Virtual Worlds Need REAL Governance of Privacy and Safety

Kavya Pearlman
XR Safety Initiative
San Francisco Bay Area, CA
kavya@xrsi.org

Joel Scharlat
IVEA Consulting
Ashburn, VA
joel@iveaconsulting.com

ABSTRACT

As the world becomes increasingly connected, and immersive technologies gain wider adoption both within government and businesses as well as in the consumer market, a framework for security and privacy for these devices is required. From headsets to other wearables and related sensors, eXtended Reality (XR) technologies are now capable of gathering untold quantities of biometric data about users, potentially everything from a user's location and skin color to their eye and hand positions at any given time. The National Institute of Standards and Technology (NIST) has offered basic guidance, while regional laws such as General Data Protection Regulations (GDPR), Children's Online Privacy Protection Rule (COPPA), and Family Educational Rights and Privacy Act (FERPA) govern some forms of data in specific locations. Despite the existing guidelines and regional laws, comprehensive protections are not in place to protect individuals and stakeholders in XR. With this in mind, the XR Safety Initiative (XRSI) developed a privacy and safety framework that sets a baseline set of standards, guidelines, and best practices that are regulation agnostic. It incorporates privacy requirements drawn from the GDPR, NIST guidance, FERPA, COPPA, and other evolving laws. The framework is designed to adapt and include novel requirements as new regulations come into effect. With version 1.1 expected to be published in 2021, this paper provides an overview of the framework, how it was developed, and highlights changes and additions in version 1.1. We also discuss who can benefit from it and offer guidance to organizations, developers, and service providers on how to implement the framework for added security and privacy designed into their product or service. This paper also provides government customers with an understanding of what a security posture should include-beyond traditional authorities to operate (ATOs)-as more organizations look to adopt emerging technologies and stake a claim in the Metaverse.

ABOUT THE AUTHORS

Kavya Pearlman Well known as the "Cyber Guardian", founder and CEO of the XR Safety Initiative (XRSI), Kavya Pearlman is an award-winning cybersecurity professional with a deep interest in immersive and emerging technologies. She recently launched a novel XRSI Privacy and Safety Framework for the XR domain. Kavya is constantly exploring new technologies to solve cybersecurity challenges. She has been named one of the Top Cybersecurity influencers for three consecutive years, 2018-2019-2020 by IFSEC Global. Kavya has previously advised Facebook on third-party security risks during the 2016 US presidential elections and worked as the head of security for the oldest virtual world, "Second Life" by Linden Lab. Kavya is the co-host of the immersive podcast "Singularity Watch" and one of the Top 50 speakers in the cybersecurity industry. Kavya has founded The CyberXR Coalition that now focuses on diversity and inclusion and the cross-section of Cybersecurity and XR, helped launch a trustworthy XR news platform, ReadyHackerOne, and established a Medical XR Advisory Council.

Joel Scharlat is the President and CEO of IVEA Consulting, a premier consulting firm focused on the security of immersive technologies. A retired Marine, Joel has been working with immersive technologies for over 15 years having conducted extensive research on how to operationalize immersive environments. Joel is also a Cybersecurity Advisor with XRSI, a global nonprofit dedicated to creating safe and secure experiences in XR, and the Director of Operations for Cyber Bytes Foundation, a 501(c)(3) focused on cybersecurity education, innovation, and outreach. Joel has a Master of Science degree from The Naval Postgraduate School where he wrote his thesis on using immersive environments to conduct information operations.

Virtual Worlds Need REAL Governance of Privacy and Safety

Kavya Pearlman
XR Safety Initiative
San Francisco Bay Area, CA
kavya@xrsi.org

Joel Scharlat
IVEA Consulting
Ashburn, VA
joel@iveaconsulting.com

INTRODUCTION

Extended Reality (XR) is a fusion of all the simulated realities–including Augmented Reality (AR), Virtual Reality (VR), and Mixed Reality (MR)–which consists of technology-mediated experiences enabled via a broad spectrum of hardware and software, including sensory interfaces, applications, and infrastructures. XR is often referred to as immersive video content, enhanced media experiences, and interactive and multi-dimensional human experiences. With the mass adoption of XR, technologies such as Brain-computer interfaces (BCI) and autonomous systems enabled via Machine Learning are increasingly interacting with XR devices to produce astonishing virtual worlds. These mirror worlds built by extending realities and fueled by massive amounts of data are going to shape many different industries, including many highly sensitive and regulated ones, such as healthcare and the military. Due to the uncharted and evolving nature of these brave new worlds, individuals and organizations are currently not fully aware of the irreversible and unintended consequences of XR on the digital and physical world (O'Brolcháin et al., 2016). With the mass adoption of emerging technologies including the US Army's procurement of the Hololens 2-based IVAS program (*U.S. Army to Use HoloLens Technology in High-Tech Headsets for Soldiers*, 2021), it is imperative to understand privacy and safety concerns and proactively address them. This framework by the XR Safety Initiative (XRSI) provides a baseline approach enabling better engineering practices that support privacy by design concepts and help organizations protect data and safeguard the XR platforms, applications, and associated ecosystems. The framework is the work of several interdisciplinary experts and serves as a tool for improving privacy, security, and safety through human-centric design, pragmatic decision-making, and proactive risk management.

