ITIF | INFORMATION TECHNOLOGY & INNOVATION FOUNDATION

Balancing User Privacy and Innovation in Augmented and Virtual Reality

ELLYSSE DICK | MARCH 2021

AR/VR devices create novel issues for user privacy due to the scope, scale, and sensitivity of the information they collect. To mitigate harms, policymakers should reform the current patchwork regulatory landscape for data privacy, which fails to address some risks while over-regulating in response to others.

KEY TAKEAWAYS

- AR/VR devices collect similar data as other consumer technologies but raise new privacy issues due to the variety of technologies involved, the sensitivity of the information they collect—and because the data is what makes the devices function.
- AR/VR devices collect extensive biometric data, which can identify individuals and infer additional information. This data can create better immersive experiences but also exacerbate privacy risks.
- The immersive nature of AR/VR makes it difficult to mitigate risks by applying existing privacy policies and practices from other digital media. It requires innovative new approaches to transparency, choice, and security.
- The current regulatory landscape for AR/VR consists of a patchwork of state and national policies, which leaves critical gaps for some privacy risks while over-regulating in response to others.
- Regulating AR/VR or individual technologies they use to deliver immersive experiences will leave policy a step behind innovation as the technology evolves. Policymakers should instead regulate based on actual harms tied to user data.
- Policymakers should create an innovation-friendly regulatory environment for user privacy in AR/VR by clarifying, updating, and harmonizing existing rules and introducing comprehensive national privacy legislation.

INTRODUCTION

In an increasingly digital world, the old saying that "your reputation precedes you" may or may not hold true—but some sort of information about you usually does. User data enables dynamic, personalized experiences with technologies from digital communications platforms to smart devices. But without necessary safeguards, widespread collection and processing of this information, especially by less careful or scrupulous organizations, can expose individuals to privacy risks. Devices and applications for augmented and virtual reality (AR/VR)—immersive technologies that enable users to experience digitally rendered content in both physical and virtual space—are a growing part of this ecosystem.

AR/VR includes applications on mobile devices that combine digital elements with images from external-facing cameras; heads-up displays that overlay digital elements on a user's view of the physical world; and headsets that allow users to navigate fully virtual spaces. In order to deliver these experiences, AR/VR devices and applications gather significant amounts of personal data, including information provided by users, information generated by users, and information inferred about users.

AR/VR raises new user privacy considerations for three reasons:

1. AR/VR devices are composed of a number of different information-gathering technologies, each presenting unique privacy risks and mitigation approaches;
2. Much of the information AR/VR devices collect is sensitive data not used in most other consumer technology devices; and
3. This comprehensive information gathering is critical to the core functions of AR/VR devices.[1]

When broken down, AR/VR technologies are essentially a collection of sensors and displays that work in concert to create an immersive experience for the user. To create the illusion of virtual elements in three-dimensional physical space, or even entirely virtual worlds, these technologies require certain basic user-provided information as a starting point, and then a constant stream of new feedback data that users generate while interacting with their virtual environments. This baseline and ongoing feedback information could include biographical and demographic details, location and movement, and biometrics. Advanced functions, such as gaze-tracking and even brain-computer interface (BCI) technologies that interpret neural signals, continue to introduce new consumer data collection practices largely unique to AR/VR devices and applications. Not only might these data streams contain multiple forms of personal, identifying, or otherwise sensitive information, AR/VR devices also might combine this information to reveal or infer additional details about individual users.

Policymakers should address privacy in AR/VR by considering the different types of information these devices collect and establishing appropriate safeguards to protect users against actual harms that may arise from this data collection.

Taken together, the scope and scale of the user data collection necessary to the core functions of AR/VR distinguish these technologies from other consumer devices and applications. Even so, the types of information collected, the privacy risks, and the potential for direct harms in the absence of safeguards mirror those of other digital technologies and connected devices—many of which have already gained widespread consumer adoption. The unique challenges AR/VR technologies present, therefore, arise from the risks of aggregating sensitive information and the challenge of adapting mitigation measures that were designed for other consumer technologies into immersive, three-dimensional environments.

Because of the wide range of information AR/VR devices collect, policy responses that approach AR/VR as a monolith will almost certainly result in overregulation of certain types of data collection, while also leaving critical gaps in protections for others. At the same time, regulating the individual technologies that are used to deliver immersive experiences will leave policy a step behind innovation as new capabilities and use cases continue to emerge. Instead, policymakers should address privacy in AR/VR by considering the different types of information these devices collect and establishing appropriate safeguards to protect users against actual harms that may arise from this data collection. The goal should be to ensure a comprehensive and technology-neutral regulatory framework that allows space for companies building AR/VR devices to continue to innovate, while mitigating harms. Specifically, this report proposes:

- Relevant federal regulatory bodies should provide guidance and clarification on the ways existing laws, such as the Health Information Portability and Accountability Act (HIPAA) and the Children's Online Privacy Protection Act (COPPA), apply to AR/VR devices and applications;
- Congress should reform privacy laws, such as COPPA and HIPAA, that would unnecessarily limit the use of AR/VR technologies in certain sectors or by certain users;
- Congress and relevant rulemaking bodies should create rules to safeguard against the potential for harm that arises from new forms of data collection, such as biometric identification and personal information inferred from biometric data, through transparency and choice requirements;
- Lawmakers should enact federal privacy legislation to harmonize compliance requirements at the national level rather than rely on state-by-state and sector-specific regulations; and
- Government agencies and industry should develop voluntary guidelines for AR/VR developers to secure users' privacy through transparency and disclosure practices, user privacy controls (including opt-out mechanics), information security standards, and considerations for the unique risks presented by biometric identifying and biometrically derived data.

This report provides a foundational overview of user data collection in AR/VR as it relates to the broader landscape of information-gathering and privacy protections in digital technologies. It reviews the four types of personal data these technologies gather (observable, observed, computed, and associated), the AR/VR data collection practices that fall within these categories, and the privacy concerns and established mitigation approaches for each data type. It then considers the unique challenges immersive technologies present to user privacy protections beyond those present in more established digital technologies, including the role of biometric data, limits to established mitigation approaches, and the potential for vulnerable users to experience exacerbated harms. Finally, it examines the existing regulatory framework for user privacy, identifying laws and regulations that apply to AR/VR as well as policy gaps, and it concludes with recommendations to address the unique challenges AR/VR technologies present to user privacy.

USER INFORMATION COLLECTED IN AR/VR

AR/VR devices rely on information from multiple sources to deliver an optimal user experience and achieve functions other consumer devices cannot. In AR/VR and other information-driven technologies, user information collection can be broadly categorized as one of four types of data:

- **Observable:** information about an individual that AR/VR technologies as well as other third parties can both observe and replicate, such as digital media the individual produces or their digital communications;
- **Observed:** information an individual provides or generates, which third parties can observe but not replicate, such as biographical information or location data;
- **Computed:** new information AR/VR technologies infer by manipulating observable and observed data, such as biometric identification or advertising profiles; and
- **Associated:** information that, on its own, does not provide descriptive details about an individual, such as a username or IP address.[2]

In some instances, particularly in complex technologies such as AR/VR, certain information could contribute to multiple data types depending on how it is collected and processed. For example, baseline health and fitness measurements (e.g., heart rate) are observed data, but calculated health information (e.g., estimated calories burned during an activity) is computed.

Each type of data contributes to the construction of immersive, interactive virtual spaces and objects in different ways, presenting unique privacy considerations and thus a need for best practices to mitigate new and exacerbated privacy concerns. (See table 1.)

Observable Data

Some information can be consistently and directly observed by third parties. With this observable data, other individuals can perceive the same information about the user firsthand. When considering digital privacy concerns, this could include personal correspondence, media shared by the user, or media recorded by third parties.[3] AR/VR devices use observable data to enable users to construct a virtual presence, whether in fully virtual spaces created in VR or physical spaces enhanced with virtual elements through AR.

