Enabling Safer Augmented Reality Experiences: Usable Privacy Interventions for AR Creators and End-Users

Shwetha Rajaram
University of Michigan
United States
shwethar@umich.edu

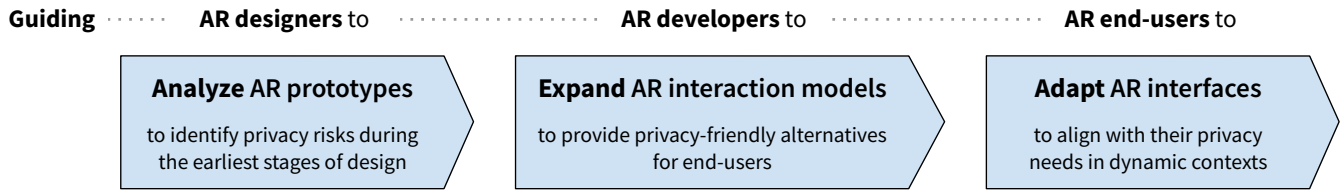

Figure 1: Overview of Research Agenda to equip AR designers, developers and end-users with a privacy mindset.

ABSTRACT

Augmented reality (AR) is approaching everyday usage, but poses novel privacy concerns for end-users and bystanders due to how AR devices capture users and process physical environments. To enable the benefits of AR while balancing privacy goals, my dissertation develops tools and frameworks to guide AR creators and users to address privacy risks that can arise with AR. First, I explore how to enable AR designers to interactively analyze potential risks in their prototypes through implicit threat modeling within AR authoring tools. Next, through elicitation studies with AR and privacy experts, I contribute frameworks to expand AR interaction models with privacy-friendlier alternatives to traditional AR input, output, and interaction techniques. Lastly, I develop a suite of AI-enabled Privacy Assistant techniques to raise users' awareness of privacy risks and help them adapt AR interfaces accordingly. Ultimately, my dissertation promotes an AR ecosystem with privacy at the forefront by equipping AR creators and users with a privacy mindset.

CCS CONCEPTS

• **Human-centered computing** → **Mixed / augmented reality**; **Accessibility**; • **Security and privacy / Privacy protections**;

KEYWORDS

augmented reality, usable privacy, authoring tools, elicitation

ACM Reference Format:
Shwetha Rajaram. 2024. Enabling Safer Augmented Reality Experiences: Usable Privacy Interventions for AR Creators and End-Users. In *The 37th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '24), October 13–16, 2024, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3672539.3686708

UIST Adjunct '24, October 13–16, 2024, Pittsburgh, PA, USA

ACM ISBN 979-8-4007-0718-6/24/10
https://doi.org/10.1145/3672539.3686708

1 INTRODUCTION

Ongoing advancements in AR technologies will soon enable end-users to leverage AR to support their daily tasks across a variety of contexts. Commercial and research applications are establishing the benefits of AR across various domains, such as enhancing communication in telepresence systems [25], providing in-situ guidance to healthcare professionals [14], and empowering instructors to teach spatial concepts in new ways [23, 27]. Meanwhile, innovations in AR hardware and sensing techniques are making AR devices more suitable for everyday use, with trends towards lightweight AR glasses[1] and techniques for adapting AR interfaces to various physical environments [7, 15].

However, AR can pose novel privacy concerns to end-users and bystanders, due to how AR devices capture users' interactions and process their physical surroundings [9, 29]. For example, biometric information (e.g., speech or motion data) can be used to infer users' identities or health conditions [13, 19]. Environmental sensing techniques (e.g., object detection and 3D reconstruction) can capture bystanders without their awareness or consent [10, 29]. As AR devices are increasingly used in always-on scenarios, these risks may be exacerbated and require users to frequently adjust their usage of AR to maintain their desired level of privacy.

To mitigate privacy risks across the AR development and usage lifecycle, my research develops tools and frameworks that equip AR creators and end-users with a privacy mindset. I envision an AR ecosystem where designers and developers not only strive to create novel AR interactions, but are also skilled at integrating privacy best practices into their designs. In parallel, I want to empower end-users to make informed decisions about their privacy and exert granular control over the configuration of AR interfaces to support their privacy needs.