HOW WE GOT HERE

From headsets to other wearables and related sensors, XR technologies are now capable of gathering untold quantities of user biometric data either by direct data stream or by inferring from combined data sources, potentially including everything from a person's location and skin color to their eye and hand positions at any given time. Some, if not most, of this information has tremendous value to individuals and organizations like the US military, but comprehensive regulations are not in place to protect XR stakeholders including consumers, private and public organizations, researchers and governments, etc. The National Institute of Standards and Technology (NIST) has offered basic guidance, while regional laws such as General Data Protection Regulation (GDPR), Children's Online Privacy Protection Act (COPPA), and the Family Educational Rights and Privacy Act (FERPA) govern some forms of data in specific contexts. XRSI's framework integrates them, adopting a more comprehensive approach that ties together regulations, practices, threat analyses, and specific use cases. The framework goes beyond the regulations and provides a greater understanding of these complex ecosystems. The XRSI Framework is not intended to be used as a compliance checklist, but as a baseline risk assessment model for building trust in the XR domain. This approach strongly relies on a deep understanding of the least visible layer of the XR experiences: data collection and usage.

Data Collection in Immersive Technologies

Many organizations are developing XR technology to build all-day wearable XR glasses that are spatially aware to deliver AR and VR experiences more immersive and integrated into the physical world. To achieve this outcome via the use of AI algorithms, a large amount of data collection is necessary. This is where most XR organizations and regulators face what we consider a Collingridge dilemma.

"The social consequences of a technology cannot be predicted early in the life of the technology. When change is easy, the need for it cannot be foreseen; when the need for change is apparent, change has become expensive, difficult and time-consuming."
-David Collingridge

According to the Collingridge dilemma, the XR domain, just like any other emerging technology research and development, faces a paradox or double-bind problem:

- An information problem: impacts cannot be easily predicted until the technology is extensively developed and widely used.
- A power problem: control or change is difficult when the technology has become fully entrenched. (Tech Liberation, 2021)

The dilemma equally applies to data collection in XR. Data is the heartbeat of XR. Meaningful experiences can only be realized when we use data *about* an individual *for* that individual's experience. This is the immersive nature of the technology. This inherently means developers need to know information about people – where you are, what you're looking at, if you're moving or not, etc. Developers may likely opt for a "more is better" approach when it comes to data collection. But the concerns for excessive data collection are heightened because of the large amount of real-time data collection and the potential inferences that are possibly made. While some inferences are necessary and even welcomed by individuals such as curated and customized shopping preferences many others such as Biometrically-Inferred Data (BID) are not and may cause harm to humans. BID is a collection of datasets resulting from information inferred from behavioral, physical, and psychological biometric identification techniques, and other nonverbal communication methods. For example, Figure 1 shows how just one type of sensor collecting just one type of data on XR devices can lead to inferences such as biometric and gender identity, mental workload, mental health status, cognitive abilities, religious and cultural background, physical health, geographic origin, and many other skills, abilities, personality traits, and more. Based on different jurisdictions, organizations may be mandated to protect BID and require adopting a data governance framework like XRSI's. This will prevent excessive and unwarranted data collection at the hardware, operating system, API, and software level, leading to responsible research and innovation in the XR domain. Below are just a select few examples of the data tracked and collected by and for XR devices.

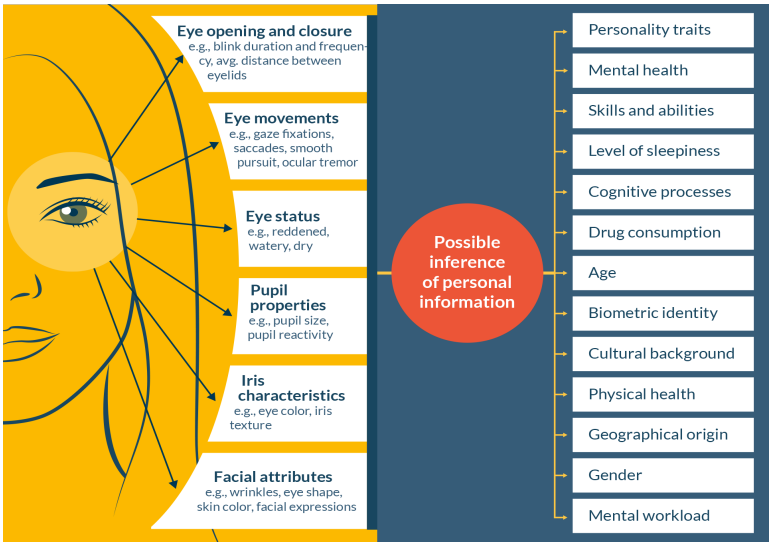

Figure 1. Biometrically Inferred Data commonly captured by eye trackers and sensors. CC BY-NC-SA 4.0

Data Collection Examples

A 2018 research study by Stanford University highlighted that "with VR, in addition to recording personal data regarding people's location, social ties, verbal communication, search queries, and product preferences, technology companies will also collect nonverbal behavior–for example, users' posture, eye gaze, gestures, facial expressions, and interpersonal distance" (Bailenson, 2018). According to the eminent psychologist Paul Ekman, "actions speak louder than words." Nonverbal behavior is indeed largely automatic, meaning that very few people can consistently regulate subtle micromovements and gestures such as sidelong glances or genuine smiles. In this sense, nonverbal data are uniquely telling and crucial for advertising, politics, and in general for identifying a person and their patterns.

