific settings. Careful integration between these two layers is critical to prevent user confusion from redundant prompts or inconsistencies, enhancing security and usability [70].

This approach is reflected in mobile ecosystems: Android and iOS require runtime permission requests for sensitive data such as location, camera, or contacts [5, 6], while applications such as Google Maps [23] and Facebook [19] supplement system permissions with additional in-app privacy options, such as location history tracking and data-sharing preferences beyond the system-level controls.

Summary This layer details the technical foundations of XR privacy, focusing on how sensor data is collected and managed. A key aspect is distinguishing between raw and inferred data, as even basic motion data can reveal personal attributes or uniquely identify users. We also discuss different permission mechanisms: OS-managed vs. app-managed permissions. While some standalone headsets provide OS-level permissions, PC-based XR systems often lack centralized privacy controls. Researchers advocate for a hybrid approach, combining OS-level baseline protections with app-managed granular controls, to balance user flexibility with security and consistency when addressing privacy risks in XR systems.

4.3 Layer 2: Permission Models & Control Granularity (Functional Layer)

Fine-Grained vs. Coarse-Grained Permissions. Due to the myriad of sensors and data types in XR, one debate is how granular permissions should be. Fine-grained permission models would ask users to authorize each specific data type or sensor (e.g., distinguish between head position, hand tracking, eye gaze, voice, etc.), whereas coarse-grained models bundle permissions (e.g., one switch to allow "VR tracking data" all at once). The advantage of fine granularity is greater user control and awareness. Users can say 'yes' to necessary data and 'no' to others. Previous work emphasizes the need for such precision; Roesner et al. [70] stressed that AR/VR platforms should consider more fine-grained access control to mitigate misuse of sensor data, even though this increases interface complexity.

In practice, we see some movement toward finer granularity: Meta's Quest Pro, for instance, separates *eye tracking* into multiple data tiers, keeping raw eye images private (processed on-device) and only sharing abstracted gaze vectors with applications that the user has explicitly permitted [53]. This two-tier model (raw vs. processed) exemplifies a fine-grained approach motivated by differing sensitivity levels. However, excessively fine permissions could overwhelm users with choices.

Coarse-grained permissions simplify decisions (fewer prompts) but at the cost of transparency. A single all-encompassing "allow VR tracking" permission might authorize more than the user realizes (e.g., granting access to body, eye, and voice data together). The trade-off is noted in discussions like the W3C's workshop on permissions for the immersive web, where experts debated whether it's better to *"bundle XR permissions or prompt separately"*, acknowledging that bundling can reduce prompt fatigue but might hide details from users [14]. Achieving a balance between granularity and user-friendliness in permission systems remains a significant challenge in user interface design. Current research trends indicate a preference for more granular permission models, complemented by good UI designs to effectively manage users' cognitive load [27, 62].

"All-at-Once" vs. "Just-in-Time" Permission Requests. Another design choice is when to prompt users for permission requests in XR environments. An "all-at-once (AAO)" model would request all needed permissions upfront, typically during installation or at first launch (e.g., Android versions prior to 6.0 Marshmallow). This approach has the benefit of informing users from the beginning about everything the app will access. Still, it can lead to *privacy fatigue* or users blindly clicking "Allow" on a series of prompts just to start using the app [79].

In contrast, "just-in-time (JIT)" permission models ask for access at the moment it is needed in context. For instance, if an XR game wants to use the microphone, it only asks when the user initiates

voice chat. JIT permissions are contextually relevant and can improve understanding ("I see that the app wants mic access because I tried to talk to friends."), but *in XR they come with the risk of breaking presence if a prompt appears mid-experience*. The immersive nature of XR amplifies the interruption cost of JIT prompts more than on mobile or PC.

The W3C immersive web report explicitly raised the question of AAO vs. JIT consent for XR, without a definite answer, noting that both strategies have pitfalls in terms of user comprehension and comfort [14]. User experience research suggests that JIT tends to align better with user expectations of why a permission is needed (since it is immediately relevant), which can improve informed consent. In general mobile app studies, contextual requests have been shown to increase user attention to permissions compared to one-time installation prompts [80, 81, 86]. However, XR UX considerations may require hybrid strategies that ensure careful prompt timing so that users' sense of presence is not unduly interrupted or broken [27, 51]. Designing JIT prompts that are brief, and perhaps presented as in-headset overlays, seamlessly integrated into the virtual world, is an active area of exploration in HCI [65].

Long-Term Control Privacy is not a one-time set-and-forget choice, especially as XR applications update and user comfort levels change. Therefore, providing users with ongoing controls, the ability to revisit and revoke permissions or change data sharing settings, is crucial for sustaining trust [71, 80].

Current VR systems are beginning to include such features; for example, the Meta Quest has a "Permissions" section per app (accessible in its settings menu) where a user can disable previously granted accesses like camera or microphone for that app [54]. However, the usability of these controls varies, and users may not know they exist unless educated.

Researchers suggested continuous consent mechanisms. One interesting concept is giving users a physical or virtual "kill switch", a control that allows them to instantly revoke access to sensitive permissions. This concept was previously explored in the form of a tangible mechanism for smart home environments [17]. In XR, a kill switch could translate to a *safe word*, a mutually agreed-upon verbal or gestural signal used to immediately revoke ongoing consent, triggering protective system-wide actions such as disabling microphone input, stopping eye-tracking, or halting the collection of biometric data (imagine saying "mute all mics" to your headset, which the OS then enforces). While not yet standard, these ideas resonate with broader HCI findings that users want *empowerment* in controlling their data after initial consent [48]. Long-term transparency also matters: users should be able to find out what data has been collected about them over time. Some proposals in the XR community include personal data logs that the user can review. For instance, an app might provide a log like "In the last month, we accessed your location 5 times and eye tracking 20 times." This is analogous to how some operating systems now show a history of sensor access (e.g., Android's privacy dashboard that shows which apps used the camera in the past seven days [39]).

In XR, making such information available can build trust as users don't feel that data practices are hidden. If a user notices something unexpected in a log, say a game accessing their camera when it is not obvious why, they can question or adjust that. Research by Balebako et al. [8] on privacy feedback indicates that users appreciate this kind of awareness, and it can encourage more mindful permission management (users might revoke access if they see it is overused).

Summary The functional layer of the *Multi-Layered Privacy Framework* explores core trade-offs in designing permission systems for XR: fine-grained versus coarse-grained permissions, all-at-once (AAO) versus just-in-time (JIT) permission requests, and the importance of long-term user control. Fine-grained permissions offer greater transparency and user control by allowing granular data permission, but may overwhelm users with complexity. Conversely, coarse-grained permissions simplify decisions but may obscure the scope of data being shared. Similarly, choosing between AAO and JIT permission requests involves balancing comprehensive initial disclosures against contextual relevance. JIT prompts can provide clearer rationale and relevance, but may disrupt the immersive XR experience if not designed carefully. More research on hybrid approaches and context-sensitive privacy UI design is needed to optimize UX without compromising neither privacy nor presence.

Finally, effective privacy management in XR requires mechanisms for ongoing user empowerment. *Long-term control means both the ability to change one's mind (revocation) and the ability to stay informed and in control as time passes.* These features turn privacy from a single moment (click "Allow") into an ongoing dialogue between the user and the system, promoting user agency.

4.4 Layer 3: User Interaction & Experience (Interaction Layer)

Onboarding and Privacy Tutorials Educating users about privacy in an immersive context is non-trivial, and several researchers have explored novel onboarding strategies [10, 82, 83]. Traditional text-based privacy policies or settings menus are often ineffective [4]. Not only do users tend to ignore long texts *in general*, they are especially unlikely to seek out 2D privacy pages in the middle of an engaging XR experience. This has led to proposals for interactive, in-XR privacy tutorials that teach users about data practices and privacy controls in a more engaging way. For example, Lim et al. [45] present a role-playing privacy tutorial game in VR where the user learns how their personal data might be "mined" and misused by actively experiencing a scenario. The idea of embedding privacy education into VR is also echoed by recent user studies – a qualitative study on VR privacy concerns suggests that interactive tutorials could help users better grasp what data is collected and how to protect themselves [75]. By introducing privacy concepts during onboarding (e.g., when a user first sets up the device or an app), baseline awareness can be raised. There are also simpler onboarding practices, like initial setup wizards that explicitly walk through key privacy settings (e.g., asking the user to choose defaults for data sharing, explaining what each sensor does). While empirical evaluations of privacy onboarding in XR are limited, research on security tutorials, particularly for two-factor authentication (2FA), indicates that well-designed tutorials can significantly enhance users' understanding and adoption of protective features [3, 4]. In summary, XR platforms should not assume users know how to manage privacy on their own; instead, they should be guided through it in an accessible manner through the UI.

Notification Timing, Interruptions, & Immersion A critical UX consideration is how and when to notify users about privacy

events (such as an app requesting access to data, or an indicator for data recording) without unduly harming the immersive experience. XR is often praised for its sense of presence, the feeling of "being there" in the virtual environment, which can be fragile. Poorly-timed modal dialogues or abrupt notifications can pull a user out of that state, an effect sometimes called *breaking presence* [28, 73]. Prior research on XR design highlights that timing and modality of notifications should be carefully calibrated [12, 28, 73]. One approach is to present privacy prompts at natural breakpoints in the experience, such as between game levels, or when the user is already in a menu, rather than in the middle of intense interactions [12, 30]. Researchers have also explored techniques such as diegetic alerts - messages that are part of the virtual world, like a ringing virtual smartphone to indicate a permission request - which can be more seamlessly integrated and less presence-breaking [24, 51].