Towards these goals, I outline the three stages of my research agenda (Fig. 1) in the following sections:

[1] *Project Aria* glasses (Meta): https://www.projectaria.com/glasses/

(1) Developing threat modeling tools for AR designers to **analyze privacy risks** and brainstorm mitigation techniques directly within their prototypes [28];
(2) Eliciting design frameworks for AR developers to **expand their interaction models** with more privacy-friendly interaction techniques for users and bystanders, through studies with AR and privacy experts [26];
(3) Developing a Privacy Assistant to help users understand risks in dynamic contexts and **adapt AR interfaces** to better serve their privacy needs.

Ultimately, my dissertation will contribute to an AR ecosystem with privacy considerations at the forefront, by embedding privacy expertise into AR creators' and users' workflows and bridging isolated research in the AR and privacy communities.

2 ANALYZING RISKS: THREAT MODELING WITHIN AR PROTOTYPING TOOLS

Promoting *Privacy by Design* [6] across the AR ecosystem requires **disseminating expertise in identifying privacy threats among the population of AR designers and developers**. One systematic method for analyzing privacy risks is *threat modeling*, which involves brainstorming how a set of critical threats could manifest for a specific technology and prioritizing which threats to mitigate, based on their severity and plausibility [30].

We identify two key challenges with embedding privacy expertise in AR designers' workflows. First, a significant population of novice AR designers are being empowered to create and deploy AR applications through authoring tools that lower the technical barrier to entry [21, 22] (e.g., by providing no-code techniques to leverage multimodal AR sensing capabilities [11, 16]). However, these designers may lack the formal training in both AR and privacy that is required to understand threats, and today's AR authoring tools fall short of raising designers' awareness of the potentially harmful impacts of their AR interaction designs.

Second, existing privacy educational tools are separate from designers' workflows, thus limiting their adoption. Our research takes inspiration from how prior tools provide simpler abstractions to teach designers about privacy threats (e.g., exploring threat models through ideation cards [1, 18], visual depictions of threats in privacy comics [31]). To support rapid iteration towards safer AR interactions, **we explored how to integrate an implicit threat modeling process within AR prototyping tools**.

Research Overview: We developed Reframe [28], an AR storyboarding system that enables designers to interactively explore potential threats directly within their prototypes. Our key innovation to make privacy threats more visible and understandable for novice AR designers is a *character-driven analysis approach* that personifies threats as bystander and adversary characters. Characters are automatically inserted in the storyboard to demonstrate when and where threats could occur, based on the user's location and interaction modalities. For example, our *Graffiti Spammer* character prompts consideration of sharing policies and access control in AR (Fig. 2). Appendix A.1 shows an overview of Reframe's authoring interface and characters.

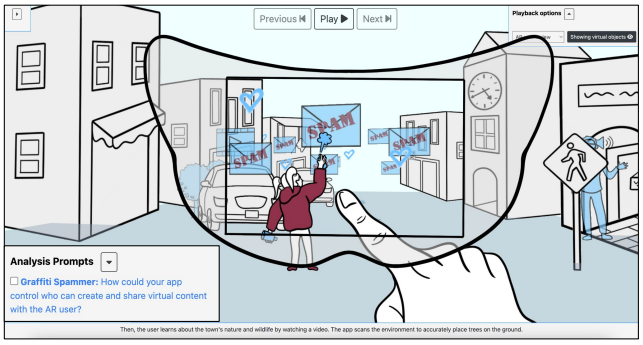

Figure 2: Threat Modeling with Reframe. Reframe enables designers to author and preview storyboards as a sequence of AR application states in a 2.5 simulation environment, using either hand-held or head-worn AR. Then, our character-driven analysis tools insert personified representations of threats and analysis prompts to help designers brainstorm mitigation techniques. Designers can review their threat modeling via screen & audio recording tools.

We evaluated Reframe in two steps: *(1)* investigating how novice AR designers' storyboards evolved when using Reframe to identify and mitigate threats; *(2)* assessing the quality of the designers' threat modeling through a review with privacy experts. Designers expressed that observing characters' interactions within 2.5D simulation environments provided effective scaffolding for brainstorming threats. The privacy experts found the designers' threat modeling as enabled by Reframe to be of good quality, compared to their own baseline analysis. To further iterate on mitigation techniques to cover a wider range of potential threats, the experts saw promise in Reframe to facilitate co-design between AR creators and privacy professionals.

Outcomes and Impact: Since publishing this work at UIST 2023[2], we made Reframe available as an open-source tool for interaction designers. Since 2022, my advisor, Prof. Michael Nebeling, and I have been using Reframe to teach graduate students about privacy considerations in Introductory AR/VR Design and Development courses at the University of Michigan School of Information.