Nonverbal behaviors are fundamental to creating immersive experiences via XR technologies because they only work if the system measures body and eye movements to make the environment respond accordingly. For example, in VR, people turn their physical heads around to make eye contact with other virtual reality users, use their legs to walk in the physical room to get across a virtual room and move their physical arms to grasp virtual objects. These tracking

data can be recorded and stored for later examination. In 2018, commercial systems typically tracked body movements 90 times per second to display the scene appropriately, and high-end systems recorded 18 types of movements across the head and hands. Consequently, spending 20 minutes in a VR simulation leaves just under two million unique recordings of body language (Bailenson, 2018). The largest part of the data collected can be included in three categories: hand tracking, eye tracking, and gait tracking, with the last two being the most interesting in terms of inferred data.

Eye-Tracking

Eye-tracking enables sophisticated data collection to tap into non-conscious processes governed by our biases and preferences. A combination of gaze position, pupil dilation, eye blinking rate, vertical and horizontal eye movements, and various other characteristics make eye-tracking an attractive tool for qualitative and quantitative research.

Eye-tracking allows for the quantification of various key metrics that can be utilized to gain new insights and discoveries. Some of the common metrics that can be derived from eye-tracking include fixations and gaze points, heatmaps, Areas of Interest (AOI), Time to First Fixation, Time spent (Dwell time), Ratio, Fixation sequences, Revisits, First Fixation Duration, Average Fixation Duration (*Understanding Human Behavior - A Physiological Approach - IMotions*, 2016).

Embedded eye-tracking sensors within these XR devices hold the keys to understanding the complex cognitive processes used during combat scenarios, allowing the simulation training to include cognitive performance metrics to better understand and improve the physical and mental responses of warfighters. Visual Recognition and Identification can be used to understand cognitive performance gains and the overall effectiveness of warfighters' training. Analyzing and evaluating the cognitive load and functions as well as characterizing deterioration over time provides insight into understanding and improving proper decision-making.

Scanning Pattern Proficiency and inferences gained from eye-tracking technologies provides insight into the cognitive effects of long-duration combat missions, executing multiple operations in one day, and recovering from tactical tasks. Military leadership can better assess the risk of accidents and fatalities under these conditions using the technologies and available data. Still, there needs to be a baseline understanding of what happened to all the data collected during these missions and exercises and where the data collection or sharing puts humans at risk. Massive unchecked use of such technologies involving the use of BID increases the dangers to privacy, civil liberties, and free agency. This data in the hands of an adversary represents a significant national security concern.

Gait Tracking

Gait analysis is "the systematic study of animal locomotion, more specifically the study of human motion, using the eye and brain of observers, augmented by instrumentation for measuring body movement, body mechanics, and the activity of the muscles" ("Gait Analysis," 2021). There are numerous uses for gait analysis, both good and bad and now many XR devices inherently collect data for gait analysis.

Traditionally conducted using external cameras, gait tracking has been very useful in diagnosing illnesses and issues that affect people's ability to walk. Having this data available, especially data that can show a change in gait of soldiers wearing combat loads over existing limits, and without the use of an elaborate camera setup would help military planners better understand the consequences of adding additional weight to the already overloaded ground Soldier or Marine. Another way the data is valuable is in trend analysis. When collected over time, the data can show changes in an individual's health and can aid in proving service-connected disabilities for ground troops and speed up the Veterans Administration (VA) disability rating process. While camera systems are more reliable, instrumented floors and inertial measurement units (IMUs) consisting of accelerometers, gyroscopes, and magnetometers, are contributing data for analysis, and are becoming increasingly more reliable as gait analysis tools. In the case of XR devices, IMUs are mostly standard equipment used to help the system know where a user is looking, how fast they are moving, etc.

Gait data can also be very useful in accurately identifying people. Termed gait biometrics, studies have shown that the level of accuracy of identifying people is between 80-100% (Connor & Ross, 2018). The same study also pointed out that gait is hard, if not impossible to mimic, making it a potentially viable biometric authentication method. While an individual's gait can be temporarily faked (i.e., adding a fake limp), another person trying to accurately imitate another's genuine gait is near, if not completely, impossible. Because of this, gait may make a great biometric authentication method, especially for wearables and XR devices.