For JIT permission requests that require immediate attention, adaptive and plausible notifications [87] could be a potential approach to reduce disruption, or by a subtle cue first, e.g., a gentle blinker in the periphery or a brief audio chime, then await user pause to show the full prompt. Evidence from general permission studies shows that if notifications come at inconvenient times, users are likely to ignore them or feel frustrated [14]. Another strategy, as investigated by Chen et al. [12], is using VR-specific affordances such as eye gaze. If the system detects that the user is busy (high motion, focused gaze), it can delay non-critical notifications until they slow down or look at a particular UI element. Ultimately, maintaining presence while maintaining privacy is a delicate balance; therefore, HCI research continues to investigate well-timed, context-aware privacy-related notifications, ensuring that users remain informed while preserving the fluidity of their virtual experience. Moreover, if a permission is not time sensitive, giving users the possibility to temporarily dismiss a consent prompt and "put a pin on it" could introduce an active, revisitable consent mechanism that preserves users' flow while deferring the decision to a moment they see fit. In XR, such requests can be anchored in the environment, for example, as a floating icon or marker attached to a virtual object, allowing users to return to them at their convenience.

Interaction Modality XR devices support a variety of input modalities, including traditional controllers, hand gestures, and voice commands, each with distinct advantages and limitations when used for privacy permissions. Traditional input methods, such as laser pointers and handheld controllers, provide precise selection accuracy but require additional cognitive and physical effort, such as aiming or navigating abstract UI controls, potentially disrupting the user's immersive experience [46, 47].

Emerging interfaces and prototypes increasingly explore natural input modalities like gestures and voice commands, which could offer more intuitive and frictionless interactions. For instance, instead of manually selecting "Allow" on a floating 2D panel, users might consent with a head nod or thumbs up gesture, or verbally deny permissions by simply saying "no, not now." These embodied interactions leverage the physical affordances of XR, allowing users to respond naturally without disengaging from their current task.

Previous research supports the viability of voice-based privacy controls. Malkin et al. [48] examined spoken commands for privacy management in always-listening devices, such as muting a smart speaker with a specific phrase. Similar voice interactions could be implemented in XR contexts, where a user might verbally confirm or deny a prompt following an initial wake-word to prevent accidental inputs. Likewise, advancements in hand tracking now reliably recognize specific gestures, such as pinching the index finger and thumb, providing a quick and subtle confirmation mechanism.

These natural input modalities also introduce specific limitations. Voice commands generally require "always-listening" features, potentially raising user concerns about continuous audio monitoring and discomfort with verbalizing privacy choices in shared physical spaces or public settings [48]. In XR scenarios specifically, users might hesitate to express privacy preferences aloud if uncertain about the presence or proximity of others in the real-world [13].

Effective interaction design must also address clear and immediate feedback to confirm that user actions are recognized and recorded [21]. This feedback does not necessarily have to mirror the input modality; spoken commands do not necessarily need to trigger spoken responses. Depending on context, auditory cues (e.g., notification tones), visual indicators (such as color changes or symbolic icons), or even subtle haptic responses may be more appropriate and less disruptive.

Privacy experts emphasize that consent mechanisms should be explicit, intuitive, and easy to learn by users [65]. Clear guidelines and user education play a key role in ensuring that users can confidently exercise their privacy choices without ambiguity.

Privacy Summaries and Dashboards To support the above-mentioned continuous control, the idea of a *privacy dashboard* in XR has gained attention. A privacy dashboard is a centralized UI where users can see and manage their privacy settings and data at a glance. On an OS level, companies like Meta have deployed privacy centers (e.g., the Oculus Privacy Center [53, 54]) where users can view which apps are currently accessing, or have accessed (i.e., summaries) their data and control them.

Offering standardized privacy dashboards, either through the OS or apps, might make them more accessible and routinely used. In VR, this might be implemented as a special virtual room or console that a user can call up, e.g., a VR home environment where one wall is a "control panel" showing icons for camera, microphone, eye tracking, etc., with indicators if they are currently active, which app is using them, and a summary of past collected data. The user could point at any icon to get more information or to revoke access.

Summaries can also provide users with an overall "privacy check" of how their data was accessed or used by different apps. Since users might not constantly check dashboards, providing periodic summaries or notifications about data use can keep them informed, as shown by Balebako et al. [8]. For example, a monthly pop-up in VR could display: "Here's your privacy report: 3 apps used your voice chat this week, your fitness app tracked 5 hours of exercise data," etc., along with prompts to review settings if anything appears unusual. Such summaries should be designed not to alarm but to maintain awareness. Researchers have found that such mechanisms can gently nudge users to reflect on their settings; however, they might not necessarily lead to behavior change [8, 20]. More research is needed to optimize their format, appearance, frequency, and delivery to make them more engaging and actionable, encouraging meaningful user interaction without causing notification fatigue.

Bystander Privacy and Shared Awareness Beyond the active user, data capture of XR devices raises parallel concerns for bystanders — individuals who may be recorded, tracked, or inferred without actively participating in the experience. This extends privacy beyond a single user–system relationship toward a shared and relational dimension, where systems must communicate awareness and consent between XR users and others sharing the same physical or virtual context. O'Hagan et al. [62] show that bystanders exhibit varying needs for awareness and consent depending on context, social proximity, and purpose. *BystandAR* [15] introduces a real-time, on-device visual data filtering mechanism that detects and anonymizes bystander faces to prevent unintentional capture and storage, operationalizing privacy through technical mediation. Perez et al. [66] propose a wearable system that allows bystanders to signal their facial privacy preferences to nearby XR devices, offering a reciprocal channel for bystander agency.

Awareness cues are equally important in the opposite direction, ensuring that XR users remain cognizant of the presence of others. The *Bystander Awareness Notification Systems (BANS)* [51] demonstrates that subtle, ambient indicators (such as outlines or proximity-based overlays representing nearby people) can enhance VR users' awareness of bystanders without breaking immersion. Such cues not only improve safety but also reduce the risk of unintentional privacy violations, paving the way for reciprocal privacy signaling systems, where both users and bystanders remain informed about mutual presence and potential data capture activities.

User-Specific Privacy Profiles XR devices are often shared among multiple users, each with distinct privacy needs and preferences, or a single user may engage with multiple XR devices and platforms. Rather than device-level defaults, user-specific privacy profiles could enable users to store and manage individualized privacy settings. Switching between user privacy profiles rather than reconfiguring privacy settings each time can help reduce repetitive decision-making and maintain consistent protection. Such profiles can also strengthen safeguards for vulnerable user groups, for example, through child accounts managed by parents who predefine appropriate data-sharing boundaries. Additionally, creating a new user privacy profile can serve as a cue for privacy onboarding or contextual tutorial, ensuring that new users understand how their data is collected, used, and protected.

Summary Effective user interaction and experience design are crucial to address privacy concerns in XR. Onboarding tutorials can improve user awareness and understanding of complex privacy concepts, while context-aware, well-timed privacy notifications are critical for maintaining a sense of presence in the virtual environment while managing privacy decisions. Different input modalities, such as controllers, gestures, and voice commands, impact how users manage permissions, with natural interactions potentially reducing friction. Additionally, privacy dashboards, periodic summaries, and user-specific privacy profiles can help maintain awareness over time, allowing users to track data usage, adjust settings as needed, and ensure that personalized privacy settings persist across devices and sessions. However, optimizing the timing and modality of privacy notifications remains a challenge, as it requires balancing presence with privacy awareness. Furthermore, the usability of privacy dashboards and the impact of periodic summaries on user behavior require further evaluation. Equally important, XR interfaces must consider bystander privacy, i.e., designing interaction cues and signaling systems that communicate mutual awareness and consent between users and non-participants, ensuring privacy protection extends beyond the active user.

4.5 Layer 4: User Perception & Cognitive Load (Cognitive & Perceptual Layer)

Privacy Awareness and Mental Models Users' understanding of what data an XR system collects and how it is used, their mental model, greatly affects their privacy behavior. Studies focusing on VR and AR have found that privacy awareness is often limited. Hadan et al. [27] conducted a large-scale survey of XR users and discovered that many were unaware of the full range of sensors and data being captured by their devices, especially *invisible* data like eye gaze patterns or physiological responses. Participants in that study were often surprised to learn about the sensitivity of certain data streams. Similarly, an interview-based study by Adams et al. [2] revealed that early VR adopters mostly thought about obvious data (such as username, profile info, or camera feed) but did not consider less visible data such as motion logs or biometric inferences; their primary concerns were often physical safety or comfort rather than information privacy. Recent qualitative work confirms that there is a spectrum of understanding: some users assume the headset is only doing "tracking" in a transient way and do not realize it can record or transmit those movements, whereas others, often more technically savvy, have concerns about how such data could be stored or shared [75].