Observable Data in AR/VR

A user's avatar, or virtual representation of themselves, may be considered observable personal information, particularly if that avatar is a hyper-realistic representation. Even less-realistic avatars a user creates to reflect their physical appearance can reveal certain information such as race and gender. Unlike two-dimensional images, such as profile pictures or digital photographs, three-dimensional avatars such as those in fully immersive VR experiences are a digital embodiment of an individual, including their physical appearance, gestures, and mannerisms.[4] Users experience these virtual bodies as they would their own in physical space—making this particular form of observable data more intimate than similar two-dimensional information.[5]

Table 1: Types of data that AR/VR technologies rely on to create user experiences

Data Type	Examples in AR/VR	Utility in AR/VR	Privacy Considerations	Mitigation Approaches
Observable	Virtual personas or likenesses (i.e., avatars); digital communications or messages; real-time in-app/in-world interactions; identifying in-app/in-world assets (e.g., screenshots, recordings, virtual objects)	Generates virtual presence unique to the user and allows them to interact with virtual spaces and objects	User anonymity and autonomy	Disclosure and user consent; user privacy settings; encrypted communications; limits on law enforcement use; laws against personal autonomy privacy violations
Observed	Location and spatial data (e.g., geolocation, lidar); motion/hand/eye tracking; raw inputs from BCI data; user-provided biographical and demographic information (e.g., name, age, interests); linked social media profiles; user-generated behavioral data and activity logs	Creates and enhances immersive experience; positions user in virtual space; enables advanced functions (e.g., interacting with virtual objects, gesture controls, and more realistic avatars)	User anonymity and autonomy; security of sensitive provided information; potential for discriminatory use of provided information by third parties	Disclosure and user consent; access controls; encryption or local storage for certain data; limits on law enforcement use; laws prohibiting discrimination based on certain information
Computed	User profiles (e.g., for recommendations or advertising); biometric identification; biometrically derived information	Improves services and enables advanced functions	Security of sensitive inferred information; potential for discriminatory use of inferred information by third parties	Disclosure and user consent; users able to contest or correct information; encryption or local storage for certain data; laws prohibiting discrimination based on certain information
Associated	Login credentials; contact information; payment information; friend lists; non-identifying virtual assets; device IP address	Associating content and preferences with specific users or devices; identifying devices and allowing for Internet-enabled functions; enhancing services with additional information	Fraud or malicious misuse; harms from combining with other forms of user data	User authentication; disclosure and user consent when combining with other data; laws establishing standards for information security

In addition to virtual representations of the user's physical self, AR/VR devices also collect certain observable data about their social interactions and affiliations in-world (in VR) or in-app (in AR). As with other technologies, certain forms of communication such as instant messages constitute observable data. Video, images, or screenshots that identify an individual participating in certain activities—whether gathered for redistribution purposes (e.g., an event recording) or more malicious purposes (e.g., recording or capturing images of individuals in sensitive spaces without their consent)—are also observable data from AR/VR. Further, because they are fully virtual (and therefore processed and rendered), a user's presence and interactions are also observable data. This could include recordings of in-world or in-app conversations or gatherings by AR/VR providers as well as other individuals.

This observable data is necessary to create interactive experiences in AR/VR, because immersive experiences require a simulated virtual presence. While single-user applications, such as a single-player game or individual productivity application, may not require individuals to create virtual avatars, the multiuser applications that capture the full collaborative and interactive potential of these technologies do. Without these virtual representations, other users would not have a fully immersive experience.

Privacy Concerns From Observable Data

These types of data can present notable privacy risks for users if not properly limited or secured. Most of the privacy concerns from observable data have to do with anonymity and personal autonomy—that is, individuals' ability to control how much, or how little, others are able to identify and observe about them.[6] The limits of when observable data should be considered private are largely subjective and highly contextual: Some users may choose to allow third parties to observe an extensive amount of information, while others may prefer to make available only the minimum observable information necessary to use a service and interact with other users. Similarly, users may feel comfortable sharing observable information with certain groups (e.g., close friends or romantic partners) but keep it private from others (e.g., employers or professional contacts).

Observable avatars, social activities, and in-world assets can reveal information with varying levels of sensitivity. In turn, user comfort levels and preferences relating to observable information will also vary.

It is important to differentiate between these privacy preferences and actual privacy risks for observable data. Privacy risks arise when the collection, recording, and distribution or replication of this data by a third party reveals private or otherwise damaging information. Because observable data can be perceived firsthand by multiple parties, this sensitive information (e.g., intimate photographs or recordings of an individual shared without their knowledge or consent) can expose the subject to significant personal and reputational harm. Similarly, distributing images or recordings that identify an individual at a sensitive event or location (e.g., at a confidential support group or a medical office) can reveal details about their life that were meant to be shared with only a limited group of people.

Observable data can also lead to personal-autonomy harms when it is used to impersonate an individual or manipulate images of them. For example, a malicious actor can use another individual's image to impersonate them on communications platforms. In fully immersive experiences, a malicious actor could even use an individual's likeness to create a fully interactive avatar. This can lead to reputational or emotional damage when this impersonation makes it look like individuals engaged in activities or interactions when they did not. It can also enable fraud and identity theft, exposing individuals to potential economic harm.

The privacy implications of observable data in two-dimensional digital media also exist in the immersive worlds of AR/VR. Observable avatars, social activities, and in-world assets can reveal information with varying levels of sensitivity. In turn, user comfort levels and preferences relating to observable information will also vary. For example, some AR/VR users may prefer to use avatars that reflect their physical appearance, while others may prefer to use avatars that obfuscate their appearance or identity. This risk is particularly salient for vulnerable users such as children, those who use AR/VR devices for health or medical purposes, and those who share particularly sensitive information in AR/VR.

Mitigation Approaches for Privacy Risks From Observable Data

Mitigation approaches for privacy risks from observable data focus on user control of how and when this data is viewed and distributed, securing this data against unauthorized access, and establishing laws and regulations that protect against misuse or involuntary distribution of observable data. For example, AR/VR devices as well as the individual applications that rely on them can mitigate risks by providing users with transparency and choice about how they collect and share observable data, and the ways in which they use the data. Because users have varying preferences in terms of this use, individual privacy settings enable them to restrict third-party access to observable data they believe is sensitive or private. For example, they may choose to select which users can view their photos or other media. In AR/VR, this could include restricting access to virtual assets as well as the ability for third parties to observe, overhear, or record users' in-app/in-world social interactions.

Technical measures can also protect observable information such as digital communications. For example, end-to-end encrypted messaging ensures only the intended recipients are able to view written communications or shared media (although this does not prevent them from then sharing the contents of that message). Technical restrictions can also prevent third parties from capturing information without the user's knowledge or consent, such as by taking screenshots, although there may be limitations to such technical measures. For example, a social platform for a particularly vulnerable group of users (e.g., children), or a dating application, can prevent others from capturing a screenshot of users' profiles, posts, or private communications—although these measures would not prevent someone from taking a photo of this content displayed on their mobile device. Similar controls can also be implemented in AR/VR applications. While they do not secure this information completely, such measures introduce more friction to potentially privacy-violating practices.

Finally, laws and regulation can play a role in addressing malicious misuse of observable data. In the United States, there are a number of laws that prohibit using an individual's observable data for malicious purposes or collecting and sharing this information without their knowledge. This ranges from surveillance and recording regulations—for government and law enforcement as well as businesses and individuals—to laws prohibiting the distribution of intimate photographs or recordings.[7]

Observed Data

Much like observable information, observed data is descriptive information about the user that they either provide or generate. However, unlike observable data, a third party could not replicate the observation that created this information.[8] Such information could include basic biographical information, personal preferences and behavioral data, affiliations and identity traits, geolocation or other metadata, and other user-provided or user-generated information. Observed data is often used to shape a user's experience with a digital product. It is especially crucial to AR/VR products and services, which rely heavily on this information to provide a fully immersive experience to the user.

Observed Data in AR/VR

A significant amount of AR/VR data falls into this category due to the reliance on sensors to replicate physical experiences in virtual spaces. AR/VR applications must be able to position the user in physical space. AR applications need to understand where the user is standing in relation to geographic locations and physical objects in order to display the relevant digital overlays (e.g., placing instructions on a machine or displaying a virtual sign in front of a building). Meanwhile, VR devices and applications need information about a user's location and physical surroundings for their physical safety: In order for applications to alert users when they approach an object or cross preset boundaries, they rely on a VR device's constant awareness of where both the user and any hazardous objects or physical barriers are located. To achieve this spatial awareness, AR/VR devices collect a broad range of observed data about a user's location. This includes information about their position in physical space, such as Global Positioning System (GPS), (inertial measurement unit) IMU, and gyroscope or accelerometer data, as well as information about their surroundings collected through a mobile device's camera or external-facing cameras on a head-mounted or heads-up display, lidar, and other spatial sensors.

In addition to collecting information about a user's position in physical space, AR/VR devices also track certain movements and collect observed biometric data in order to replicate a user's actions in virtual space. This form of observed data is particularly important in VR, where the user is fully immersed in a virtual environment. Because their users lack any tangible landmarks to orient themselves as they would in physical space, fully virtual spaces must reconstruct physical experiences. The sensation of standing on solid ground, touching objects, shifting fields of vision based on head and eye position, moving freely in any direction, and any number of other seemingly mundane human experiences must be translated into a digitally rendered environment. The more precise this reconstruction is, the more immersive the experience for users.