For all the good that gait biometrics can do, it can also be used as a surveillance tool. The Associated Press (AP) reported in 2018 that China has developed the capability to identify an individual up to 165 feet away, and without having to see their face. The CEO of one company working on the capability has even been quoted as saying "You don't need people's cooperation for us to be able to recognize their identity..." (Kang, 2018)

Intersections of XR and Other Emerging Technologies

XRSI has been observing the technology space, specifically, the intersections with other emerging technologies such as 5G, BCI, AI, haptics, and multi-sensory. While new intersections continue to emerge and bring novel privacy and safety risks along with them, AI and BCI are the most concerning due to the excessive but necessary data collection.

Artificial Intelligence (AI)

Artificial intelligence is the study/domain of problem-solving, pattern recognition and developing systems armed with the intellectual characteristics of humans, such as the ability to reason, discover meaning, generalize, or learn from experience. Currently, the AI-based systems identify objects but do not fully understand the context of inter-object relationships, whether it is hierarchical or flat, or interdependence of objects and their relationship with one another influencing events in real-time. The convergence of AI and XR will enable the next stage of assigning behaviors and understanding events by linking real-world and virtual objects. This intersection of two evolving domains not only combines the associated risks but also exposes novel risks.

XR stakeholders must prepare in advance to handle the evolving risks brought on by the convergence of new technologies such as AI and XR. While these technology convergences can be used to mitigate some of the emerging risks, they also introduce threats of their own. Liu et al. stated that "machine learning models, especially deep neural networks, have been recently found to be vulnerable to carefully-crafted input called adversarial samples." (Liu et al., 2020) An example of an adversarial sample would be someone posting a picture on a stop sign that changes it to a different sign with the intent of crashing a self-driving car. This same kind of problem must be taken into consideration during the building process of any new XR system and experience. Likewise, just as humans have both conscious and unconscious biases, many of these artificial intelligence systems learn the same biases that humans have regarding gender, race, and attitudes toward people's differences. For example, gendered language can be introduced into these systems because machine learning models are only a reflection of the world in which we live (Hoyle et al., 2019). Recent innovations have shown that these biases can be mitigated proactively by conducting new research specifically to address these concerns.

In AI-driven XR ecosystems, identification of an individual's head and hand movements can lead to health inferences and identification of the mental state and health conditions that otherwise remain confidential. Unchecked and excessive data collection about military members can negatively impact their long-term mental health and can even potentially be the basis for discrimination regarding performance and promotions.

Brain-Computer Interface (BCI)

Sometimes called neural-control interface (NCI), mind-machine interface (MMI), direct neural interface (DNI), or brain-machine interface (BMI), a Brain-Computer Interface is a direct communication pathway between an enhanced or wired brain and an external device (*Research & Standards–The CyberXR Coalition 2021*, n.d.). A BCI allows for bidirectional information flow opening up the possibility for "writing back" to the brain. BCIs are often used for researching, mapping, assisting, augmenting, or repairing human cognitive or sensory-motor functions.

BCI translates a user's brain activity directly into a computer control signal to activate a computer or other external devices. The brain signals are usually measured using electroencephalography (EEG) and processed by neural interfaces. A BCI in XR is ideally situated to supplement and improve conventional XR modalities, by both expanding the actions possible with XR, and providing a more flexible and immersive user experience overall. (Scherer et al., 2010)

Interaction modalities of XR technologies are improving rapidly. Prevailing interaction methods such as hand gestures and voice recognition continue to evolve to meet the needs of XR environments, beginning with performing common tasks (e.g., object selection, menu navigation, and others). An ideal interaction method that robustly and naturally translates a user's intention into 2D and 3D environmental controls is what a direct Brain-computer interface (BCI) system is ideally situated to accomplish. BCI technology provides a solution to maximize XR's potential, affording users real-time mental selection via dry electroencephalography (EEG). The risks and benefits of these systems depend

mainly on their level of invasiveness. This technology is by nature intrusive, and privacy and security risks must be assessed and mitigated as they emerge.

Unintended Consequences and Human Risk

As already discussed, the unprecedented amount of individuals' data collected and processed by XR devices is redefining the concept of anonymity. The amount of first-hand and inferred information that can be extracted from data can create unique and unmistakable patterns that permit identifying an individual no matter the anonymization of their data. Furthermore, data in aggregate can begin to show likes and dislikes, and people with access to the data can use that data – and the intimate (literally) "in your face" access to manipulate a person's behavior.

The design and color choices, button placement, and app behavior all lend themselves to encouraging people to act in a specific way. For instance, Disney's subsidiary Pixar Animation Studios uses specific colors at specific times to elicit emotions in movie scenes (Rogers, 2021). App developers strategically place (or hide) buttons and menu items to steer you down a specific path. Taken a step further, it is feasible for someone to combine these techniques above with personal information about an individual, injecting that into a visual environment for nefarious reasons. This injection of personal information can be timed and placed correctly to elicit emotions, foster trustworthiness, and ultimately sway decision a major decision (e.g., a Presidential election) or to create AI-based Agents capable of extracting sensitive information out of unwitting Service-members (Scharlat, 2007).