Users may have misconceptions about their privacy in virtual environments. For example, a user might assume that their movements are private and known only to them, when in reality, an XR application may transmit this data to cloud servers. These gaps in mental models mean that users might not seek out protections that are actually needed. To address this, researchers have suggested improving transparency and incorporating privacy hints into the UI to continually inform users. Designing interfaces that gently educate and correct false assumptions without unnecessarily scaring users is a goal derived from these studies [27, 65, 75].

Overall, the literature highlights a need to boost privacy awareness in XR because, without it, users cannot make informed decisions or effectively use the privacy features provided.

Impact of Presence on Privacy Decision-Making. Presence, the subjective feeling of "being there" in the virtual environment, can influence how users perceive and act on privacy choices. High presence can potentially distract users from outside concerns; for example, a user deeply engaged in a VR game might be more likely to quickly click "Allow" on a permission prompt just to maintain the flow of the experience, without fully considering the implications.

Research suggests that users' priorities shift in immersive contexts. Adams et al. [2] noted that while privacy was acknowledged, it was often secondary to immediate experiential factors like physical safety (e.g., not tripping over real objects) and enjoyment. As a result, privacy-related decisions, such as granting data access, might not get the careful deliberation they would if the user were in a non-immersive setting.

In social XR, however, presence can amplify feelings of vulnerability or trust. For example, users in social VR settings express concerns about unintended exposure or surveillance due to the extensive data collection capabilities of XR devices, which can heighten their privacy concerns despite the virtual context [27]. These concerns are not limited to system-level data handling but extend to interpersonal dynamics within virtual spaces. In a recent study, users of social VR applications emphasized the need for meaningful consent mechanisms not only with platforms, but also with other participants, especially in scenarios involving body tracking, voice, and gaze interactions [77].

Li et al. [40] explored users' perceptions of security in VR and found that users wanted security measures and authentication actions to be seamlessly integrated — or even entertaining — within the virtual experience. This suggests that privacy prompts perceived as overly "technical" or external to the immersive narrative may be more easily ignored, dismissed, or resented.

A similar dynamic appears in research on seamless authentication. Studies have shown that users prefer security measures like authentication to be integrated unobtrusively or even entertainingly into the VR experience, highlighting users' low tolerance for disruptions to immersion [40]. Such seamless authentication often relies on implicit data, including kinetic signatures, movement patterns, and gaze behaviors - data types that can uniquely identify users [41–44, 67]. While these methods enhance usability by eliminating explicit interruptions, they simultaneously blur the line between convenience and privacy invasion.

This blurring of boundaries can also create unique conditions where users are more susceptible to manipulation. Recent studies have shown that XR platforms can take advantage of user sensory immersion to introduce deceptive patterns, i.e., design choices that subtly steer users toward decisions they may not otherwise make if fully informed, mirroring findings in web-based environments, where interface design patterns, such as visual emphasis, pre-selected options, and misleading button labels significantly increase consent rates [9]. In XR, techniques such as urgency signals, limited navigation options, or immersive role-play scenarios can manipulate users into consenting to invasive data collection without deliberate reflection [26]. Krauß et al. [37] identified perception, spatiality, physical/virtual barriers, and XR device sensing, such as eye tracking, body movement, and location, as XR-specific properties that allow deceptive patterns. These features not only support the application of known deceptive patterns but also give rise to novel manipulative strategies specific to XR environments. For example, spatially imbalancing options, such as placing favorable choices physically closer than less favorable ones, or requiring users to perform emotionally or ethically distressing actions, e.g., harming a virtual creature in order to opt out of a service, are emergent patterns that leverage XR's affordances. Adding to this, research found that some XR design practitioners rationalize or normalize such tactics under business or usability pressures, even when they undermine informed consent [89].

This introduces a critical tension: the same immersive mechanisms designed to maintain presence and flow can obscure the user's awareness of continuous data collection. Users immersed in XR environments, experiencing seamless interactions, may not consciously perceive the privacy risks associated with the implicit data being continuously collected. Efforts to maintain uninterrupted experiences through seamless interactions inherently involve privacy trade-offs, necessitating careful consideration of how to balance users' sense of presence with their ability to carefully deliberate on privacy decisions.

Moreover, beyond cognitive and behavioral factors, emotional responses also shape how users perceive privacy in XR. As Seberger et al. [78] highlight, users often *normalize affective discomfort*, a persistent sense of "creepiness" when technology feels intrusive, that may lead them to tolerate invasive data practices despite unease. Recognizing and mitigating such affective normalization is crucial in XR, where continuous sensing and embodiment can amplify these emotional undercurrents.

These cognitive and emotional dynamics point to the need for XR systems to support users in making more deliberate privacy decisions, even while they are immersed. This could involve subtle reminders or privacy notifications that blend naturally into the virtual environment, or designing experiences where critical privacy prompts appear at moments when users can pause, reflect, and reestablish a sense of agency.

Cognitive Load and Understanding Permission Dialogues. The cognitive demands on an XR user are already relatively high. They may be managing spatial navigation, interacting with virtual objects, and processing rich audio-visual stimuli at the same time. Introducing a permission dialogue in this context can add significantly to the cognitive load.

If a prompt is too complex, e.g., lots of text explaining data use, the user may not bother to read or understand it, as their mental resources are focused on the primary task [4]. In traditional interfaces, as highlighted by McDonald and Cranor [52], users often do not read privacy policies or lengthy consent forms, as it would take hours to read all the privacy policies encountered.

In VR, reading paragraphs of text is even more impractical, as display resolution can hinder readability. Kim [35] specifically calls out the futility of text-based informed consent in VR, arguing that new methods are needed because expecting users to page through legalistic text in a headset is unrealistic.

High cognitive load can also arise from the format of the permission request. A study on user perceptions of Internet of Things (IoT) devices by Zheng et al. [90] noted that users favor simple, glanceable information when deciding on privacy trade-offs. This aligns with proposals to use icons or *privacy nutrition labels* [34] instead of verbose descriptions. For XR, one could imagine a concise overlay, e.g., an icon of an eye with an ON/OFF indicator to show an app is accessing eye-tracking, rather than a text box saying "This app is now accessing eye-tracking data." Such approaches reduce the cognitive burden of comprehending the situation.

Another important aspect is memory load. If a user grants permission once (often at first use), they might forget what they consented to over time. For example, an XR platform might request access to certain data during initial setup, but it does not offer reminders later on. As a result, later, you may not remember that the app continues to have access to specific information.

Researchers have discussed the idea of offering information at multiple levels: a brief prompt for immediate action and the option to "dig deeper" for those seeking details [76]. Thus, the cognitively

light path is default, but the system is transparent to those who wish to have more detailed information.

In summary, minimizing cognitive load in permission dialogs means simplifying content using visuals and clear language, integrating the dialog into the XR context to leverage spatial familiarity, and not overloading the user's working memory.

Summary User perception and cognitive load play critical roles in privacy decision-making in XR. Users often have incomplete or incorrect mental models regarding the types and sensitivity of data collected by XR systems, leading to misunderstandings or underestimations of privacy risks. Presence, while essential to the XR experience, might hinder users in making deliberate privacy decisions or create an illusory sense of privacy. Furthermore, the high cognitive demands inherent to XR, such as spatial navigation and interaction with rich stimuli, extend these challenges, making traditional text-heavy privacy notifications ineffective. To mitigate these issues, research advocates for integrating clear, concise, and contextually appropriate privacy indicators directly into the immersive experience. This includes leveraging visual and symbolic cues, minimizing disruptions, and offering multiple layers of information tailored to varying user engagement levels.

5 The User Privacy Journey Model

While prior research has identified various privacy risks in XR and advocated for principles such as privacy-by-design and data minimization [11, 65, 88], existing solutions often treat privacy as a one-off transaction, lacking a structured, user-centered approach that covers all aspects of a user's *privacy journey* in immersive environments. To bridge this gap, we propose the *User Privacy Journey Model*, which operationalizes the *Multi-Layered Privacy Framework* introduced in section 4. This model conceptualizes privacy in XR as a temporal and interactive process, mapping how users encounter and manage privacy permissions throughout the course of an XR experience. It provides a scaffold for supporting user awareness, transparency, and control — not just at the entry point, but throughout the entire user journey.

Figure 2 provides an overview of the stages in this model, where each stage corresponds to a specific phase of user interaction, with privacy considerations and interface elements tailored to that phase. The model consists of an ordered sequence of stages through which the user progresses: **Onboarding & Privacy Education** ($S_{onboard}$), **Contextual Permission Prompts** ($S_{in\_context}$), **In-Experience Transparency & Control** (S_{in_use}), and **Post-Experience Review & Feedback** (S_{post_use}). Next, we describe each of these stages in detail and discuss the design considerations for each.

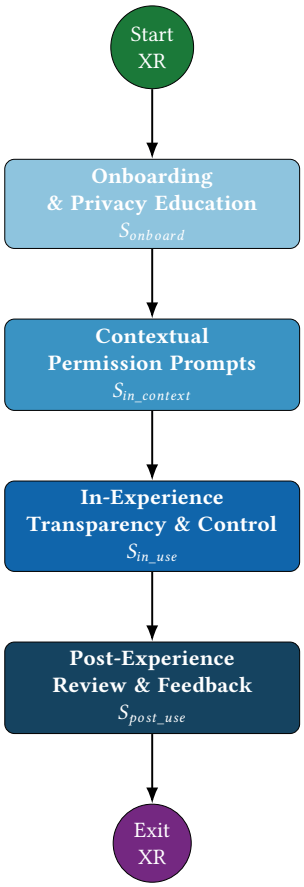

Figure 2: Flowchart of the *User Privacy Journey Model*, illustrating the stages from onboarding to post-experience review.