Devices and applications achieve this immersive simulation by collecting observed data about a user in real time. Head-mounted displays and controllers track head and arm movements and replicate them in the virtual environment. Basic headsets can track movement using three degrees of freedom (3DoF), or three types of head rotation (i.e., looking side to side, up and down, left and right). Those that can duplicate movement in six degrees of freedom (6DoF)—that is, the three rotational directions for head movement as well as whole-body movements (i.e., standing up and sitting down, shifting left and right, and moving forwards and backwards)—offer a more realistic reconstruction of a user's actual movements in three-dimensional space.[9] Some devices also use external sensors or cameras to detect and track hand and finger movements, allowing users to interact with their device's virtual interface and virtual objects without using controllers.[10]

Eye-tracking technologies, which use internal cameras to collect observed data such as where a user is looking, changes in their pupil size, and whether their eyes are open or closed, can be used to create even more-realistic immersive experiences.[11] For example, this data allow programs to display more-authentic avatars that reflect the user's actual eye motion and expressions. Gaze-tracking capabilities also allow VR displays to use foveated rendering, which simulates human field of vision by lowering the resolution of displays that would appear in a user's real-world peripheral vision.[12] Not only does this make the experience more realistic and reduce eyestrain, it can also allow developers to render higher-quality images at the focal point and reduce latency—a major contributor to motion sickness in immersive experiences.[13]

Even more advanced than eye tracking are BCI technologies. The term encompasses any use of sensors that measure brain activity to respond and adapt directly to a user's neural signals. This includes external sensors, such as head-mounted electroencephalography (EEG) sensors, as well as neural implants that could emerge in the future. In the context of AR/VR, the consumer BCI technologies developers envision are typically sensors embedded in a headset or other wearable device. While BCI technologies are not yet deployed in mainstream consumer AR/VR devices, they offer exciting possibilities for AR/VR. An immersive service or game could use these signals to adjust experiences to more-precisely meet a user's unique needs in real time.[14] A BCI-enabled wearable AR device, such as smart glasses, could be controlled inconspicuously by the user with barely noticeable gestures. In order to accomplish this, AR/VR devices will have to collect observed raw data from neural activity sensors, either embedded in head-mounted displays or gathered from other wearable sensors.

Finally, AR/VR devices and applications supplement sensor-based observed data with other information to optimize the user experience. User-provided biographical information (e.g., age, gender, affiliations, and interests) allows services to deliver experiences tailored to individual users' needs, while identifying information (e.g., names or linked social media profiles) verifies unique users and further merges their virtual environments with the real world by allowing them to build and expand social networks. In addition to this user-provided information, many AR/VR devices and applications will also collect observed data about in-world or in-app behavior and activities. Information about what a user does with their AR/VR devices or applications, how long they spend on certain activities, and what experiences they seek out and participate in can reveal personal information. For example, a user may join a support group in a social AR/VR app, or attend an event meant for individuals with specific interests or affiliations. Any records of their participation in this kind of activity is observed data.

Not only do these capabilities enable and enhance immersive experiences, they also expand the potential for their use. Studies have shown that eye tracking can aid mental health practitioners in diagnosing certain brain disorders, and including this technology in AR/VR devices opens new opportunities for their use in medical fields.[15] Market researchers can also use sensor-enabled VR to gather data analytics about consumer attention and interaction with products.[16] Meanwhile, device-observed data can bolster the safety and security of AR/VR devices and experiences: Biometric information, such as iris signatures, can be highly effective for user authentication, while biographical or behavioral information can bolster in-app or in-world safety (e.g., by ensuring a user can only access age-appropriate content).[17] As use cases—and user bases—continue to grow and diversify, so too will the potential uses of this observed data.

Privacy Concerns From Observed Data

As with observable data, a primary privacy concern from observed data relates to individuals' anonymity and autonomy. This data can reveal a significant amount of information about users' lives, with varying levels of sensitivity. From biographical or health information to web-browsing or shopping history, the observed data that users provide can directly reveal or easily infer details they may expect to keep private, such as demographic information, where they live, or how they spend their free time. Also, as with observable data, the sensitivity of observed information will vary among users. For example, some users may find sharing geolocation data with friends and family beneficial and low risk, while this same information could endanger particularly vulnerable groups such as children or abuse survivors.[18] Similarly, some users may want to provide directly identifying information, while others may prefer to remain anonymous.

Simply concealing, anonymizing, or restricting collection of this data would drastically reduce the quality of these services—or indeed render them effectively useless—and impede innovation of any technology that might require user-provided information.

For some users, exposing this information could also lead to harmful discrimination. This creates a notable privacy concern for individuals who are vulnerable to discriminatory practices, for example, in employment or access to critical services, due to attributes such as gender, age, race, disability, sexual orientation, and others.[19] Without safeguards in place to protect against such discrimination, these types of observable information can cause significant harm if they become known to others.

Many of these same risks arise from the observed data gathered by AR/VR devices and applications—and due to the volume of information collected, there is even greater variation in its sensitivity and potential to lead to harm. Risks to users will depend on how, where, and for what purpose they use AR/VR technologies. The observed biographical and health data of a patient using AR/VR therapies would likely be considered more sensitive than the same information provided by a user on a VR gaming or fitness platform; and observed information about a user's location or surroundings would be more sensitive in their living room than in a public park or a shopping mall.

Mitigation Approaches for Privacy Risks From Observed Data

Because of the broad range of observed data collected, there is no one-size-fits-all approach to mitigating the privacy concerns from this type of information. Further, this information adds significant value to many digital products, and in many cases is necessary for their core functions. Because of this, simply concealing, anonymizing, or restricting collection of this data would drastically reduce the quality of these services—or indeed render them effectively useless—and impede innovation of any technology that might require user-provided information. It is important to balance privacy concerns with functionality when considering mitigation approaches for observed data.

Transparency, disclosure, and user consent play an important role in mitigating the potential for harms from observed data. When users understand how and why various AR/VR services collect and share their data, they can make informed decisions about the information they choose to share either directly or through opt-out/opt-in options. As with their observable information, users can adjust individual privacy settings based on their own perceptions of what information they would prefer remains private. However, not all observed information can be restricted without degrading or disrupting the service. For example, AR/VR devices require some motion-tracking information in order to replicate physical movements in virtual space, while a search platform may use geolocation data to deliver more-relevant results. In these cases, ensuring users understand how devices and applications utilize their data to provide different services gives them more control over their personal risk.

Clear guidelines for how AR/VR applications store and process data, as well as how, when, and by whom it can be accessed can also mitigate privacy risks from observed data. In many instances, particularly with highly sensitive information (e.g., biometric identifiers or sensitive health information), the AR/VR application may process the data entirely on the local device—and not transfer any data to a third party—or only store data in the cloud when the user fully controls the encryption keys, thereby reducing the risk of misuse. In others, data can be shared or stored externally without being observed by human eyes. The products and services gathering this observed data can establish thresholds for different levels of access and storage, which can also be included in transparency and disclosure measures.

Finally, laws and regulations can protect users against possible harms from misuse of observed data. There are laws in place that prohibit discrimination against certain protected classes, regardless of how the discriminating party discovered details about an individual. For example, various laws in the United States protect individuals from employment discrimination based on race, gender, age, disability, and other attributes.[20] Other laws prohibit discrimination in health insurance eligibility, housing, and other services.[21] Rules are also in place to govern user privacy from government surveillance for law enforcement purposes. Fourth Amendment rights protect users from unlawful searches, and recent caselaw has extended this to include observed data such as location information.[22] These rules reduce risk when people disclose certain observed identifying information while using digital services.

Computed Data

Unlike observable and observed data, computed data is not provided directly by users. Rather, it is the result of manipulating the observable and observed information users generate in order to derive new information.[23] Computed data analyzes information, sometimes from multiple sources, to make inferences or predictions about users. Because of this, it can provide a more complete picture that can be used to tailor products and experiences to individual users. It can also be wrong. Computed information may include biometric identification, advertising profiles composed of various individual activities, or any other information that is inferred or interpreted rather than directly provided. AR/VR devices rely on these computational processes to derive necessary information from—and make use of—the flood of raw data gathered through their various sensors and user inputs.

Computed Data in AR/VR

AR/VR technologies collect a significant amount of observable and observed data, which can then be interpreted to provide more advanced capabilities and tailored experiences. Descriptive information about users, such as demographic information, location, and in-app/in-world behavior or activities can be combined and analyzed to tailor advertisements, recommendations, and content to individuals. For example, an app might infer, correctly or not, that users who play with virtual dogs in a game have an interest in pets, and target ads for pet-sitting services at this group of people. Similarly, this information can be used to generate user-facing analytics, such as estimating calories burned during a workout based on user-provided demographic information and observed information about physical activities.

AR/VR devices can also generate computed data from the various sensors included in these devices. For example, hand-tracking technologies use observed images of a user's hand and machine-learning technologies to estimate important information such as the size, shape, and positioning of users' hands and fingers.[24] Users can also secure their applications and devices with computed biometric identification, such as iris scanning or facial recognition. And in future BCI-enabled devices, computed data will result from the interpretation of neural signals into actionable commands.[25]

Privacy Concerns From Computed Data

Computed data is distinct from observed and observable information in that it is largely intended only for the parties that have produced it—that is, it is neither accessible to nor replicable by third parties. Because of this, the privacy concerns from computed data are less direct, but also more complex. As with other types of information, users may have varying levels of comfort with how their information is compiled and processed; however, the potential for harm arises from how the information is used, rather than simply who can view or access it. The inferences or predictions that comprise computed data can reveal more sensitive or potentially damaging information about a user than the unique observed and observable information used to generate them. For example, observed biometrics can generate additional details about a user's physical traits or health information.[26] Because of the sensitive nature of this biometrically derived information, the security of computed data and the potential for third parties to access it is a notable privacy concern.