3 EXPANDING AR INTERACTION MODELS: PRIVACY ADAPTATION TECHNIQUES

Reframe makes progress towards enabling designers without formal training in privacy to identify potential threats; however, developing appropriate techniques to mitigate these threats remains a challenge. AR users' perceptions of risk may vary based on location [29], types of data collected [13], and concerns for bystanders' privacy [2, 10]. For AR developers, this raises a **need to provide a range of techniques for users to adapt AR interfaces to meet their privacy needs**. For example, if users are not comfortable using 3D reconstruction to the full extent in their personal home

[2]**Reframe paper:** https://shwetharajaram.github.io/paper-pdfs/reframe-uist23.pdf; **Reframe video demonstration:** https://youtu.be/kGkDWZSr_2k

(Fig. 3), they could choose to only reconstruct specific surfaces at a coarse-grained quality to disclose less detail.

However, the landscape of privacy-preserving techniques for AR interfaces is not well understood. While the HCI community has explored adapting AR visual output to optimize for usability objectives (e.g., to improve reachability [4] or minimize users' cognitive load [17]), adapting AR interfaces with privacy as the objective can involve a wide range of AR input, output, and interaction modalities. For example, in personal environments, users may prefer marker-based tracking that only requires camera data, rather than marker-less approaches that take both camera and depth data as input. When using AR in public environments, AR users may prefer subtler interaction techniques (e.g., using microgestures over mid-air gestures) to prevent bystanders from inferring their activities.

As steps towards establishing a design space of privacy adaptation techniques, **my next set of projects contribute design catalogs to expand AR interaction models with alternate, privacy-friendly interaction techniques**.

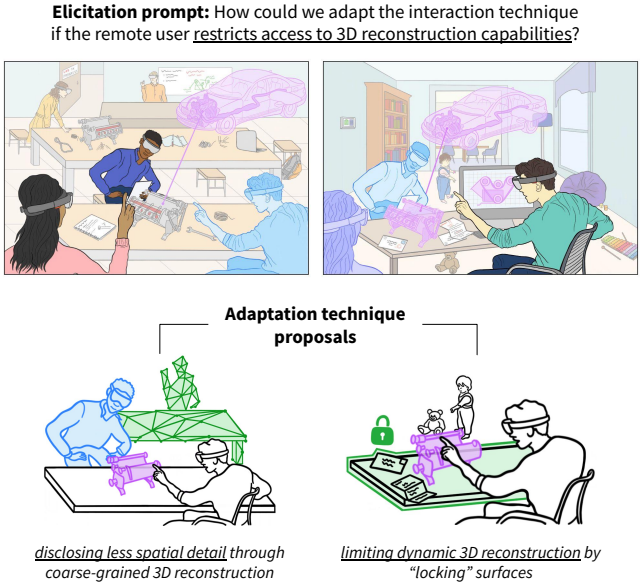

Figure 3: Scenario-based Elicitation Method. Based on sketches or prototypes of AR applications (e.g., engineering lab, navigation system), we prompted AR experts to redesign the AR interactions if end-users chose to restrict access to a core AR sensing capability, while still achieving the functionality to the highest possible extent.

Research Overview: We conducted two elicitation studies with researchers with expertise in AR and privacy to explore the design space of privacy adaptation techniques for AR. *User-driven elicitation* is an HCI method for designing interaction techniques where users are prompted with a system effect (e.g., sharing AR content), then produce interaction proposals to accomplish that effect [20, 32]. We extended traditional elicitation studies to holistically consider usability and privacy goals by *(1)* incorporating AR usage scenarios as a basis for analyzing context-dependent privacy risks (Fig. 3); *(2)* crafting elicitation prompts around a threat model or permission model; *(3)* facilitating co-design between AR developers and privacy experts to study the interplay of usability and privacy goals.

Our first study with 12 AR and 4 privacy experts focused on how to extend traditional AR interaction techniques (e.g., gestures, voice commands) to provide fine-grained access control over AR content in multi-user scenarios [26]; Appendix A.2 shows our resulting design catalog. Our second study with 10 AR experts focused on how to accomplish traditional AR sensing techniques (e.g., 3D reconstruction, object detection) in more privacy-preserving ways. This resulted in a catalog of 62 adaptation techniques, which we operationalized into a web-based visualization tool[3] to help AR developers analyze and expand their interaction designs with more granular controls for end-users (Fig. 4).