5.1 Stage 1: Onboarding & Privacy Education

The first stage occurs when the user is preparing to enter an XR application or experience for the first time. At this **Onboarding & Privacy Education** ($S_{onboard}$) stage, the system introduces its privacy-related features and obtains initial consent preferences. Instead of presenting a dense privacy policy text, the model recommends an interactive tutorial or clear, short notices about what data the XR application will access (such as location boundaries, microphone, eye tracking, etc.). Key principles like why each permission is needed are communicated, aligning with the idea of informed consent.

As ongoing research and in-the-wild deployments continue to reveal new privacy risks, onboarding should not be treated as a one-time intervention. Instead, onboarding materials and privacy notices should serve as living artifacts, periodically updated and proactively surfaced to existing users when significant changes occur. This ensures that emerging risks, mitigations, and new data practices are transparently communicated, maintaining the integrity of informed consent over time.

By proactively educating users about data practices (as exemplified by prior work on in-VR privacy tutorials [45]), this stage sets the foundation for trust. The user can also be offered a chance to adjust default privacy settings (for example, choosing a strict, moderate, or relaxed privacy level) before the immersive experience begins. This reflects the privacy-by-design notion of privacy as the default [11]: if the user takes no action, the system should operate in a privacy-protective mode until additional consents are given.

5.2 Stage 2: Contextual Permission Prompts

The second stage focuses on requesting privacy permissions contextually, that is, just-in-time, as the user engages with specific features of the XR experience. Instead of prompting users for all permissions at launch (which users might forget or not fully understand initially), the model uses *contextual* prompts that appear at the moment when a particular feature requires access to sensitive user data.

For example, if a VR game attempts to use the microphone to enable voice chat, a contextual prompt could appear as a pop-up panel or a subtle overlay in the virtual environment, requesting microphone access. The prompt should explicitly state the reason for the request, such as: "This app needs access to your microphone to enable real-time voice chat with other players. Do you want to allow this?" By presenting the request at the moment it becomes relevant, users are more likely to understand the purpose of the data access and make an informed decision in context, rather than relying on abstract permissions granted earlier in the experience.

To preserve the sense of presence, permission prompts should be carefully timed and contextually integrated. One strategy is to trigger prompts during natural pauses in the experience. For example, in a VR game, a request might appear between levels or while transitioning to a new location, when user engagement is momentarily reduced. Alternatively, prompts can be embedded directly into the XR environment in a non-disruptive way. In an AR app, this could take the form of a contextual tag anchored to a real-world object the user is interacting with, such as a smart speaker on a desk. These approaches help maintain the user's sense of presence and continuity while still ensuring that explicit consent is obtained for access to sensitive data.

To avoid repetitive interruptions, the system may remember the user's choices for similar future interactions or group-related permissions into a single request. For instance, one consent action might apply to all avatar body tracking data needed for the session, rather than prompting separately for head tracking, hand tracking, and eye tracking. This aligns with findings from mobile and web interfaces, where batching permissions and providing context improve user acceptance [88]. It is crucial, however, that the user retains both control and understanding over their privacy settings. The model's contextual prompts aim to balance privacy and usability, ensuring sensitive data access is always explicitly authorized, while avoiding overwhelming the user with excessive or poorly timed interruptions.

5.3 Stage 3: In-Experience Transparency & Control

Once permissions are granted and the XR experience is underway, the third stage, **In-Experience Transparency & Control** (S_{in_use}), addresses the need for real-time feedback and user agency within the XR environment. In contrast to traditional platforms, where privacy settings are typically static and managed outside the core experience, XR environments demand in-situ mechanisms that preserve presence while enabling users to remain aware of and responsive to ongoing data collection.

Subtle indicators can provide passive feedback on active data streams. For example, if the microphone or camera is in use, a small icon or ambient highlight in the user's peripheral vision could indicate activity. Likewise, if spatial mapping is underway, a boundary cue or visual overlay might indicate that the system is actively scanning the environment. These cues translate familiar metaphors from mobile operating systems (e.g., sensor icons) into the spatial and embodied context of XR, where screen real estate and user attention operate differently. Importantly, these signals must be designed to align with the aesthetics and logic of the virtual world to avoid breaking presence.

Beyond passive indicators, users should have access to in-experience on-demand controls that allow them to modify privacy settings without exiting the virtual environment. In XR, leaving the immersive context to adjust settings — such as pausing a game to access a 2D menu — can be cognitively disruptive and cumbersome. Instead, users might invoke a privacy dashboard via a gesture, gaze, or voice command, enabling quick access to permissions such as muting the microphone, disabling eye or body tracking, or temporarily suspending data logging. This aligns with user expectations around agency and supports the need for immediate remediation if a user feels uncomfortable [50].

This stage reinforces the *principle of user control*, where privacy decisions are reversible and user-driven. Granting a privacy permission should not be a one-time irreversible action; users must retain the ability to update their choices as contexts evolve. In XR, where the system may track body movements, eye gaze, or even infer emotional states, preserving this kind of user autonomy is imperative for building trustworthy XR applications.

5.4 Stage 4: Post-Experience Review & Feedback

The final stage of the journey, **Post-Experience Review & Feedback** (S_{post_use}), occurs after or upon ending the XR session. In this stage, we introduce the concept of a *privacy reflection prompt*, where the system presents users with a summary of what data was accessed, how it was used, and by which components. This moment of reflection is not merely a log, but an opportunity for the user to re-evaluate and adjust their future privacy preferences, transforming privacy from a one-time decision to a *dynamic feedback loop*. This information can be presented to the user directly within XR, for instance, through a post-session menu or dashboard, or through a companion mobile or desktop app linked to the XR platform. To foster user motivation to engage with these prompts, the system needs to clearly communicate the benefits upfront, e.g., by presenting the reflection as an opportunity to gain more control over their personal data and tailor future experiences to their preferences, rather the simply responding to system-driven requests.

The goal of this stage is to provide users with clarity and closure. For example, a message might state: "This session accessed: room boundary data, microphone (voice chat), and hand tracking. All data was processed locally and not stored on external servers." To enhance transparency, this summary could be accompanied by a simple visualization, such as icons or charts, representing the types and frequency of data accessed, giving users more contextual insight into their digital footprint. Alongside the summary, users could be offered options to clear cached data, adjust specific permissions for future sessions, or revisit the privacy policy. This kind of post-session feedback not only meets transparency and accountability

expectations rooted in legal frameworks like GDPR and CCPA, but also puts the principle of data minimization into practice by prompting users to critically reflect on the necessity of collected data [88].

Importantly, this stage can also serve as a touchpoint for experiential reflection, i.e., a time when users can evaluate their comfort level with the privacy dynamics of the session. A simple prompt such as "How did you feel about the privacy aspects of this session? [Good / Concerned / Tell me more]" can capture user sentiment and signal whether the privacy experience aligned with their expectations. By this stage, users have had a chance to experience the application with a specific set of permissions, giving them a more informed basis for evaluating what data access feels acceptable or intrusive. Collecting such feedback (with user consent) can inform ongoing improvements to privacy UX/UI, contributing to systems that better align with user comfort, expectations, and understanding. Moreover, this data could also be used to adapt future onboarding processes, for example, by presenting alternative onboarding scenarios or walkthroughs tailored to the types of interactions the user previously found intrusive or unclear, thereby helping them better understand how and why specific data is used in those contexts.

Reminders about available controls (e.g., "You can change what data this app accesses at any time in Privacy Settings") can further reinforce a sense of agency and informed choice. Ultimately, by closing the loop with reflection and review, this stage transforms privacy from a one-time agreement into a lived, ongoing relationship, reinforcing trust and accountability as central to the XR experience.

5.5 Mapping the Journey to the Framework

While the previous sections detailed the four sequential stages of the *User Privacy Journey Model*, each stage also corresponds to specific layers within the *Multi-Layered Privacy Framework* introduced earlier. Figure 3 illustrates this alignment, showing how ethical, technical, functional, interactional, and cognitive considerations are embedded throughout the user's privacy journey. This mapping reinforces the idea that privacy in XR is not only sequential but also structurally layered, requiring coordination across multiple levels of design and implementation.

6 Using the Multi-Layered Privacy Framework and the User Privacy Journey Model

Grounded in the *Multi-Layered Privacy Framework* and the *User Privacy Journey Model*, we present a structured privacy checklist to support implementation across different stages of XR development. This is followed by a detailed example scenario situated in the context of a social XR application.

6.1 The XR Privacy Checklist

The *XR Privacy Checklist* synthesizes each framework layer into a structured set of questions, designed to assist developers, researchers, and policymakers in assessing privacy readiness and informing design decisions at both strategic and technical levels.