Disclosure of such information without a user's consent or knowledge can lead to significant reputational harm or embarrassment when the nature of the inferred information is particularly sensitive or highly personal. Further, there is the potential for direct harm from computed data being used to unfairly deny an individual certain services or opportunities. This includes discrimination in housing, employment, insurance, benefits, and other services based on accurately inferred details about an individual, such as their age, gender, sexual orientation, or health conditions. These harms can also result when inaccurate computed information, such as an incorrect credit score, bars an individual from accessing services they would otherwise qualify for.

These privacy risks from computed data are particularly acute in the context of AR/VR. Not only do these technologies gather a broad range of observable and observed information, including sensitive biometric information, unrestricted use of this information can reveal significant and highly sensitive additional information about a user. This includes inferred data about preferences from involuntary or subconscious movements or reactions, as well as identifying personal, demographic, and health information.[27] Unrestricted gathering and processing of this information could be used to discover highly specific information about individuals. The extensive potential of AR/VR technologies to generate computed data, while beneficial for their use and technological advancement, can also put users at risk of harm if necessary safeguards are not in place.

Mitigation Approaches for Privacy Concerns From Computed Data

Privacy harms from computed data arise largely due to unintended use, unauthorized access, or malicious misuse. Mitigation approaches seek to secure inferred information that has not been directly provided or generated by the user, and offer remedies for these harms. Much like other information types, disclosure and user consent form the foundation of any mitigation approach for computed data. Users should understand what information can be inferred from the data they provide and how it is used. For inferred data that is not essential to the core functions of a device or application, user privacy preferences can allow individuals to opt out of certain data aggregation and computing. When this information is necessary for functionality or quality of services, transparency and disclosure around how and why inferred information is used can ensure users understand the data practices in place. Further, as with observed data, clear guidelines outlining how data is stored and how, when, and by whom it can be accessed protect users from potential personal or reputational harms that could come from unauthorized access.

Beyond these practices, other mitigation approaches can address the actual harms users might experience from computed data, such as discrimination. Because computed data could reveal information users may not wish to disclose, laws against discrimination based on such information are particularly important. However, nondiscrimination laws do not always address the risk of harms from incorrectly inferred information, such as adverse actions taken based on an inaccurate credit score. Allowing users to correct information that is untrue or out of date can further mitigate these risks.

As AR/VR devices become more widely adopted, it is particularly important to consider the mitigation approaches for computed data in uses beyond personal use or entertainment. For example, workplaces may require employees to complete training that utilize this technology, or universities may provide students with headsets for educational enrichment. These activities could generate significant computed data about employees or students in the process. When use is mandatory, users cannot provide meaningful consent to data collection. Instead, limits on how and when third parties (e.g., service providers as well as employers or instructors) can access or use this information may also be necessary.

Associated Data

A final category of user information that may present privacy risks is associated data. Unlike other types of user information, associated data does not provide descriptive information. In other words, associated data in and of itself does not offer any specific details about a user.[28] This includes information such as identification, registration, and account numbers; device and login information; and addresses or other contact information. AR/VR devices utilize and generate associated data which, when combined with other information or used for malicious purposes, can present privacy risks.

Associated Data in AR/VR

Although AR/VR technologies may offer novel use cases and experiences, they are ultimately just a variant of connected device, alongside laptops and internet of things (IoT) devices. As connected devices, whether they be wearable headsets or AR-enabled mobile devices, AR/VR technologies both collect and generate associated data about their users. User-provided associated data includes details such as login information (i.e., usernames and passwords) for applications and services; user contact information, such as email, phone number, and home address; and user payment information, such as credit card and bank account numbers.

Associated data may also be generated through users' in-app or in-world activities. For example, lists of "friends" or other connections on social or multiuser AR/VR platforms are associated data that can reveal information about a user's social connections and activities. Similarly, non-identifying digital media and virtual assets are also associated information. This includes screenshots and recordings of fully or partially virtual spaces by users themselves, which they may choose to share publicly or keep private. For example, a user may capture a photo of a virtual object in their home using a smartphone or heads-up AR device. In addition to AR/VR versions of digital media, fully virtual objects can also constitute associated information. This includes partially or fully virtual spaces that users create, virtual objects or overlays that they share through AR, and objects users can interact with through immersive experiences.

Such information is necessary to associate users with their unique accounts, user preferences, and virtual assets. For example, in order to provide information such as recently used applications, social interactions, and other user-specific details, a device must have a way of recognizing and authenticating that user—typically through a username or email address and password. Similarly, associated payment information is necessary for any device or application that allows users to make purchases, such as of paid content.

AR/VR devices also have associated data, such as registration numbers and IP addresses. Much like a username and password for an individual, this information identifies a device and enables it to conduct functions that require an Internet connection. Such device information is also associated with the user, but taken alone does not directly identify or describe them.

Privacy Concerns From Associated Data

Associated data, when on its own and used only by authorized parties, presents few significant privacy risks to individuals. However, as part of a larger ecosystem of user information, associated data can lead to direct harms. It can be linked with other descriptive information a particular device or service collects to reveal additional details a user may not wish to share. For example, a screen name or user ID, when combined with information that reveals a user's identity, could link an individual to certain activity, such as browsing history. Similarly, in some instances, a device IP address can reveal a user's identity and connect them to certain Internet-enabled activities.[29] Depending on the nature and sensitivity of this linked information, this exposes users to potential reputational harm or embarrassment, as well as personal-autonomy harms. Associated information may also lead to harms when it is misused by malicious actors. For example, while a username and password may not provide much information about individual users, they can be used to gain access to private accounts from social media to financial services. This can lead to economic as well as reputational harms from fraud.

Without sufficient safeguards, associated information can cause significant harms that are often difficult to reverse. Many of these risks are also present in AR/VR, and in some cases may be exacerbated by the extent of information that can be combined with associated data. For example, a malicious actor could use associated credentials to not only access a user's account, but also to impersonate them in virtual space.

Mitigation Approaches for Privacy Concerns From Associated Data

Effective mitigation approaches for risks of harm from this data focus on ensuring authorized access to or use of this information, as well as limiting the extent to which it can be combined with other descriptive or identifying information. In some instances, combining identifying data with associated data can protect users against these harms through stronger user authentication. For example, if a username and password is linked to a biometric identifier such as a fingerprint, it will be more difficult for malicious actors to use that associated information to perpetrate fraud.

Laws and regulations addressing information security can also protect users from malicious misuse of their associated data, or mitigate harms if malicious actors gain access to this information. This includes both standards for information security and data protection, and data-breach notification requirements to alert users when potentially sensitive associated data has been compromised. Disclosure, transparency, and consent standards can also mitigate risks of harm by notifying users when a product or service combines their associated data with other information, including identifying information.

UNIQUE USER PRIVACY CONSIDERATIONS FOR AR/VR

Existing frameworks and mitigation approaches for privacy concerns offer a valuable foundation for addressing user privacy in AR/VR. However, these considerations alone are not sufficient to fully address the novel risks immersive technologies present. AR/VR devices and applications' immersive, multimodal user data collection and processing differs from other digital media in both scale and levels of sensitivity. For example, while both AR/VR platforms and other digital media can share and record observable videos in real time, immersive recording offers more-advanced capabilities—and requires more-extensive data collection—such as gaze and motion tracking and integrating associated information through virtual assets.[30]

AR/VR devices and applications' immersive, multimodal user data collection and processing differs from other digital media in both scale and levels of sensitivity.

Further, collaborative AR/VR, whether in the form of multiplayer games, social experiences, training simulations, virtual classrooms, or remote office solutions, will require new norms and expectations around privacy and conduct that are not present in other platforms such as social media and videoconferencing (or one-on-one and small-group "real-world" interactions). Interactions on AR/VR platforms will generate significant amounts of observed and observable data, from the details of conversations to individuals' movements and even physical responses in virtual space. [31] As one study noted, the more immersive the experience, the greater the possibility that a platform "may give the illusion of greater privacy … than is actually the case," with individuals "forgetting that their actions might become known to many more people than expected."[32]

In collecting and using this data, AR/VR presents novel concerns that cut across all four data types and warrant careful consideration. First, the extensive collection of biometric data, and the potential for that data to reveal personal information beyond simply identifying a user, raises privacy challenges that other biometric information technologies do not. Second, the extensive nature of this data collection and the experiences it enables leave many existing mitigation approaches insufficient or impractical for use in this context. Finally, the sensitive data collection as well as the immersive nature of AR/VR experiences exacerbate existing risks for acute harms vulnerable users face.