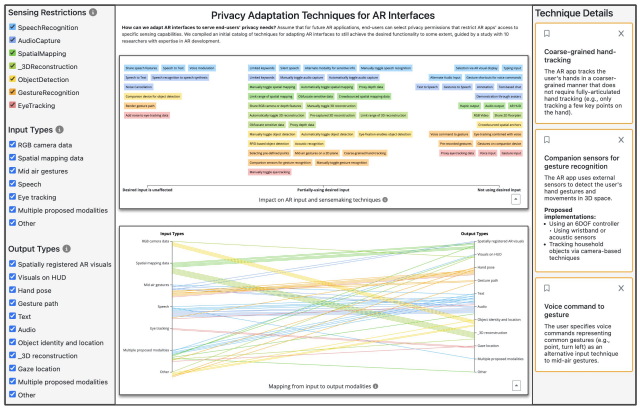

Figure 4: Visualization Tool for Privacy Adaptation Techniques. We created a visualization tool to help AR developers navigate the design space of possible techniques for adapting AR input, output, and interaction techniques to serve end-users' privacy needs. The techniques are organized along design dimensions (e.g., system-driven vs. user-driven adaptations) to help developers weigh the usability and privacy implications of different techniques.

Outcomes and Ongoing Work: We published insights from our first elicitation study at CHI 2023[4]. A submission around our second study is under review. We plan to deploy our design catalogs and visualization tools tool to support AR developers in the future.

To further study the tradeoffs between usability and privacy when applying privacy adaptation techniques, I will supervise a master thesis project from September 2024–April 2025. We plan to conduct studies around a functional AR remote assistance application, where pairs of participants complete collaborative tasks with the remote assistance app and are prompted to choose adaptation techniques in response to privacy-oriented stimuli (e.g., a bystander entering their field-of-view). We will probe into participants' rationale for choosing specific adaptation techniques and study the impact on task performance.

[3] **Privacy Adaptation Techniques Visualization Tool:** https://www.youtube.com/watch?v=7ypc-KAi8Nw
[4] **Eliciting S&P-Informed Sharing Techniques for Multi-User AR:** https://shwetharajaram.github.io/paper-pdfs/elicitation-chi23.pdf

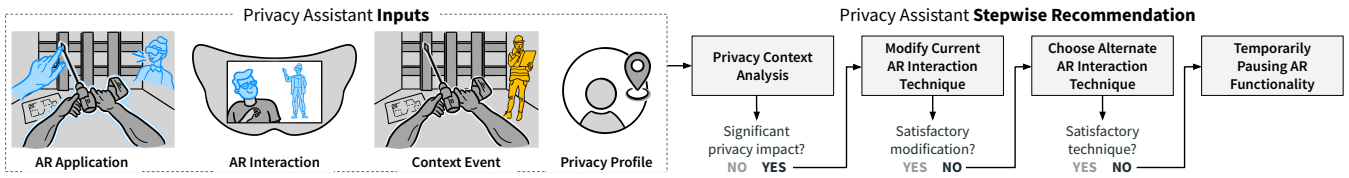

Figure 5: AR Privacy Assistant Workflow. Our LLM-enabled Privacy Assistant takes contextual data as input, such as the type of AR app, current AR interaction, change in the user's context that warrants a re-evaluation of their privacy needs, and a persona representing the user's privacy preferences [12]. Based on these inputs, the Privacy Assistant recommends how to adapt the AR interface to meet the user's privacy needs. This could range from minor modifications (e.g., reducing the sensing range) to more drastic changes to the interaction model (e.g., switching to an alternate AR sensing modality).

4 ADAPTING AR INTERFACES: PRIVACY ASSISTANT FOR AR END-USERS

Our design catalogs take important steps towards offering end-users privacy-friendlier alternatives to traditional AR interaction techniques. However, **further scaffolding is needed to make these privacy controls feasible for end-users to apply in everyday, always-on AR scenarios**. Weighing the impact of different AR interface configurations on both privacy and usability may pose challenges for users without technical expertise in AR. Further, as users increasingly leverage AR across dynamic contexts, maintaining their desired level of privacy would require them to frequently reassess their perceptions of risk and adapt AR interfaces accordingly. However, users already have limited *compliance budgets* to take privacy precautions for mobile and web interfaces [3], tending to abandon practices that require recurring effort [33].