Layer 0: Ethics & Governance (Ethical & Legal Layer)

- ☐ Does the XR system comply with relevant data protection regulations (e.g., GDPR, CCPA)?
- ☐ Have the ethical implications of data collection been thoroughly considered?

Layer 1: System & Data Foundations (Technical Layer)

- ☐ Have all sensor data sources been identified and documented? How are raw/inferred data managed and protected?
- ☐ Are OS level/app-specific permissions clearly outlined and managed appropriately?
- ☐ Is data storage compliant, secure, and clearly structured?

Layer 2: Permission Models & Control Granularity (Functional Layer)

- ☐ Does the permission offer an appropriate choice of granularity (fine vs. coarse)?
- ☐ Have contextual ("just-in-time") permissions been effectively integrated where appropriate?
- ☐ Can users easily modify or revoke permissions after initial consent?

Layer 3: User Interaction & Experience (Interaction Layer)

- ☐ Does the system support user-specific privacy profiles to manage privacy preferences across different users, devices, and platforms?
- ☐ Have users been educated effectively about privacy features at first launch (onboarding)?
- ☐ Are privacy notifications carefully timed to avoid breaking presence? Has the appropriate modality been chosen (e.g., diegetic alerts, subtle cues)?
- ☐ Can the user easily obtain an overview of all active permissions and ongoing data collection during the experience (in-experience transparency)?
- ☐ Does the system include mechanisms to detect and notify non-users (bystanders) who may be recorded or tracked, and, where appropriate, obtain consent from them?
- ☐ Is post-session feedback provided to help users review data collection and adjust privacy decisions for future sessions (post-experience review)?

Layer 4: User Perception & Cognitive Load (Cognitive & Perceptual Layer)

- ☐ Are users assisted in developing accurate mental models about data collection and use?
- ☐ Do privacy notices match user expectations and cognitive capacities?
- ☐ Are privacy interactions reviewed for deceptive or manipulative consent-steering patterns?

Cross-Layer Integration and Monitoring

- ☐ Has a scenario-based privacy walkthrough been performed (e.g., using the *User Privacy Journey Model*) to assess privacy interactions across a typical user journey within the XR app?
- ☐ Is there regular monitoring and updating of privacy practices in response to user feedback, regulatory changes, or technological developments?

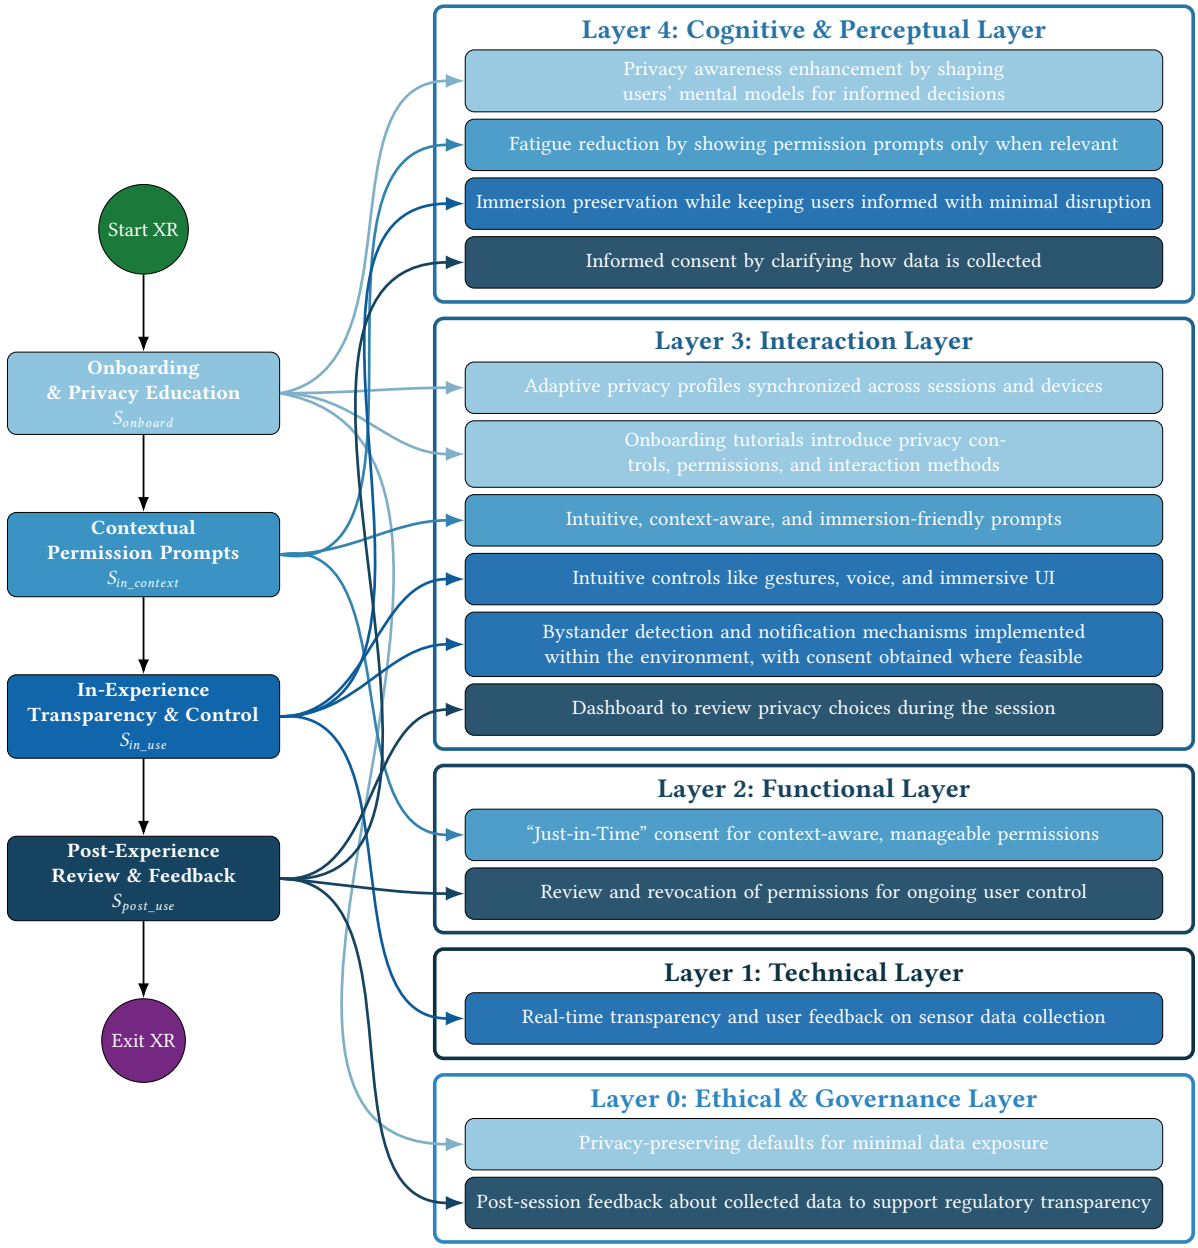

Figure 3: The diagram illustrates how each stage of the *User Privacy Journey Model* maps to corresponding layers of the *Multi-Layered Privacy Framework*, highlighting how ethical, technical, functional, interactional, and cognitive and perceptual considerations are integrated throughout the user's privacy journey in XR.

6.2 Example Scenario

A development team is preparing to launch *GatherXR*, a cross-platform social XR application where users create avatars, join virtual rooms, and communicate via spatial voice chat and hand gestures. Throughout the development process, the team continuously applies the *XR Privacy Checklist* to guide design decisions, ensure compliance, and iteratively address privacy considerations at each stage of development.

Layer 0 The team confirms that *GatherXR* complies with GDPR and CCPA by mapping out all data flows, ensuring explicit consent is obtained for sensitive user data. Internal documentation outlines ethical considerations around inferred data, such as behavioral profiling, and includes mitigation strategies for risks related to bias, discrimination, or exclusion.

Layer 1 Developers inventory all data sources, including gaze, hand gestures, voice input, and movement patterns. They implement OS-level permission dialogues for hardware access (camera, mic), and build app-level controls for toggling features like voice chat or gesture sharing. Data storage is sandboxed per user, encrypted at rest, and structured for efficient deletion on request.

Layer 2 Instead of prompting users for all permissions at login, *GatherXR* adopts a just-in-time model: for instance, users are asked to grant microphone access only when they enter a voice-enabled room. Permissions are structured into fine-grained controls, such as voice input for communication (e.g., speaking to other users), voice logging (e.g., recording for playback or moderation), and voice sentiment analysis (e.g., detecting emotional tone to animate avatars), allowing privacy-conscious users to tailor their preferences. Revocation controls are embedded into the user dashboard, allowing permissions to be changed mid-session or globally reset at any time.

Layer 3 An interactive onboarding tutorial in the virtual lobby teaches users about privacy settings via animated, scenario-based tutorials. Notifications for active data collection, such as gesture tracking, are presented as diegetic in-world cues, such as a glowing avatar's hands indicating motion capture is active. Users can access a floating control panel via hand gesture to view and modify permissions without leaving the virtual experience. After each session, users receive a brief post-experience summary (e.g., visual log of accessed sensors and data collection) and options to modify settings before their next use.