Biometric Data and Personal Autonomy

AR/VR poses new challenges to user control over personal autonomy and anonymity. Unlike other forms of digital media, AR/VR devices and applications must fully or partially translate multiple aspects of a user's identity and activity from the real world into virtual space. This includes not only aspects of an individual's identity such as biographical information or interests and affiliations, but also details about their location, movements, appearance, physical responses, and other information. When combined, this information effectively integrates "real" identity details provided directly by users with more-subtle features observed through their in-app/in-world activity that may reveal undisclosed information about their activities, interests, and preferences.[33]

Rather than creating a separate, alternative, or parallel identity, AR/VR users, particularly those in fully immersive experiences, identify with their virtual representations as a part of themselves. As photorealistic avatars become more widely used, particularly for use cases outside of entertainment, virtual identities may more-closely mirror physical reality.[34] This representation becomes even more accurate with the use of motion, gesture, and gaze-tracking, which replicate a user's physical responses within their virtual environment. While the ability to mirror an individual's expressions and reactions enhances interactions within virtual space, it also requires users to share—and allow devices to gather, track, and process—much more information than they would with other digital media platforms that simply transmit audiovisual information. This tracked "nonverbal data"—the subtle, subconscious movements sensors can detect—is virtually impossible for users to consciously control.[35]

These immersive identities mean users are unable to navigate virtual spaces with complete anonymity. For example, once an application connects identifying information (e.g., a full name) with biometric identification data (e.g., from eye-tracking cameras), it is nearly impossible to fully anonymize a user even if the identifying information is removed.[36] This exacerbates the potential for harm from observable, observed, computed, and associated data by inextricably combining them all.

Although biometric data collection is not unique to AR/VR, the scope of information gathered and the potential for additional information to be inferred is not seen in other consumer devices intended for use outside of controlled settings.

Perhaps the greatest distinction between immersive technologies and other digital media is the former's reliance on biometric information to replicate physical experiences in virtual space. Taken individually and without context, this information presents few concerns policymakers and privacy experts have not already raised in other technologies such as mobile phones and IoT devices. However, AR/VR devices' ability to aggregate and process extensive biometric information creates new risks of harm. Researchers at the Stanford Virtual Human Interaction Lab have estimated users generate "just under 2 million unique recordings of body language" in one 20-minute session in VR.[37]

If permitted, AR/VR devices and applications can infer significant additional information that reveals identifying biographical and demographic information, even if a user has not elected to provide these details. For example, the observed data gathered through eye-tracking technologies not only reveals where an individual is looking or what they are focusing on, it can also serve as an indicator of personal details such as age, gender, and race.[38] Other information, such as movements captured from motion or hand tracking, can uniquely identify users: In one study, five minutes of 6DoF tracking data while standing was sufficient to re-identify individuals with up to 95 percent accuracy, including across sessions.[39] Applications can also use biometric data to infer details about a user's physical and emotional responses to stimuli, as well as sensitive health information. Motion and eye tracking can capture a user's subconscious reactions, such as pupil dilation, which can in turn reveal inferred information about their interests and preferences—from favorite foods to sexual orientation.[40]

Although biometric data collection is not unique to AR/VR, the scope of information gathered and the potential for additional information to be inferred is not seen in other consumer devices intended for use outside of controlled settings. Without sufficient safeguards, the psychographic profiles generated from this information could cause direct harms, including discrimination and autonomy violations.[41] They can reveal sensitive information and other private details an individual did not choose to disclose. Further, additional harms can arise from unauthorized access to a user's biometric identifying information. Although biometric information (e.g., a fingerprint) may not reveal personal or identifying information on its own, it can be associated with an individual user for identification or authentication. Given the extent of biometric data involved, information security presents a notable privacy concern in AR/VR.

Limits to Existing Mitigation Practices

Because of the immersive nature of user data collection in AR/VR, many of the standard mitigation approaches used to limit harms from different types of data will fall short in the context of these experiences. First, user-focused measures of consent, transparency, and disclosure are more complicated in AR/VR than other digital and connected technologies. Users interact with fully or partially virtual spaces in different ways than they do on two-dimensional platforms, meaning standard consent practices may need to be reimagined for immersive experiences.[42] For example, users may not have the ability to easily click on hyperlinks that lead to additional information or otherwise indicate consent within a fully immersive experience.[43] In addition, while privacy preferences can allow users to opt out of sharing or choose not to disclose certain information, there are significant limits to this approach as sensitive or potentially identifying data is necessary for the core functions of immersive technologies.

Users interact with fully or partially virtual spaces in different ways than they do on two-dimensional platforms, meaning standard consent practices may need to be reimagined for immersive experiences.

Second, data anonymization is particularly difficult given the extent of identifying information provided and generated by users. Even if tracking data is de-identified by removing names, the raw biometric data can relatively easily re-identify users based on their unique movements.[44] Additional practices beyond simply removing names are necessary to truly de-identify sensitive biometric and biometrically derived data.[45] This makes secure storage and clear limits to access particularly important—which in turn raises questions about the benefits and risks of different approaches to data management. Most important are the questions of whether the user or a third-party has access to their data and whether proper safeguards are in place for the data.[46]

Finally, existing legal safeguards may not be sufficient to address the risks from various data gathered in AR/VR. For example, while there are laws in place to discourage nonconsensual pornography, they do not usually protect users from anonymity and autonomy harms from virtual reproductions of themselves or virtual assets. This policy gap is evident in the proliferation of "deepfakes," or synthetic media that replicates an individual's likeness, which has raised concerns about both personal autonomy and rights to publicity from digital replicas.[47] However, this risk is even more pronounced in immersive experiences, wherein the technical sophistication of synthetic media may not even be necessary: Once an unauthorized user gains access to or control of another user's virtual presence (e.g., by gaining access to their personal user account), it is relatively easy for them to impersonate the user or make it appear as though they did or said something without their consent. And malicious actors can also falsify a virtual presence without gaining such unauthorized access—with sufficient observable information, they can create replica avatars and other virtual assets. It is not difficult to imagine such a replica being used for fraudulent activities and to cause emotional, reputational, and economic harms. Further, laws protecting Fourth Amendment rights in the United States have yet to be applied explicitly to government requests for data from AR/VR devices and applications.[48]

Exacerbated Risk of Harm to Vulnerable Users

Any user data collection has the potential to disproportionately harm particularly vulnerable users, including children, older adults, and other marginalized and vulnerable populations. Risks of harm from personal information, such as discrimination or autonomy violations, are exacerbated for these individuals, while at the same time, they are also often less equipped to manage their personal privacy risks or give fully informed consent to data collection. Given the extent of information collected in AR/VR and the potential for misuse, it is worth noting the unique risks to its most vulnerable users. This is particularly important to consider when examining potentially sensitive use cases, such as in health care, child development, education, and certain workforce-training applications.

First, users who may already be susceptible to harms from discrimination or loss of anonymity or autonomy are particularly at risk in AR/VR. For example, individuals who use AR/VR as participants in health care research or as part of therapy for mental illness may generate both observed and computed data that could lead to discrimination in health care or employment should it be shared with service providers or employers. Higher risks to autonomy and anonymity are also a serious concern. For example, without sufficient safeguards, individuals who could face discrimination or even physical harm based on age, sex, race, sexual orientation or certain health conditions may generate observed biometric data that could be used to infer this information without their consent.

Second, it is important to consider the extent to which child users are able to provide full consent to the data collection that takes places in immersive experiences. As discussed, translating standard consent mechanisms into immersive, three-dimensional systems is already a challenge. The same is true for traditional approaches to age-appropriate content and children's safety on digital platforms, such as parental controls, age verification, and limits on personal data collection.[49] Not only is it unreasonable to expect children to fully grasp the extent and purpose of personal data collection (and thereby give informed consent to its use), they may also lack the ability to fully differentiate between real-world and virtual elements within immersive experiences, which could further expose them to harm from sharing personal information. Further, because children cannot fully consent to the risks to personal autonomy and anonymity that sharing identifying or biometric information in AR/VR present, failure to mitigate these risks at a young age can expose them to long-term harm.[50] For example, motion-tracking information from a device someone used as a child could later be used to re-identify them on a new device or application.

THE REGULATORY LANDSCAPE FOR AR/VR USER PRIVACY

Although the technology itself may be unique, AR/VR devices and applications are already subject to a number of laws and regulations governing individual privacy and user data in the United States. However, the current regulatory landscape addresses only some of the risks from AR/VR, and certain requirements complicate the data collection necessary to provide robust and secure immersive experiences across sectors. Further, just as AR/VR presents new considerations for user privacy, AR/VR technologies also introduce unique challenges to developing policies to address these concerns. Because of this, current legal and regulatory frameworks for these technologies leave critical gaps in policy around some concerns, while requiring unnecessarily complex responses to others.