THE XRSI FRAMEWORK

The XRSI Privacy and Safety Framework also referred to as "The XRSI Framework" is a free, globally accessible baseline rulebook curated by the XR Safety Initiative, designed to provide the technical foundations for the regulatory guidance and build trust in XR ecosystems.

The XRSI Framework is a community-built, agile, iterative, and functionality-driven approach to safety and privacy in XR. The development process is managed by a steering committee composed of multidisciplinary professionals, from cybersecurity to medical experts, from developers to child safety advocates, and more. Input was solicited and received from over 1000 experts from across the world. Figure 2 further identifies the locations and areas of expertise of XRSI's contributors and advisors. This cohort is still active in evolving the Framework.

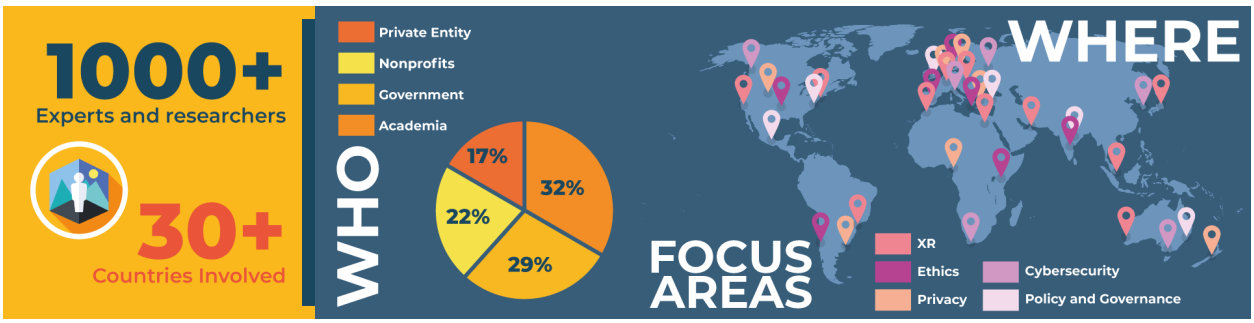

Figure 2. The XRSI Framework includes contributions from over 1,000 industry, academic, and government experts worldwide.

The main focus areas are separated into discrete functions, and each function has specific controls. For instance, the assessment area (Assess-AS) is separated into Mapping and Analysis (AS.MA) and Risk Assessment (AS.RA) functions. Within the AS.RA function, the framework proposes many different controls, such as "AS.RA2: Identify and Assess informational and financial risks that potentially impact individual privacy and safety (e.g., social engineering, tracking payment, and expenditure as well as other privacy-related risks)" or "AS.RA4: Identify and Assess datasets and input/output mechanisms for biases against humans, specifically vulnerable populations and underrepresented communities." The list of functions and controls is constantly evolving, with each iteration made public.

The XRSI Framework creates a baseline set of standards, guidelines, and best practices that are regulation-agnostic. It includes privacy and safety requirements drawn from the GDPR, National Institute of Standards and Technology (NIST) guidance, FERPA, COPPA, and a few other evolving laws. The XRSI Framework is a living document designed to quickly adapt to new requirements, regulations, and approaches as soon as they come into effect. The framework is not a law or standard; it is a free tool that is continually evolving.

Overview of The XRSI Framework

The XRSI Framework is composed of three layers as indicated in Figure 3. These layers include focus areas, a set of functions, and granular controls. Each component reinforces how organizations and institutions achieve privacy and safety goals through aligning strategy, roles, and responsibilities.