Layer 4 Icon-based privacy indicators and plain-language explanations are used instead of dense text to reduce cognitive burden during consent moments. Post-session, a privacy summary report presents users with a snapshot of what data was accessed to support building more accurate mental models of in-app data collection.

Cross-Layer Integration and Monitoring As part of the rollout process, the *GatherXR* team conducts a full privacy walkthrough by simulating a typical user journey: a player joins a public game room, enables voice chat, interacts using hand gestures, and later explores a private breakout space. The team verifies that just-in-time prompts appear at appropriate moments (e.g., mic access is requested when joining voice chat, not before), and that the floating control panel reflects real-time data usage. After the session, a post-experience summary correctly lists which permissions were used and offers quick access to change them. Additionally, the team collects anonymized user feedback on privacy (dis)comfort or confusion through in-app prompts and uses these insights to refine timing, language, and presentation of permissions. A privacy lead is assigned to oversee these reviews and ensure the checklist is consistently applied with each update, helping privacy remain an evolving and integral part of product development.

7 Rethinking Privacy in XR

As XR technologies increasingly blur the boundaries between the digital and physical self, privacy can no longer be treated as a static checkbox or a one-time legal formality. Instead, it must be recognized as a dynamic and context-sensitive process that is continuously shaped throughout the user's interaction with the XR system.

Our proposed *Multi-Layered Privacy Framework* and *User Privacy Journey Model* call for a paradigm shift in how we design for privacy in XR. This shift reframes consent not as a momentary agreement but as a continuous dialogue between user and system — unfolding across various stages of interaction, and shaped by perceptual, technical, legal, and cognitive factors. A key component of this shift is the introduction of a *privacy reflection prompt*, a post-experience summary that informs users of what data was accessed and how it was used, while offering them the opportunity to reflect and adjust their future privacy settings. By transforming post-use interaction into a moment of reflection and control, this mechanism reinforces privacy as an ongoing, user-driven process across the XR lifecycle.

Despite the conceptual and practical contributions of our work, several open questions remain. How can privacy notifications be made seamlessly adaptive to user state and context without compromising presence? What are some effective mechanisms for communicating complex data practices in an embodied environment? How might regulatory frameworks evolve to support continuous, multimodal, contextual consent in immersive technologies?

While our framework promotes continuous, in-context consent as a more adaptive and user-centered approach, it is important to acknowledge that real-world privacy dynamics are often more complex and contested than what design models can fully anticipate. Our example scenario assumes a relatively smooth process of user permission and system response, but in practice, *revocations*, *renegotiations*, and even *user resistance* may introduce friction or expose conflicts between user expectations and system affordances. Hence, longitudinal user studies and real-world deployments of privacy-aware XR systems are needed to practically validate the *Multi-Layered Privacy Framework*, refine the *User Privacy Journey Model* and the *XR Privacy Checklist*, and uncover unforeseen challenges in its practical application. Features like the proposed *safe word* gesture or phrase can offer one pathway for immediate withdrawal, but future frameworks must also account for escalation paths, fallback modes, and recovery mechanisms when trust is broken or boundaries are crossed.

8 Conclusion

While XR introduces unprecedented opportunities for immersive experiences, it also requires a fundamental rethinking of user privacy, not as a static, one-off consent event, but as a continuous and contextual user-centered dialogue.

To support this shift, we introduced the *Multi-Layered Privacy Permission Framework* encompassing five interdependent layers — from ethics and governance to system infrastructure, permission models, user experience design, and perception and cognition — as a structured lens for embedding privacy throughout the XR lifecycle. We operationalized this framework through the *User Privacy Journey Model*, which maps privacy touchpoints across four key stages: onboarding and education, contextual permission prompts, in-experience transparency and control, and post-experience review and feedback, ensuring that consent is not front-loaded or forgotten, but *renegotiated* as user understanding evolves. To support practical application, we further introduced the layer-by-layer *XR Privacy Checklist*, a tool that guides the implementation of privacy-by-design practices across these stages. Together they establish a

new paradigm for privacy in XR, where consent is not a legal checkbox, but a *continuous, context-aware process* that evolves alongside the user's understanding and interaction.

As XR technologies become more embedded in daily life, future research must continue to explore how such models of renegotiated consent can be made usable, trustworthy, and experientially integrated, empowering users to intuitively understand, control, and reflect on their data interactions.

Acknowledgments

This work has received funding from the German Research Foundation (DFG) under grant agreement no. 521584224. We thank Florian Mathis for the valuable discussion on the initial draft of this paper and Sarah Prange for her thoughtful feedback that helped clarify and strengthen our arguments. We are also grateful to the three anonymous reviewers from the NSPW program committee for their constructive input, and to Filipo Sharevski for his guidance throughout the shepherding process and for helping us sharpen and more clearly communicate the key ideas of this work. Lastly, we thank the workshop attendees for their invaluable feedback, and Jennifer Vander Loop for shepherding the post-workshop paper revision.

Acknowledgment of AI Use

GPT-4 was used to paraphrase and reword text, improve writing style, and perform grammar and spelling checks in order to improve the overall text quality.

References

[1] Melvin Abraham, Pejman Saeghe, Mark Mcgill, and Mohamed Khamis. 2022. Implications of XR on Privacy, Security and Behaviour: Insights from Experts. In *Nordic Human-Computer Interaction Conference* (Aarhus, Denmark) *(NordiCHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 30, 12 pages. doi:10.1145/3546155.3546691

[2] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 427–442. https://www.usenix.org/conference/soups2018/presentation/adams

[3] Elham Al Qahtani, Lipsarani Sahoo, Yousra Javed, and Mohamed Shehab. 2022. "Why would Someone Hack Me out of Thousands of Students": Video Presenter's Impact on Motivating Users to Adopt 2FA. In *Proceedings of the 27th ACM on Symposium on Access Control Models and Technologies* (New York, NY, USA) *(SACMAT '22)*. Association for Computing Machinery, New York, NY, USA, 139–150. doi:10.1145/3532105.3535013

[4] Yusuf Albayram, John Liu, and Stivi Cangonj. 2021. Comparing the Effectiveness of Text-based and Video-based Delivery in Motivating Users to Adopt a Password Manager. In *Proceedings of the 2021 European Symposium on Usable Security* (Karlsruhe, Germany) *(EuroUSEC '21)*. Association for Computing Machinery, New York, NY, USA, 89–104. doi:10.1145/3481357.3481519

[5] Android Developers. 2024. Request App Permissions. https://developer.android.com/guide/topics/permissions/overview Accessed: 2025-04-25.

[6] Apple Inc. 2024. Privacy Control. https://www.apple.com/privacy/control/ Accessed: 2025-04-25.

[7] Samantha Aziz and Oleg Komogortsev. 2025. Exploring the Uncoordinated Privacy Protections of Eye Tracking and VR Motion Data for Unauthorized User Identification. arXiv:2411.12766 [cs.HC] https://arxiv.org/abs/2411.12766

[8] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. 2013. "Little brothers watching you": raising awareness of data leaks on smartphones. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (Newcastle, United Kingdom) *(SOUPS '13)*. Association for Computing Machinery, New York, NY, USA, Article 12, 11 pages. doi:10.1145/2501604.2501616

[9] Nataliia Bielova, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Hary. 2024. The effect of design patterns on (present and future) cookie consent decisions. In *Proceedings of the 33rd USENIX Conference on Security Symposium* (Philadelphia, PA, USA) *(SEC '24)*. USENIX Association, USA, Article 158, 18 pages.

[10] Lal Bozgeyikli, Andrew Raij, Srinivas Katkoori, and Redwan Alqasemi. 2017. Effects of Instruction Methods on User Experience in Virtual Reality Serious Games. In *Virtual, Augmented and Mixed Reality*, Stephanie Lackey and Jessie Chen (Eds.). Springer International Publishing, Cham, 215–226.

[11] Ann Cavoukian. 2009. *Privacy by Design: The 7 Foundational Principles.* Technical Report. Information & Privacy Commissioner of Ontario, Canada. Tech. Rep..

[12] Kuan-Wen Chen, Yung-Ju Chang, and Liwei Chan. 2022. Predicting Opportune Moments to Deliver Notifications in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 186, 18 pages. doi:10.1145/3491102.3517529

[13] Shreya Chopra and Frank Maurer. 2020. Evaluating User Preferences for Augmented Reality Interactions with the Internet of Things. In *Proceedings of the 2020 International Conference on Advanced Visual Interfaces* (Salerno, Italy) *(AVI '20)*. Association for Computing Machinery, New York, NY, USA, Article 20, 9 pages. doi:10.1145/3399715.3399716

[14] World Wide Web Consortium. 2018. *Report from W3C Workshop on Permissions and User Consent.* Technical Report. World Wide Web Consortium. https://www.w3.org/Privacy/permissions-ws-2018/report.html Accessed: March 09, 2025.