Existing Laws and Regulations for Privacy in AR/VR

The current regulatory landscape for user privacy in the United States is a patchwork of national- and state-level legislation addressing various concerns. Privacy regulations at the national level address both specific types of particularly vulnerable users and sensitive data, and requirements for data collection and management by certain actors. COPPA, the Family Educational Rights and Privacy Act (FERPA), and HIPAA all regulate data that may be gathered through immersive experiences. However, these regulations address only specific purposes of information, rather than more general information types. For example, COPPA restricts the collection and storage of observed and observable data such as biographical information, recordings, and geolocation information—but only in products and services intended for children.[51]

Laws are also in place to regulate government use of digital information, including data gathered in AR/VR. The Privacy Act of 1974 regulates federal agencies' management of records about individuals, and could include data collected during any agency use of AR/VR technologies.[52] Nearly a century of case law provides additional safeguards for law enforcement's use of personal and digital information, including audiovisual recordings and certain forms of metadata.[53] However, there have been no legal decisions that address the scope and scale of data gathered and inferred from an AR/VR device.

At the state level, there are several laws related to biometric privacy and data protection more broadly. In California, the California Consumer Privacy Act (CCPA) includes compliance requirements for companies that collect, process, and share personal data, including biometric data. While CCPA does not address AR/VR specifically, companies providing these technologies have adopted a mixture of regulation-specific and more-general compliance measures as a response.[54] Illinois, Texas, and Washington have all introduced laws that specifically target biometric data collection and facial-recognition technologies, with specific requirements for notice and user consent when such data is gathered.[55] This further complicates the regulatory environment and compliance requirements for AR/VR providers, which, by nature of the technology, will likely have to collect some form of biometric information from their users.

Policy Challenges in Addressing AR/VR User Privacy

Given the extent of information required to operate AR/VR devices and applications, and the sensitive nature of much of that data, these technologies introduce new challenges to policy debates about user privacy. First, there is no clear baseline for the amount of data necessary to enable core functions (as opposed to additional benefits that allow users to determine whether they are willing to share personal data in order to utilize them). This complicates any consent requirements that may be imposed on AR/VR providers. Indeed, the XR Safety Initiative, a multi-stakeholder initiative working on privacy and safety in immersive experiences, noted that "there might be a question over whether subjects have a 'real choice' to refuse the processing [of sensitive data] and whether it is possible to draw the line between necessary and unnecessary data."[56] Further, because AR/VR is still a relatively new technology, there is still little consumer awareness of the purpose of its extensive data collection. A challenge for policymakers going forward will be separating negative reactions to these unknowns—which can be naturally addressed as public awareness and understanding grows—and the actual potential for harm from data collection. Most users are willing to exchange some data for free or low-cost online services.[57]

Many of these regulations do not translate directly to immersive experiences, instead prohibitively restricting some uses while leaving other risks unaddressed.

Another challenge facing policymakers is defining the scope of identifying information and sensitive data. Current legal and policy frameworks for securing personal, identifying, or even biometric data do not cover the extent of sensitive information collected in AR/VR or its potential uses beyond user identification or authorization.[58] Further, data protection regulations in the United States and around the world address privacy concerns of individuals as data subjects, but do not extend definitions of "personal" or identifying information to include fully virtual spaces or assets.[59] This makes it difficult to apply existing understandings of user privacy that form the foundation for privacy law and policy in the context of AR/VR. While the overarching goals of privacy regulations can certainly be extended to AR/VR devices and applications, the mechanisms to achieve those objectives laid out in current laws such as the General Data Protection Regulation (GDPR) may not directly apply to—or indeed be sufficient for—immersive experiences.

Finally, the patchwork regulatory environment currently in place presents a long list of standards and compliance requirements that apply to select aspects of AR/VR technologies. Because of the comprehensive data collection required, restrictions on one type of information are effectively restrictions on the technology as a whole. Many of these regulations do not translate directly to immersive experiences, instead prohibitively restricting some uses while leaving other risks unaddressed. This presents a critical policy challenge: first, in clarifying application of these regulation in the context of AR/VR technologies; and second, in harmonizing these various regulations to ensure they do not overly restrict the development and use of AR/VR devices and applications.

RECOMMENDATIONS

The scope and scale of user data collection in AR/VR present important questions about how to protect users from harms while encouraging continued innovation of this rapidly developing technology. Policymakers and developers alike should pursue robust solutions that put necessary safeguards in place and set standards for user privacy that can apply to existing and future use cases. This will require careful consideration of the component parts of AR/VR technologies, the types of information they collect, and the potential for harm from that information. The approaches put in place today will impact how AR/VR devices and applications are developed for consumer, enterprise, and even government use well into the future.

Create an Innovation-Friendly Regulatory Environment to Address User Privacy Concerns in AR/VR

Policymakers should ensure that the regulatory environment for AR/VR development provides a framework for user privacy protections while allowing developers to explore additional safeguards. As this report has described, many of the standard mitigation approaches for privacy concerns from digital technologies and connected devices are not easily transferred into immersive experiences. Policy measures should not restrict the ability of AR/VR developers to create innovative mechanisms to mitigate privacy risks, including for user consent, transparency, and choice.

Provide Guidance and Clarification on Existing Privacy Laws' Applications to AR/VR

Many of the existing laws on data privacy will apply to certain data in AR/VR. However, the extent to which AR/VR data collection practices align with these rules remains ambiguous. This leaves companies developing AR/VR devices and applications in a state of regulatory uncertainty in which it is unclear when and how their products are expected to comply with federal and state regulations. A lack of clear guidelines will discourage innovation, particularly in nascent fields with less-straightforward regulatory guidance (e.g., tracking and BCI technologies) or strictly regulated fields wherein existing interpretations are not directly applicable to AR/VR (e.g., health data and products developed for children).

The relevant federal agencies and regulatory bodies that oversee existing privacy regulations should provide explicit guidance on their application in immersive contexts. Such guidance should synthesize existing regulatory frameworks consistently at the federal level and discourage further state-by-state fragmentation. In doing so, regulators should carefully consider the data types outlined in this report and their necessity to relevant AR/VR devices and applications. For example, the Department of Health and Human Services may consider which observable, observed, and resulting computed data gathered by AR/VR constitute "protected health information" under HIPAA when these technologies are used in health care. Similarly, the Federal Trade Commission (FTC) can issue additional clarification on COPPA compliance for AR/VR data collection, such as when collecting audiovisual recording or geolocation data may be appropriate for the core functions of AR/VR devices and applications, as the agency previously did for voice recordings.[60]

Reform Privacy Laws That Would Unnecessarily Limit AR/VR

Many laws developed with a specific technology or use case in mind impose unnecessary limits on AR/VR innovation. Because AR/VR devices and applications require extensive information to perform even basic functions, regulations intended for standalone technologies can impose compliance standards that are difficult or impossible for AR/VR providers to maintain. At the state level, policymakers should review data privacy laws that are meant to address specific use cases or data collection practices and revisit those that may impose prohibitive restrictions on AR/VR technologies. Particularly noteworthy here are state laws governing biometric data, which are generally written to address biometric identification but could impede AR/VR functions that require observed biometric information, such as eye and motion tracking.

At the federal level, Congress should address the risk of fragmented state privacy laws by establishing a unified national privacy framework that preempts state laws.[61] In addition, highly specific data privacy laws such as COPPA warrant similar consideration. Because it is nearly impossible to limit collection of certain types of data without limiting the quality and functionality of an AR/VR device or application as a whole, restrictions on broadly defined personal information including biographic information, voice recordings, and geolocation, could severely limit the possibilities for AR/VR technologies to be used in child-focused contexts. At the same time, such restrictions will not necessarily protect children against many of the potential harms and privacy risks that are unique to or exacerbated by AR/VR. For example, children may still experience harassment, access inappropriate content, or share identifying or otherwise harmful information with strangers in multiplayer environments. In addition to providing specific guidance on COPPA compliance for AR/VR, the FTC should work with developers to establish best practices for safeguards and technical measures that protect children from harm in immersive experiences. Such best practices could inform any future proposed changes to COPPA regulations.

Create Rules to Safeguard Against New Risks of Harm

The current regulatory landscape for AR/VR not only places unnecessary restrictions on these technologies, it also leaves notable gaps in protections against acute potential harms. As this report has discussed, the information AR/VR devices gather from sensors, including eye and motion tracking, poses privacy risks as both raw observed metadata and inferred computed data. It also serves a number of purposes beyond biometric identification, and can be used to infer significant personal and identifying information about users. Yet existing definitions of personal and biometric data do not account for this widespread collection and processing of biometric information beyond purposes of identification.[62] This could expose users to harms from unauthorized access to or malicious use of data that falls outside of existing definitions, such as deriving computed psychographic data to infer sensitive personal information.

Policymakers should establish a clear definition of personal and identifying information to include highly sensitive biometric data and necessary uses beyond user authentication that both encourage greater protection of this data and allow for a variety of use cases. They should consider not only observed biometric information, but also the ways in which this information can be manipulated to reveal additional details about a user. Importantly, any rules pertaining to biometric and biometrically derived information should not directly ban its collection. Devices and applications may require varying levels of information to function. For example, while collecting precise computed information about a user's reactions and preferences from gaze tracking may be seen as extraneous for a social experience, this same information could be necessary in a market research context.