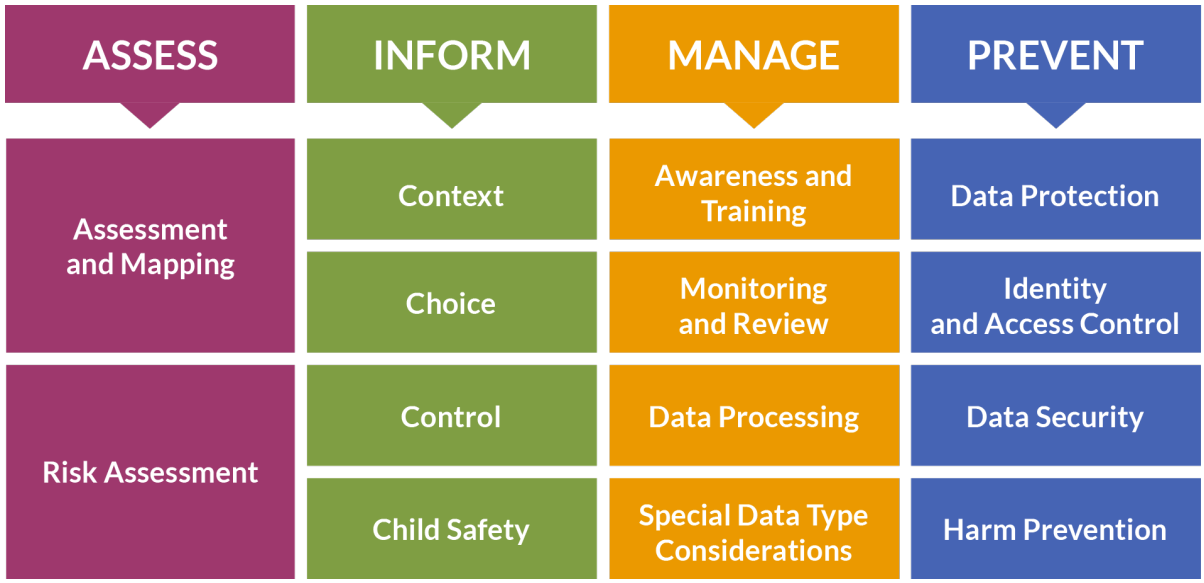

Figure 3. Areas of Work and Subcategories in the XRSI Framework. Reprinted from The XRSI Privacy Framework version 1.0, XRSI, 2020, p.9. Reprinted with permission

Focus Areas-The Focus Areas provide a foundation for outlining the scope of work that enables institutions to incorporate human-centric privacy and safety by design and default into their academic practices and development for XR.

Functions-Functions are the subcategories for outlining groups of privacy and safety-focused activities tied to the focus areas and corresponding granular controls.

Controls-Controls are the activities that are carried out to achieve specific privacy and safety outcomes about operations. They provide a set of results and help support the achievement of the intended outcome in each of the focus areas.

Within The XRSI Framework, four foundational focus areas have been identified for building inclusive, safe, private, and responsible XR systems for all. The first area, Assess, is crucial for the identification of the risks. The second and third areas, Inform and Manage, collectively address organizations' capacity to build trust by informing individuals and managing privacy and safety risks within the ecosystems. The fourth focus area, Prevent, enhances safety by outlining proactive controls needed for privacy and safety within the XR ecosystems.

The XRSI Framework Focus Areas

The XRSI Framework is a voluntary tool for managing privacy and safety risks in XR. It is intended to serve organizations and institutions of all sizes. The XRSI Framework is strategically designed to be compatible with existing domestic and international legal and regulatory regimes and usable by any type of organization to enable

widespread adoption. It explicitly considers key regulations such as GDPR, CCPA, COPPA, FERPA, and a few others, as previously mentioned.

The XRSI Framework's purpose is to help organizations and institutions utilizing XR technologies, manage privacy and safety risks.

When used as a risk management tool, The XRSI Framework can help an institution build trust, achieve transparency, and create accountability during research, development, and innovation while minimizing unintended consequences for individuals or collectives. It can also help organizations and institutions using XR assess the scope of their impact on individual and collective rights, facilitate those rights, and comply with various international and state regulations.

The XRSI Framework remains flexible, complements existing risk strategy, and leaves its application decision to the organization or institution itself. For example, the US Department of Defense already has a robust risk management program. Still, it may use the framework to analyze novel privacy and safety risks that the introduction of XR may create. Likewise, service components can use the focus areas and functions as a reference to understand and communicate privacy and safety needs and expectations to its materiel solutions providers. Government contractors and other support organizations can also use the XRSI Framework to understand their responsibility toward individual privacy and safety while building an experience, application, or platform suited for operational use by warfighters.

Caution must be taken when using the XRSI Framework as a compliance tool because that is not its intended function. There is no inherent expectation for organizations to "comply with The XRSI Framework." Instead, they should use it as the baseline measure to optimize trust and safety efforts in minimizing risks within the XR Data Processing Ecosystem.

AS-Assess

This area is focused on developing an understanding to manage privacy and safety risks for individuals and organizations arising from data processing and collection. Assessment of risks is fundamental to building responsible and human-centric XR ecosystems. Develop a comprehensive understanding of the organizations' risks associated with data collection, processing, analysis, and its impact on the individuals. This enables an organization to understand the business environment in which it operates and prioritize risk mitigation accordingly.

Mapping and Analysis: Data processing systems, products, and/or services are mapped and analyzed for potential privacy and safety risks.

Risk Assessment: Privacy and safety risks related to the organizations are assessed to determine the impact on their operations, mission, and functions taking into account other risk factors, including human, societal, informational, financial, and legal risks.

While performing a risk assessment for the XR ecosystems, the following questions, taken directly from the XRSI Framework, can help organizations assess their privacy data exposure vulnerabilities (*The XRSI Privacy Framework* (2020)):

- What are the various types of data required by the XR platform, service, or app?
- What are the various types of data being collected, processed, and shared?
- What is the legal basis for storing personal and sensitive XR data?
- Which third parties will the data be shared with, and how will they be processing the data??
- What processes are in place to communicate to customers, collaborators, and regulators what data is being collected and why?
- What processes are in place to ensure the data is stored securely?
- What processes are in place for timely responding to a data breach or any privacy incident?
- What is the data collection pipeline?
    - What is collected by the device?
    - What is stored locally on the device?
    - What data is shared with:
        - Other individuals?
        - Third-party applications?