[15] Matthew Corbett, Brendan David-John, Jiacheng Shang, Y. Charlie Hu, and Bo Ji. 2023. BystandAR: Protecting Bystander Visual Data in Augmented Reality Systems. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services* (Helsinki, Finland) *(MobiSys '23)*. Association for Computing Machinery, New York, NY, USA, 370–382. doi:10.1145/3581791.3596830

[16] Michael A Cusumano, Annabelle Gawer, and David B Yoffie. 2021. Can self-regulation save digital platforms? *Industrial and Corporate Change* 30, 5 (2021), 1259–1285. doi:10.1093/icc/dtab052

[17] Sarah Delgado Rodriguez, Sarah Prange, Christina Vergara Ossenberg, Markus Henkel, Florian Alt, and Karola Marky. 2022. PriKey – Investigating Tangible Privacy Control for Smart Home Inhabitants and Visitors. In *Nordic Human-Computer Interaction Conference* (Aarhus, Denmark) *(NordiCHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 74, 13 pages. doi:10.1145/3546155.3546640

[18] European Union. 2016. General Data Protection Regulation. L119 pages. https://gdpr-info.eu Retrieved from https://gdpr-info.eu.

[19] Facebook. 2024. Facebook. https://www.facebook.com Accessed: 2025-04-25.

[20] Florian M. Farke, David G. Balash, Maximilian Golla, Markus Dürmuth, and Adam J. Aviv. 2021. Are Privacy Dashboards Good for End Users? Evaluating User Perceptions and Reactions to Google's My Activity. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Virtual Event, 483–500. https://www.usenix.org/conference/usenixsecurity21/presentation/farke

[21] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 64, 16 pages. doi:10.1145/3411764.3445148

[22] Robert Gellman and Pam Dixon. 2011. Many Failures: A Brief History of Privacy Self-Regulation in the United States. In *World Privacy Forum*. World Privacy Forum, 1–29.

[23] Google Maps. 2024. Google Maps. https://www.google.com/maps Accessed: 2025-04-25.

[24] Matt Gottsacker, Nahal Norouzi, Kangsoo Kim, Gerd Bruder, and Greg Welch. 2021. Diegetic Representations for Seamless Cross-Reality Interruptions. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, Bari, Italy, 310–319. doi:10.1109/ISMAR52148.2021.00047

[25] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "Okay, whatever": An Evaluation of Cookie Consent Interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 621, 27 pages. doi:10.1145/3491102.3501985

[26] Hilda Hadan, Lydia Choong, Leah Zhang-Kennedy, and Lennart E. Nacke. 2024. Deceived by Immersion: A Systematic Analysis of Deceptive Design in Extended Reality. *ACM Comput. Surv.* 56, 10, Article 250 (May 2024), 25 pages. doi:10.1145/3659945

[27] Hilda Hadan, Derrick M. Wang, Lennart E. Nacke, and Leah Zhang-Kennedy. 2024. Privacy in Immersive Extended Reality: Exploring User Perceptions, Concerns, and Coping Strategies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 784, 24 pages. doi:10.1145/3613904.3642104

[28] Ching-Yu Hsieh, Yi-Shyuan Chiang, Hung-Yu Chiu, and Yung-Ju Chang. 2020. Bridging the Virtual and Real Worlds: A Preliminary Study of Messaging Notifications in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376228

[29] XR Safety Initiative. 2020. The XRSI privacy framework.

[30] Shamsi T. Iqbal and Brian P. Bailey. 2007. Understanding and developing models for detecting and differentiating breakpoints during interactive tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 697–706. doi:10.1145/1240624.1240732

[31] Suman Jana, David Molnar, Alexander Moshchuk, Alan Dunn, Benjamin Livshits, Helen J Wang, and Eyal Ofek. 2013. Enabling {Fine-Grained} permissions for augmented reality applications with recognizers. In *22nd USENIX Security Symposium (USENIX Security 13)*. 415–430.

[32] Ismat Jarin, Yu Duan, Rahmadi Trimananda, Hao Cui, Salma Elmalaki, and Athina Markopoulou. 2024. BehaVR: User Identification Based on VR Sensor Data. arXiv:2308.07304 [cs.HC] https://arxiv.org/abs/2308.07304

[33] Garrett A Johnson, Scott K Shriver, and Shaoyin Du. 2020. Consumer Privacy Choice in Online Advertising: Who Opts Out and at What Cost to Industry? *Marketing Science* 39, 1 (2020), 33–51. doi:10.1287/mksc.2019.1198

[34] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (Mountain View, California, USA) *(SOUPS '09)*. Association for Computing Machinery, New York, NY, USA, Article 4, 12 pages. doi:10.1145/1572532.1572538

[35] Yeji Kim. 2022. Virtual reality data and its privacy regulatory challenges: A call to move beyond text-based informed consent. *Cal. L. Rev.* 110 (2022), 225.

[36] Fride Haram Klykken. 2022. Implementing continuous consent in qualitative research. *Qualitative Research* 22, 5 (2022), 795–810. doi:10.1177/14687941211014366 arXiv:https://doi.org/10.1177/14687941211014366

[37] Veronika Krauß, Pejman Saeghe, Alexander Boden, Mohamed Khamis, Mark McGill, Jan Gugenheimer, and Michael Nebeling. 2024. What Makes XR Dark? Examining Emerging Dark Patterns in Augmented and Virtual Reality through Expert Co-Design. *ACM Trans. Comput.-Hum. Interact.* 31, 3, Article 32 (Aug. 2024), 39 pages. doi:10.1145/3660340

[38] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2018. Towards Security and Privacy for Multi-user Augmented Reality: Foundations with End Users. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 392–408. doi:10.1109/SP.2018.00051

[39] Abner Li. 2024. Android 16 DP1: Privacy dashboard gets 7-day history. https://9to5google.com/2024/11/19/android-16-privacy-dashboard-history/ Accessed: 2025-03-09.

[40] Jingjie Li, Sunpreet Singh Arora, Kassem Fawaz, Younghyun Kim, Can Liu, Sebastian Meiser, Mohsen Minaei, Maliheh Shirvanian, and Kim Wagner. 2023. How Interactions Influence Users' Security Perception of Virtual Reality Authentication? arXiv:2303.11575 [cs.CR] https://arxiv.org/abs/2303.11575

[41] Jonathan Liebers, Mark Abdelaziz, Lukas Mecke, Alia Saad, Jonas Auda, Uwe Gruenefeld, Florian Alt, and Stefan Schneegass. 2021. Understanding User Identification in Virtual Reality Through Behavioral Biometrics and the Effect of Body Normalization. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 517, 11 pages. doi:10.1145/3411764.3445528

[42] Jonathan Liebers, Sascha Brockel, Uwe Gruenefeld, and Stefan Schneegass. 2024. Identifying users by their hand tracking data in augmented and virtual reality. *International Journal of Human–Computer Interaction* 40, 2 (2024), 409–424.

[43] Jonathan Liebers, Christian Burschik, Uwe Gruenefeld, and Stefan Schneegass. 2023. Exploring the Stability of Behavioral Biometrics in Virtual Reality in a Remote Field Study: Towards Implicit and Continuous User Identification through Body Movements. In *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology* (Christchurch, New Zealand) *(VRST '23)*. Association for Computing Machinery, New York, NY, USA, Article 30, 12 pages. doi:10.1145/3611659.3615696

[44] Jonathan Liebers and Stefan Schneegass. 2020. Gaze-based Authentication in Virtual Reality. In *ACM Symposium on Eye Tracking Research and Applications* (Stuttgart, Germany) *(ETRA '20 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 41, 2 pages. doi:10.1145/3379157.3391421

[45] Junsu Lim, Hyeonggeun Yun, Auejin Ham, and Sunjun Kim. 2022. Mine Yourself!: A Role-playing Privacy Tutorial in Virtual Reality Environment. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 375, 7 pages. doi:10.1145/3491101.3519773

[46] Tiffany Luong, Yi Fei Cheng, Max Möbus, Andreas Fender, and Christian Holz. 2023. Controllers or Bare Hands? A Controlled Evaluation of Input Techniques on Interaction Performance and Exertion in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 29, 11 (Nov 2023), 4633–4643. doi:10.1109/TVCG.2023.3320211

[47] Francisco Lopez Luro and Veronica Sundstedt. 2019. A comparative study of eye tracking and hand controller for aiming tasks in virtual reality. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado) *(ETRA '19)*. Association for Computing Machinery, New York, NY, USA, Article 68, 9 pages. doi:10.1145/3317956.3318153

[48] Nathan Malkin, Joe Deatrick, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy Attitudes of Smart Speaker Users. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 250–271. doi:10.2478/popets-2019-0068

[49] Divine Maloney, Guo Freeman, and Andrew Robb. 2021. Social Virtual Reality: Ethical Considerations and Future Directions for An Emerging Research Space. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, Virtual Conference, 271–277. doi:10.1109/VRW52623.2021.00056

[50] Divine Maloney, Samaneh Zamanifard, and Guo Freeman. 2020. Anonymity vs. Familiarity: Self-Disclosure and Privacy in Social Virtual Reality. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology* (Virtual Event, Canada) *(VRST '20)*. Association for Computing Machinery, New York, NY, USA, Article 25, 9 pages. doi:10.1145/3385956.3418967

[51] Shady Mansour, Pascal Knierim, Joseph O'Hagan, Florian Alt, and Florian Mathis. 2023. BANS: Evaluation of Bystander Awareness Notification Systems for Productivity in VR. In *Proceedings of the 2023 Symposium on Usable Security and Privacy (USEC)*. The Internet Society, San Diego, California, USA, 1–19. doi:10.14722/usec.2023.234566

[52] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.