To lower the barrier for everyday AR users to make informed privacy decisions, **we are exploring mixed-initiative approaches to anticipate and mitigate privacy risks through developing a suite of AI-enabled "Privacy Assistant" techniques**. We envision an intelligent agent that analyzes changes in AR users' contexts and makes stepwise recommendations to reconfigure AR interfaces to meet their privacy preferences (Fig. 5), by applying adaptation techniques from our design catalogs (Sec. 3).

We are conducting this research in two phases: *(1)* eliciting design guidelines for privacy recommendations through studies with privacy experts and investigating how to align today's LLMs to these guidelines; *(2)* exploring interaction techniques for end-users to communicate with AR Privacy Assistants (e.g., to specify privacy preferences or negotiate on AR adaptation techniques).

Ongoing Work: Investigating how to align an LLM-enabled Privacy Assistant to privacy experts' design guidelines. Formulating privacy recommendations for everyday AR involves competing design goals (e.g., balancing AR usability and privacy requirements, strictly aligning advice to users' privacy preferences vs. nudging users to be more privacy-conscious). Today's LLMs show promise for modeling human behavior [24] and weighing conflicting goals [8], but their sensemaking capabilities for technical concepts in AR and privacy are unclear. To elicit concrete design guidelines and assess the viability of using LLMs for privacy recommendations, we implemented an initial Privacy Assistant driven by GPT-4 (Fig. 5) to facilitate a study with privacy experts. Inspired by Find-Fix-Verify [5], we task privacy experts with reviewing the Privacy Assistant's output for different scenarios, annotating positive and negative aspects, and implementing revisions (Appendix A.3).

To investigate technical requirements for aligning LLMs with experts' design guidelines, we will use the experts' revised examples to improve our Privacy Assistant implementation via a variety of techniques (e.g., prompt engineering, fine-tuning, retrieval-augmented generation). To assess the benefits & limitations of each approach, we will conduct a comparative study with AR and privacy experts.

Future Work: Developing interaction techniques for AR users to leverage Privacy Assistants. Finally, I plan to develop interaction techniques that support users' communication with the Privacy Assistant and aid their privacy decision-making. Specifically, I am interested in *(1)* techniques to deliver privacy recommendations without distracting from the AR experience, *(2)* abstractions to explain the usability and privacy implications of AR adaptation techniques to non-technical users, and *(3)* design strategies to gracefully transition between adaptation techniques and minimize disruptions to users' current tasks. I plan to operationalize the suite of Privacy Assistant techniques into sample AR applications to evaluate their effectiveness in studies with end-users.

5 DISSERTATION STATUS & AIMS FOR UIST DOCTORAL SYMPOSIUM

I will start the 5th year of my PhD program in Fall 2024. I plan to propose my dissertation by February 2025 and complete my defense by December 2025. As such, the UIST Doctoral Symposium would come at an opportune time for me, enabling me to receive targeted feedback to improve the last few projects in my dissertation.

In particular, I would appreciate advice on how to strengthen the validity of the Privacy Assistant research, given that many aspects of the future AR ecosystem are under assumption (e.g., to what extent AR developers would expose underlying application information to enable a privacy analysis). I also look forward to mentorship on positioning my work to different academic and industrial research organizations as I prepare to go on the job market.

ACKNOWLEDGMENTS

I thank my advisor, Michael Nebeling, and all of my collaborators for their contributions to this research and support throughout my PhD. My dissertation work is supported by the Rackham Predoctoral Fellowship at the University of Michigan.

REFERENCES

[1] 2013. The Security Cards: A Security Threat Brainstorming Kit. https://securitycards.cs.washington.edu/

[2] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security, SOUPS 2018, Baltimore, MD, USA, August 12-14, 2018*. USENIX Association, 427–442. https://www.usenix.org/conference/soups2018/presentation/adams

[3] Adam Beautement, Martina Angela Sasse, and Mike Wonham. 2008. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 Workshop on New Security Paradigms, Lake Tahoe, CA, USA, September 22-25, 2008*. ACM, 47–58. https://doi.org/10.1145/1595676.1595684

[4] João Marcelo Evangelista Belo, Mathias N. Lystbæk, Anna Maria Feit, Ken Pfeuffer, Peter Kán, Antti Oulasvirta, and Kaj Grønbæk. 2022. AUIT - the Adaptive User Interfaces Toolkit for Designing XR Applications. In *The 35th Annual ACM Symposium on User Interface Software and Technology, UIST 2022, Bend, OR, USA, 29 October 2022 - 2 November 2022*. ACM, 48:1–48:16. https://doi.org/10.1145/3526113.3545651

[5] Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: a word processor with a crowd inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology, New York, NY, USA, October 3-6, 2010*. ACM, 313–322. https://doi.org/10.1145/1866029.1866078

[6] Ann Cavoukian. 2010. Privacy by Design: The 7 Foundational Principles. Revised: October 2010.