        - Other companies?
    - What data is stored?
        - On-device?
        - Distributed to other individuals?
        - On an edge cloud?
        - On a remote cloud?
        - How long will the data be retained?
        - Will the personal and sensitive XR data be encrypted, de-identified, obfuscated, and/or aggregated when storing or processing?

PR-Prevent

The strong cybersecurity background within the XR Safety Initiative drives the Framework development with a clear understanding: preventing threats is always better than responding to them. More specifically, when this principle is plugged into the immersive domain, the prevention activities must be extended to a wide set of functions, such as:

- Data Protection
- Identity Management, Authentication, and Access Control
- Data Security
- Online Harm Prevention
- Child Safety

Given the context, the nature, and the purposes of the processing of individuals' data (and the amount of personal data processed), XR providers should minimize any potential data/information exposure. Some XR providers do not ensure the adoption of certain data security measures, such as encryption or pseudonymization (standard practice in more traditional digital communication means such as instant messaging apps). Furthermore, certain XR systems also rely on third-party services or apps which do not appear to implement suitable security standards. XR providers need to implement adequate policies and security measures (e.g., physical security of data, physical security of facilities/personnel, network security, system hardening, password security, endpoint protection, patch management, remote access, etc.) to satisfy legal requirements. Many commentators are currently pushing for the identification by governments (or even XR providers' self-regulation bodies) of specific XR minimum security standards (e.g., SANS, NIST, ISO, CIS). Such standards would help XR operators provide more secure products and services, thus fostering a wider deployment of XR solutions. Moreover, XR providers must ensure restoration of systems to ordinary operation as soon as possible; adequate business continuity and disaster recovery plans should accordingly be in place, also to address incident and security breaches.

Protection of children in the XR domain is a highly sensitive aspect: when it comes to minors, prevention is the only possible approach, because managing already performed threats could not be possible, and the harm could already be put into effect on a particularly vulnerable community.

MN-Manage

The XR industry is moving fast, and some of the products, devices, applications, and experiences are not XR-native. At the same time, new data processing tools and technologies are constantly released. This means that not every possible risk can be prevented, and the preventive approach must be integrated with a clear understanding of how to manage existing threats.

The MN area focuses on organizational-level activities such as establishing organizational values and policies, addressing safety, privacy, legal and regulatory requirements, managing organizational risk tolerance to allow an organization to focus and prioritize its efforts, consistent with its overall risk management strategy, and business needs.

More specifically, the area has been divided into four critical functions, allowing organizations to understand and manage risks in the different phases of the product and data lifecycle:

- Awareness and Training: The organization's workforce and third parties engaged in data processing are provided privacy awareness education and are trained to perform their privacy-related duties and responsibilities consistent with related policies, processes, procedures, and agreements, and organizational privacy values.
- Monitoring and Review: The policies, processes, and procedures for ongoing review of existing controls and risks to develop an effective plan of action.

- Data Disclosures (Breach Notification): Data breach notification requirements were designed to empower consumers and shame companies into improving their data security practices. The precise contours of when/where/how individuals must be informed varies widely based on jurisdiction.
- Data Processing Ecosystem Risk Management: The organization's priorities, constraints, risk tolerance, and assumptions are established and implemented to support risk management decisions within the data processing ecosystem.
- Special Data Type Consideration: The organization has established and implemented the processes to identify, assess, and manage privacy risks related to special data types. The organization's priorities, constraints, risk tolerance, and assumptions are established and used to support risk decisions associated with sensitive data that may put humans at risk.

IN-Inform

Articles 13 and 14 of the European Union's law 2016/679, GDPR, clearly states that any processing of personal data should be lawful, fair, and transparent. It should be clear and transparent to individuals that personal data concerning them are collected, used, consulted, or otherwise processed, and to what extent the personal data are, or will be, processed. The right to be informed, under Articles 13 and 14 GDPR, is a key part of any organization's obligations to be transparent.

The principle of transparency requires that any information or communication relating to the processing of personal data is easily accessible and easy to understand, and that clear and plain language be used. Any information addressed to the public or the data subject must be concise, easily accessible, and easy to understand, and that clear and plain language and, additionally, where appropriate, visualization be used.

A New Definition of Personal Data

XR expands the definition of personal information that must be protected, including BID. The people whose information is tracked, collected, used, and shared, must be granted the right to be informed on the full data pipeline.

Given the potential immersion of XR experiences and the breadth of sensitive information available to XR hardware, informed consent is critical. This concept becomes especially important when the data tracking, collection, usage, and sharing, involves vulnerable individuals or categories, such as minorities, persons with disabilities, or children. The right to informed consent includes ensuring age-appropriate design and awareness for parents to increase child safety.

Following these principles, and based on the fact that individuals should be made aware of risks, rules, safeguards, and rights about the immersive experiences, the XRSI Framework is enabling easy-to-implement mechanisms to support individuals' awareness when immersing into XR technologies. Even more broadly, the individuals must be put in the position of always being the real owners of their data, meaning that they must be able to have full understanding and control over any piece of information that is tracked, stored, processed, and shared.