[53] Meta. 2024. Leveling Up the Safety, Privacy, & Convenience of the Meta Quest 3 Family of MR Headsets. https://www.meta.com/en-gb/blog/new-safety-privacy-features-mr-headset-family-friendly Accessed: 2025-03-08.

[54] Meta Quest Help Center. 2025. View and Manage App Permissions on Meta Quest. https://www.meta.com/help/quest/337925588963509/ Accessed: 2025-03-09.

[55] Abraham Hani Mhaidli and Florian Schaub. 2021. Identifying Manipulative Advertising Techniques in XR Through Scenario Construction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 296, 18 pages. doi:10.1145/3411764.3445253

[56] Mark Roman Miller, Fernanda Herrera, Hanseul Jun, James A Landay, and Jeremy N Bailenson. 2020. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports* 10, 1 (2020), 17404.

[57] Anthony Morton and M. Angela Sasse. 2012. Privacy is a process, not a PET: a theory for effective privacy practice. In *Proceedings of the 2012 New Security Paradigms Workshop* (Bertinoro, Italy) *(NSPW '12)*. Association for Computing Machinery, New York, NY, USA, 87–104. doi:10.1145/2413296.2413305

[58] Vivek Nair, Gonzalo Munilla Garrido, Dawn Song, and James O'Brien. 2023. Exploring the privacy risks of adversarial VR game design. *Proceedings on Privacy Enhancing Technologies* (2023).

[59] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2023. Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 895–910. https://www.usenix.org/conference/usenixsecurity23/presentation/nair-identification

[60] Vivek Nair, Wenbo Guo, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2024. Berkeley Open Extended Reality Recordings 2023 (BOXRR-23): 4.7 Million Motion Capture Recordings from 105,000 XR Users. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2239–2246. doi:10.1109/TVCG.2024.3372087

[61] Jonathan A. Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2020), 128–147. doi:10.1080/1369118X.2018.1486870 arXiv:https://doi.org/10.1080/1369118X.2018.1486870

[62] Joseph O'Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. 2023. Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 4, Article 177 (Jan. 2023), 35 pages. doi:10.1145/3569501

[63] Suchismita Pahi and Calli Schroeder. 2023. Extended privacy for extended reality: XR technology has 99 problems and privacy is several of them. *Notre Dame J. on Emerging Tech.* 4 (2023), 1.

[64] Viktorija Paneva, Marvin Strauss, Verena Winterhalter, Stefan Schneegass, and Florian Alt. 2024. Privacy in the Metaverse:. *IEEE Pervasive Computing* 23, 3 (2024), 73–78. doi:10.1109/MPRV.2024.3432953

[65] Viktorija Paneva, Verena Winterhalter, Naga Sai Surya Vamsy Malladi, Marvin Strauss, Stefan Schneegass, and Florian Alt. 2025. Usable Privacy in Virtual Worlds: Design Implications for Data Collection Awareness and Control Interfaces in Virtual Reality. arXiv:2503.10915 [cs.HC] https://arxiv.org/abs/2503.10915

[66] Alfredo J. Perez, Sherali Zeadally, Scott Griffith, Luis Y. Matos Garcia, and Jaouad A. Mouloud. 2020. A User Study of a Wearable System to Enhance Bystanders' Facial Privacy. *IoT* 1, 2 (2020), 198–217. doi:10.3390/iot1020013

[67] Ken Pfeuffer, Matthias J. Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. 2019. Behavioural Biometrics in VR: Identifying People from Body Motion and Relations in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300340

[68] Sarah Prange, Pascal Knierim, Gabriel Knoll, Felix Dietz, Alexander De Luca, and Florian Alt. 2024. I do (not) need that Feature! – Understanding Users' Awareness and Control of Privacy Permissions on Android Smartphones. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*. USENIX Association, Philadelphia, PA. https://www.usenix.net/conference/soups2024/presentation/prange

[69] Sarah Prange, Sven Mayer, Maria-Lena Bittl, Mariam Hassib, and Florian Alt. 2021. Investigating user perceptions towards wearable mobile electromyography. In *IFIP Conference on Human-Computer Interaction*. Springer, 339–360.

[70] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (April 2014), 88–96. doi:10.1145/2580723.2580730

[71] Franziska Roesner, Tadayoshi Kohno, Alexander Moshchuk, Bryan Parno, Helen J. Wang, and Crispin Cowan. 2012. User-Driven Access Control: Rethinking Permission Granting in Modern Operating Systems. In *2012 IEEE Symposium on Security and Privacy*. IEEE Computer Society, San Francisco, CA, USA, 224–238. doi:10.1109/SP.2012.24

[72] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J. Wang. 2014. World-Driven Access Control for Continuous Sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (Scottsdale, Arizona, USA) *(CCS '14)*. Association for Computing Machinery, New York, NY, USA, 1169–1181. doi:10.1145/2660267.2660319

[73] Rufat Rzayev, Sven Mayer, Christian Krauter, and Niels Henze. 2019. Notification in VR: The Effect of Notification Placement, Task and Environment. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play* (Barcelona, Spain) *(CHI PLAY '19)*. Association for Computing Machinery, New York, NY, USA, 199–211. doi:10.1145/3311350.3347190

[74] Jerome H. Saltzer and Michael D. Schroeder. 1975. The protection of information in computer systems. *Proc. IEEE* 63, 9 (1975), 1278–1308. doi:10.1109/PROC.1975.9939

[75] Abhinaya SB, Abhishri Agrawal, Yaxing Yao, Yixin Zou, and Anupam Das. 2025. "What are they gonna do with my data?": Privacy Expectations, Concerns, and Behaviors in Virtual Reality. *Proceedings on Privacy Enhancing Technologies* 2025, 1 (2025), 58–77.

[76] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security* (Ottawa, Canada) *(SOUPS '15)*. USENIX Association, USA, 1–17.

[77] Kelsea Schulenberg, Lingyuan Li, Caitlin Lancaster, Douglas Zytko, and Guo Freeman. 2023. "We Don't Want a Bird Cage, We Want Guardrails": Understanding & Designing for Preventing Interpersonal Harm in Social VR through the Lens of Consent. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 323 (Oct. 2023), 30 pages. doi:10.1145/3610172

[78] John S. Seberger, Irina Shklovski, Emily Swiatek, and Sameer Patil. 2022. Still Creepy After All These Years:The Normalization of Affective Discomfort in App Use. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 159, 19 pages. doi:10.1145/3491102.3502112

[79] Han Shao, Xiang Li, and Guodi Wang. 2022. Are You Tired? I am: Trying to Understand Privacy Fatigue of Social Media Users. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 378, 7 pages. doi:10.1145/3491101.3519778

[80] Mohammad Tahaei, Ruba Abu-Salma, and Awais Rashid. 2023. Stuck in the Permissions With You: Developer & End-User Perspectives on App Permissions & Their Privacy Ramifications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 168, 24 pages. doi:10.1145/3544548.3581060

[81] Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. 2014. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 91–100. doi:10.1145/2556288.2557400

[82] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Björn Hartmann. 2019. TutoriVR: A Video-Based Tutorial System for Design Applications in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300514

[83] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Bjoern Hartmann. 2020. TransceiVR: Bridging Asymmetrical Communication Between VR Users and External Collaborators. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '20)*. Association for Computing Machinery, New York, NY, USA, 182–195. doi:10.1145/3379337.3415827

[84] Joshua Tooker. 2021. *PRIVACY in the era of Constant Reality Capture: Informed Consent in Extended Reality (XR)*. Ph. D. Dissertation. University of Michigan.

[85] Hanqiu Wang, Zihao Zhan, Haoqi Shan, Siqi Dai, Maximilian Panoff, and Shuo Wang. 2024. GAZEploit: Remote Keystroke Inference Attack by Gaze Estimation from Avatar Views in VR/MR Devices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security* (Salt Lake City, UT, USA) *(CCS '24)*. Association for Computing Machinery, New York, NY, USA, 1731–1745. doi:10.1145/3658644.3690285

[86] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android permissions remystified: a field study on contextual integrity. In *Proceedings of the 24th USENIX Conference on Security Symposium* (Washington, D.C.) *(SEC'15)*. USENIX Association, USA, 499–514.

[87] André Zenner, Marco Speicher, Sören Klingner, Donald Degraen, Florian Daiber, and Antonio Krüger. 2018. Immersive Notification Framework: Adaptive & Plausible Notifications in Virtual Reality. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI EA '18)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3170427.3188505

[88] Yuxia Zhan, Yan Meng, Lu Zhou, and Haojin Zhu. 2023. Vetting Privacy Policies in VR: A Data Minimization Principle Perspective. In *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, New York City, NY, USA, 1–2. doi:10.1109/INFOCOMWKSHPS57453.2023.10225937

[89] Leah Zhang-Kennedy, Maxwell Keleher, and Michaela Valiquette. 2024. Navigating the Gray: Design Practitioners' Perceptions Toward the Implementation of Privacy Dark Patterns. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 97 (April 2024), 26 pages. doi:10.1145/3637374

[90] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 200 (Nov. 2018), 20 pages. doi:10.1145/3274469