Although not yet widely in use, BCI data presents similar definitional concerns. While any direct regulation of BCI technologies would be premature, policymakers should carefully consider how to classify both observed BCI inputs and inferred computed data from BCI-enabled devices. Any privacy regulations should include clear definitions of biometric identifying and biometrically derived data, and present transparency, consent, and choice requirements consistent with the purpose of its collection and risks of harm. While both may be considered personal information, biometric identifying data (e.g., facial recognition) will require different safeguards than biometrically derived information (e.g., inferred data about personal preferences from eye-tracking or even BCI technologies). Distinguishing between the two will ensure users are adequately protected from harm while allowing for innovative uses of this information with proper user transparency and choice.

Further, there should be clear guidelines for use of AR/VR data in legal investigations. Because the data from AR/VR devices and applications can form a comprehensive profile of an individual, it can be a valuable tool for law enforcement and legal proceedings. However, existing case law on digital information indicates that such use could violate individuals' Fourth Amendment rights in the United States. Immersive experiences also introduce a new dimension of law enforcement use: real-time virtual presence. Multiuser AR/VR platforms raise questions about the extent to which third-party doctrine and other legal frameworks extend to investigators or law enforcement officials interacting with, observing, and recording the activities of users in fully or partially virtual spaces. Policymakers should introduce new legal safeguards for comprehensive user data in AR/VR, including clear guidelines for when access to this information would first require a warrant.

Enact Federal Privacy Legislation to Harmonize Compliance Requirements and Enable Innovation

While reforming or introducing rules to address the unique nature of user privacy in AR/VR can protect users and allow for innovation in the short term, comprehensive national privacy legislation would better position regulators and developers alike to ensure necessary safeguards are consistently implemented as these technologies continue to evolve. Policymakers should enact privacy legislation that:

- Establishes clear guidelines for the collection, processing, and sharing of various types of data with consideration for varying levels of sensitivity;
- Implements user data privacy rights and safeguards against risks of harm; and
- Strengthens notice, transparency, and consent practices to ensure users can make informed decisions about the data they choose to share, including sensitive biometric and biometrically derived information.

Regulatory harmonization should also consider sector- and purpose-specific privacy regulations, such as HIPAA and FERPA. Such regulations may impose additional and potentially conflicting requirements on AR/VR technologies that could impact their potential to be used in sectors where they could offer significant value, such as health care and education. Regulators should ensure that requirements are consistent across sector-specific regulations, and any such specific requirements augment, rather than contradict or complicate, broader federal privacy legislation. Such an approach can address potentially conflicting compliance requirements and set clear standards for user privacy protections for both existing and emerging data collection practices in AR/VR.

Encourage Voluntary Practices to Secure User Privacy in AR/VR

Particularly in the absence of comprehensive federal privacy legislation, AR/VR developers and policymakers should work together to develop effective self-regulatory approaches to ensure user privacy. Clear, consistent standards and practices will enable the companies building AR/VR devices and applications to ensure their products implement appropriate safeguards, while also providing policymakers and regulatory bodies with a more complete understanding of the mitigation approaches that are both most effective and technically feasible.

In consultation with AR/VR developers, federal agencies, including the Departments of Education, Health and Human Services, and Transportation, should develop voluntary standards for user privacy protection in AR/VR for relevant use cases. Such a framework should be built off of existing standards and best practices, with consideration for the unique or exacerbated risks and potential for harm AR/VR technologies present. Importantly, any voluntary standards should include input from AR/VR developers across sectors and industries, including workforce development, education, health care, and entertainment.

There are several contributions to this area which could be used as a starting point for an AR/VR-specific framework, such as the XR Safety Initiative's Privacy Framework, the Open AR Cloud Privacy Manifesto, and the XR Association's series of Developers Guides.[63] While existing documents, such as the National Institute of Standards and Technology (NIST) Privacy Framework, offer a blueprint, the extent of data collected by AR/VR devices and applications warrants more industry-specific standards. [64] This includes:

- Transparency and disclosure standards and mechanisms for immersive experiences, including clear disclosure of how sensitive biometric data is collected and used;
- User privacy controls and opt-out mechanics for information that is not critical to core functions;
- Information security standards, including guidelines for encryption and local storage of highly sensitive information such as biometric identifiers or spatial mapping of private homes; and
- Guidelines for the collection and use of biometrically derived information used for purposes other than user identification.

Such a standards-setting exercise could also reveal areas in which policy intervention is necessary to fully protect users from harm, such as laws addressing autonomy violations and discrimination.

CONCLUSION

AR/VR devices and applications offer a glimpse into a future that is more connected, adaptive, and rich in immersive experiences. But they also introduce a level of user data collection and privacy concerns that other consumer technologies have not. In order to realize this potential while mitigating privacy risks, developers and policymakers alike should consider the actual harms that could arise from this extensive data collection.

It is critical to approach AR/VR not as one monolithic technology, but as a collection of numerous information-gathering technologies delivering a singular experience. The different observable, observed, computed, and associated data collected will vary both in sensitivity and potential for harm. Addressing these actual harms, rather than the technologies themselves, will allow policymakers and developers to differentiate between user preferences and critical privacy risks. Privacy preferences can largely be addressed by developers directly, while additional policy interventions may be necessary to mitigate risks of harm from user information.

The complexities of user privacy in AR/VR require equally nuanced approaches to mitigating the very real threats the scope and scale of data collection present. However, it is important to resist imposing reactionary measures that could obstruct innovations in AR/VR devices and applications as well as new approaches to user privacy in immersive experiences. By approaching these questions in terms of the types of information gathered and actual harms that could arise, policymakers can shape a regulatory environment that encourages this innovation while establishing necessary safeguards.

About the Author

Ellysse Dick (@Ellysse_D) is a policy analyst in tech and cyber policy at ITIF. Her research focuses on AR/VR innovation and policy including privacy, safety, and accountability. Prior to ITIF, she led communications and outreach for the Women in Public Service Project at the Wilson Center. She holds a Master of Arts in Law and Diplomacy from the Fletcher School at Tufts University and a BA in International Affairs and German Studies from the University of Colorado.

About ITIF

The Information Technology and Innovation Foundation (ITIF) is an independent, nonprofit, nonpartisan research and educational institute focusing on the intersection of technological innovation and public policy. Recognized by its peers in the think tank community as the global center of excellence for science and technology policy, ITIF's mission is to formulate and promote policy solutions that accelerate innovation and boost productivity to spur growth, opportunity, and progress.

For more information, visit us at www.itif.org.

ENDNOTES

1. XR Safety Initiative, The XRSI Privacy Framework Version 1.0, September 2020, https://xrsi.org/publication/the-xrsi-privacy-framework.
2. Daniel Castro and Alan McQuinn, "ITIF Filing to FTC on Informational Injury Workshop," Information Technology and Innovation Foundation, October 27, 2017, http://www2.itif.org/2017-informational-injury-comments.pdf.
3. Ibid.
4. Fiachra O'Brolcháin et al., "The Convergence of Virtual Reality and Social Networks – Threats to Privacy and Autonomy," Science and Engineering Ethics 22 (2016), https://doi.org/10.1007/s11948-014-9621-1.
5. Brittan Heller, "Reimagining Reality: Human Rights and Immersive Technology," Carr Center Discussion Paper Series, June 12, 2020, https://carrcenter.hks.harvard.edu/publications/reimagining-reality-human-rights-and-immersive-technology.
6. Castro and McQuinn, "ITIF Filing to FTC on Informational Injury Workshop."
7. According to the Cyber Civil Rights Initiative, 46 states as well as the District of Columbia and Guam had "revenge porn" (nonconsensual pornography) laws in effect as of February 2021. See Cyber Civil Rights Initiative, "46 States + DC + One Territory Now Have Revenge Porn Laws," accessed February 10, 2021, https://www.cybercivilrights.org/revenge-porn-laws.
8. Castro and McQuinn, "ITIF Filing to FTC on Informational Injury Workshop."
9. Kei Studios, "A Quick Guide to Degrees of Freedom in Virtual Reality," accessed February 10, 2021, https://kei-studios.com/quick-guide-degrees-of-freedom-virtual-reality-vr.
10. Oculus Blog, "Introducing Hand Tracking on Oculus Quest—Bringing Your Real Hands into VR," oculus.com, September 25, 2019, https://www.oculus.com/blog/introducing-hand-tracking-on-oculus-quest-bringing-your-real-hands-into-vr.
11. See for example HTC's VIVE Eye Tracking data collection: HTC, "Terms: Learn More," htc.com, accessed February 10, 2021, https://www.htc.com/us/terms/learn-more.
12. VIVE Enterprise, "VIVE Pro Eye Office," enterprise.vive.com, accessed February 10, 2021, https://enterprise.vive.com/us/product/vive-pro-eye-office.
13. Heller, "Reimagining Reality."
14. Scott Hayden, "Valve Psychologist: Brain-Computer Interfaces Are Coming & Could Be Built Into VR Headsets," Road to VR, March 23, 2019, https://www.roadtovr.com/valve-brain-computer-interfaces-vr-ar-gdc-2019.
15. Tia Ghose, "Eye Tracking Could Diagnose Brain Disorders," *Live Science*, September 18, 2012, https://www.livescience.com/23274-eye-tracking-gaze-brain-disorders.html.
16. Albert Liu, "Why VR Analytics is Critical to Prove ROI," cognitive3D blog, December 3, 2019, https://content.cognitive3d.com/vr-merchandising-case-study.
17. See: Avi Bar-Zeev, "The Eyes Are the Prize: Eye-Tracking Technology is Advertising's Holy Grail," *VICE*, May 28, 2019, https://www.vice.com/en/article/bj9ygv/the-eyes-are-the-prize-eye-tracking-technology-is-advertisings-holy-grail; XR Safety Initiative, "3.3.1:The XR Safety Initiative (XRSI)'s Child Safety Risks and Recommendations," The XRSI Privacy Framework Version 1.0.
18. Daniel Christensen and Diana Jimenez, "Data Protection Should Extend to Virtual Places and Data Objects," IAPP Privacy Perspectives, August 24, 2016, https://iapp.org/news/a/data-protection-should-extend-to-virtual-places-and-data-objects.
19. Castro and McQuinn, "ITIF Filing to FTC on Informational Injury Workshop."