The first crucial element an individual must be provided is context, meaning that any product or experience must clearly communicate what kind of data is managed, how the whole data funnel is shaped, and how long the data will be in every phase of this process.

In the XRSI Framework approach, the context is the minimum baseline for offering individuals a choice on what to share, how to share it, and with whom. This is not only a matter of development and programming (for instance, implementing a clear User Interface to provide opt-in and opt-out choices), but a cultural switch to a privacy-first and human-centric privacy-by-design approach, respecting privacy as a fundamental individuals' right.

Where To Find More Information

For more information about the Framework or how you can get involved, visit XRSI at www.xrsi.org or contact the authors directly via email.

CONCLUSION

XR is a new and emerging domain and no regulations or guidelines currently exist to mitigate safety and privacy risk. This creates a massive gap in terms of our ability to trust in the technologies used to enable virtual experiences. Various ways of data collection such as eye and GAIT tracking combined with various technological intersections such as AI and BCI, multiply and significantly increase risks to humans, and society in general. While we await regulations, various technologies converge, and the use of XR within critical sectors (e.g., military and healthcare) warrant us to adopt a self-governance model to address expected and unintended consequences.

With the recent acquisition of the Integrated Virtual Augmentation System (IVAS) by the Army, the military is signaling a move from using XR as a training tool to an operational device, bringing with it additional considerations. New opportunities and capabilities bring new risks. Traditionally, war preparations have focused on the physical and strategic actions essential to ensuring success on the battlefield. However, with the adoption of XR technologies into military operations, understanding the safety and privacy risks to the warfighters, organizations, and personnel involved is paramount to national security.

The XRSI Framework acts as the tool for assessing risks, informing individuals around the associated risks, managing the risks, and most importantly preventing harm with proactive technical, administrative, and physical safety and privacy controls. Adopting the XRSI Framework to establish baseline privacy and safety of XR ecosystems will allow us to continue down the path of responsible and ethical innovation to safeguard national security.

REFERENCES

Bailenson, J. (2018). *Experience on Demand*. W. W. Norton & Company.

Connor, P., & Ross, A. (2018). Biometric recognition by gait: A survey of modalities and features. *Computer Vision and Image Understanding*, *167*, 1–27. https://doi.org/10.1016/j.cviu.2018.01.007

Gait analysis. (2021). In *Wikipedia*. https://en.wikipedia.org/w/index.php?title=Gait_analysis&oldid=1014320399

Hoyle, A. M., Wolf-Sonkin, L., Wallach, H., Augenstein, I., & Cotterell, R. (2019). Unsupervised Discovery of Gendered Language through Latent-Variable Modeling. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1706–1716. https://doi.org/10.18653/v1/P19-1167

Kang, D. (2018, November 6). *Chinese "gait recognition" tech IDs people by how they walk*. https://apnews.com/article/china-technology-beijing-business-international-news-bf75dd1c26c947b7826d270a16e2658a

Liu, N., Du, M., Guo, R., Liu, H., & Hu, X. (2020). Adversarial Attacks and Defenses: An Interpretation Perspective. *ArXiv:2004.11488 [Cs, Stat]*. http://arxiv.org/abs/2004.11488

O'Brolcháin, F., Jacquemard, T., Monaghan, D., O'Connor, N., Novitzky, P., & Gordijn, B. (2016). The Convergence of Virtual Reality and Social Networks: Threats to Privacy and Autonomy. *Science and Engineering Ethics*, *22*(1), 1–29. https://doi.org/10.1007/s11948-014-9621-1

Research & Standards–The CyberXR Coalition 2021. (n.d.). Retrieved June 13, 2021, from https://cyberxr.org/research-standards/

Rogers, A. (2021, April 29). *How Pixar Uses Hyper-Colors to Hack Your Brain | WIRED*. https://www.wired.com/story/how-pixar-uses-hyper-colors-to-hack-your-brain/

Scharlat, J. (2007). *Intelligent Virtual Environment Agents (IVEAs): Conducting Information Operations in Virtual Environments*. Naval Postgraduate School.

Scherer, R., Chung, M., Lyon, J., Cheung, W., & Rao, R. P. N. (2010). Interaction with Virtual and Augmented Reality Environments using Non-Invasive Brain-Computer Interfacing. *Proc. of the 1st International Conference on Applied Bionics and Biomechanics (ICABB-2010)*, 1–4.

The XRSI Privacy Framework (2020). Retrieved June 13, 2021, from https://www.gao.gov/assets/710/702715.pdf

Understanding Human Behavior–A Physiological Approach–IMotions. (2016, March 28). Imotions Publish. https://imotions.com/blog/understanding-human-behavior/

U.S. Army to use HoloLens technology in high-tech headsets for soldiers. (2021, June 8). Transform. https://news.microsoft.com/transform/u-s-army-to-use-hololens-technology-in-high-tech-headsets-for-soldiers/