20. U.S. Equal Employment Opportunity Commission, “Laws Enforced by EEOC,” eeoc.gov, accessed February 10, 2021, https://www.eeoc.gov/statutes/laws-enforced-eeoc.
21. See for example: U.S. Department of Justice Civil Rights Division, “The Fair Housing Act,” justice.gov, accessed February 10, 2021, https://www.justice.gov/crt/fair-housing-act-1; U.S. Department of Labor Employee Benefits Security Administration, “FAQs on HIPAA Portability and Nondiscrimination Requirements for Employers and Advisers,” U.S. Department of Labor, accessed February 10, 2021, https://www.dol.gov/sites/dolgov/files/ebsa/about-ebsa/our-activities/resource-center/faqs/hipaa-compliance.pdf.
22. Louise Matsakis, “The Supreme Court Just Greatly Strengthened Digital Privacy,” *WIRED*, June 22, 2018, https://www.wired.com/story/carpenter-v-united-states-supreme-court-digital-privacy.
23. Castro and McQuinn, “ITIF Filing to FTC on Informational Injury Workshop.”
24. Shangchen Han et al., “Using Deep Neural Networks for Accurate Hand-Tracking on Oculus Quest,” Facebook AI, September 25, 2019, https://ai.facebook.com/blog/hand-tracking-deep-neural-networks.
25. Scott Stein, “Mind Control Comes to VR, Letting Me Explode Alien Heads with a Thought,” cnet, January 30, 2021, https://www.cnet.com/news/controlling-vr-with-my-mind-nextminds-dev-kit-shows-me-a-strange-new-world.
26. Jacob Leon Kröger et al., “What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking,” IFIP Advances in Information and Communication Technology 576 (2020), https://doi.org/10.1007/978-3-030-42504-3_15.
27. Ibid.
28. Castro and McQuinn, “ITIF Filing to FTC on Informational Injury Workshop.”
29. Nefi Acosta, “Are IP Addresses ‘Personal Information’ Under CCPA?” IAPP Privacy Advisor, April 28, 2020, https://iapp.org/news/a/are-ip-addresses-personal-information-under-ccpa.
30. XR Safety Initiative, “3.4.1.1: FERPA: Protection of Education Record Considerations,” The XRSI Privacy Framework Version 1.0.
31. O’Brolcháin et al., “The Convergence of Virtual Reality and Social Networks.”
32. Ibid.
33. Heller, “Reimagining Reality.”
34. O’Brolcháin et al., “The Convergence of Virtual Reality and Social Networks.”
35. Jeremy Bailenson, “Protecting Nonverbal Data Tracked in Virtual Reality,” JAMA Pediatrics, August 6, 2018, https://vhil.stanford.edu/mm/2018/08/bailenson-jamap-protecting-nonverbal.pdf.
36. Bar-Zeev, “The Eyes Are the Prize.”
37. Bailenson, “Protecting Nonverbal Data Tracked in Virtual Reality.”
38. Joseph Jerome, “Establishing Privacy Controls for Virtual Reality and Immersive Technology,” IAPP Privacy Perspectives, September 9, 2020, https://iapp.org/news/a/establishing-privacy-controls-for-virtual-reality-and-immersive-technology.
39. Mark Roman Miller et al., “Personal Identifiability of User Data During Observation of 360-Degree VR Video,” Scientific Reports 10 (2020), https://doi.org/10.1038/s41598-020-74486-y.
40. Bar-Zeev, “The Eyes Are the Prize.”
41. Ibid.
42. Bailenson, “Protecting Nonverbal Data Tracked in Virtual Reality.”
43. Erin Egan, “II.C: New Technologies and Interfaces,” in *Communicating About Privacy: Towards People-Centered and Accountable Design*, Facebook, July 2020, https://about.fb.com/wp-content/uploads/2020/07/Privacy-Transparency-White-Paper.pdf.

44. Jessica Outlaw and Susan Persky, “Industry Review Boards are Needed to Protect VR User Privacy,” World Economic Forum, August 29, 2019, https://www.weforum.org/agenda/2019/08/the-hidden-risk-of-virtual-reality-and-what-to-do-about-it.

45. Ann Cavoukian and Daniel Castro, “Setting the Record Straight: De-Identification Does Work,” Information Technology and Innovation Foundation, June 16, 2014, https://itif.org/publications/2014/06/16/setting-record-straight-de-identification-does-work.

46. Bar-Zeev, “The Eyes Are the Prize.”

47. For example, the New York State legislature proposed a bill that would extend current rights to publicity to include digital replicas—restricting the use of such replicas for trade or commercial purposes. See: New York State Senate, 2017-2018 legislative session, A8155B, https://www.nysenate.gov/legislation/bills/2017/a8155.

48. Bailenson, “Protecting Nonverbal Data Tracked in Virtual Reality.”

49. The XR Safety Initiative, “1.5.6: Child Safety,” XRSI Privacy Framework Version 1.0.

50. Outlaw and Persky, “Industry Review Boards are Needed to Protect VR User Privacy.”

51. U.S. Federal Trade Commission, “Complying with COPPA: Frequently Asked Questions,” ftc.gov, accessed February 10, 2021, https://www.ftc.gov/tips-advice/business-center/guidance/complying-coppa-frequently-asked-questions.

52. U.S. Department of Justice Office of Privacy and Civil Liberties, “Overview of the Privacy Act of 1974 (2020 Edition),” justice.gov, accessed February 10, 2021, https://www.justice.gov/opcl/overview-privacy-act-1974-2020-edition.

53. Matsakis, “The Supreme Court Just Greatly Strengthened Digital Privacy.”

54. Bailenson, “Protecting Nonverbal Data Tracked in Virtual Reality.”

55. Ross D. Emmerman et al., “New Biometric Information Privacy Cases Reveal Breadth of Potential Exposure for Companies,” Loeb & Loeb, March 2018, https://www.loeb.com/en/insights/publications/2018/03/new-biometric-information-privacy-cases-reveal-b__.

56. The XR Safety Initiative, “1.5.2: Consent,” XRSI Privacy Framework Version 1.0.

57. Bailenson, “Protecting Nonverbal Data Tracked in Virtual Reality.”

58. Heller, “Reimagining Reality.”

59. Christensen and Jimenez, “Data Protection Should Extend to Virtual Places and Data Objects.”

60. U.S. Federal Trade Commission, “FTC Provides Additional Guidance on COPPA and Voice Recordings,” ftc.gov, October 23, 2017, https://www.ftc.gov/news-events/press-releases/2017/10/ftc-provides-additional-guidance-coppa-voice-recordings.

61. Alan McQuinn and Daniel Castro, “A Grand Bargain on Data Privacy Legislation for America,” Information Technology and Innovation Foundation, January 14, 2019, https://itif.org/publications/2019/01/14/grand-bargain-data-privacy-legislation-america .

62. Heller, “Reimagining Reality.”

63. XRSI Privacy Framework Version 1.0; Jan-Erik Vinje, “Privacy Manifesto for AR Cloud Solutions,” Open AR Cloud on Medium, October 17, 2018, https://medium.com/openarcloud/privacy-manifesto-for-ar-cloud-solutions-9507543f50b6; XR Association, “Research & Best Practices,” xra.org, accessed February 10, 2021, https://xra.org/research-best-practices.

64. National Institute of Standards and Technology, NIST Privacy Framework, U.S. Department of Commerce National Institute of Standards and Technology, January 16, 2020, https://www.nist.gov/system/files/documents/2020/01/16/NIST%20Privacy%20Framework_V1.0.pdf.