[7] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. 2021. SemanticAdapt: Optimization-based Adaptation of Mixed Reality Layouts Leveraging Virtual-Physical Semantic Connections. In *UIST '21: The 34th Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 10-14, 2021*. ACM, 282–297. https://doi.org/10.1145/3472749.3474750

[8] Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2024. Enhancing AI-Assisted Group Decision Making through LLM-Powered Devil's Advocate. In *Proceedings of the 29th International Conference on Intelligent User Interfaces, IUI 2024, Greenville, SC, USA, March 18-21, 2024*. ACM, 103–119. https://doi.org/10.1145/3640543.3645199

[9] Jaybie A. de Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2020. Security and Privacy Approaches in Mixed Reality: A Literature Survey. *ACM Comput. Surv.* 52, 6 (2020), 110:1–110:37. https://doi.org/10.1145/3359626

[10] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies. In *CHI Conference on Human Factors in Computing Systems, CHI'14, Toronto, ON, Canada - April 26 - May 01, 2014*. ACM, 2377–2386. https://doi.org/10.1145/2556288.2557352

[11] Ruofei Du, Eric Turner, Maksym Dzitsiuk, Luca Prasso, Ivo Duarte, Jason Dourgarian, João Afonso, Jose Pascoal, Josh Gladstone, Nuno Cruces, Shahram Izadi, Adarsh Kowdle, Konstantine Tsotsos, and David Kim. 2020. DepthLab: Real-time 3D Interaction with Depth Maps for Mobile Augmented Reality. In *UIST '20: The 33rd Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 20-23, 2020*. ACM, 829–843. https://doi.org/10.1145/3379337.3415881

[12] Janna Lynn Dupree, Richard Devries, Daniel M. Berry, and Edward Lank. 2016. Privacy Personas: Clustering Users via Attitudes and Behaviors toward Security Practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, May 7-12, 2016*. ACM, 5228–5239. https://doi.org/10.1145/2858036.2858214

[13] Andrea Gallardo, Chris Choy, Jaideep Juneja, Efe Bozkir, Camille Cobb, Lujo Bauer, and Lorrie Cranor. 2023. Speculative Privacy Concerns about AR Glasses Data Collection. *Proc. Priv. Enhancing Technol.* 2023, 4 (2023), 416–435. https://doi.org/10.56553/popets-2023-0117

[14] Danilo Gasques, Janet G. Johnson, Tommy Sharkey, Yuanyuan Feng, Ru Wang, Zhuoqun Robin Xu, Enrique Zavala, Yifei Zhang, Wanze Xie, Xinming Zhang, Konrad Davis, Michael Yip, and Nadir Weibel. 2021. ARTEMIS: A Collaborative Mixed-Reality System for Immersive Surgical Telementoring. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*. ACM, 662:1–662:14. https://doi.org/10.1145/3411764.3445576

[15] Jens Grubert, Tobias Langlotz, Stefanie Zollmann, and Holger Regenbrecht. 2017. Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality. *IEEE Trans. Vis. Comput. Graph.* 23, 6 (2017), 1706–1724. https://doi.org/10.1109/TVCG.2016.2543720

[16] Germán Leiva, Cuong Nguyen, Rubaiat Habib Kazi, and Paul Asente. 2020. Pronto: Rapid Augmented Reality Video Prototyping Using Sketches and Enaction. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*. ACM, 1–13. https://doi.org/10.1145/3313831.3376160

[17] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. 2019. Context-Aware Online Adaptation of Mixed Reality Interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST 2019, New Orleans, LA, USA, October 20-23, 2019*, François Guimbretière, Michael S. Bernstein, and Katharina Reinecke (Eds.). ACM, 147–160. https://doi.org/10.1145/3332165.3347945

[18] Ewa Luger, Lachlan Urquhart, Tom Rodden, and Michael Golembewski. 2015. Playing the Legal Card: Using Ideation Cards to Raise Data Protection Issues within the Design Process. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI 2015, Seoul, Republic of Korea, April 18-23, 2015*. ACM, 457–466. https://doi.org/10.1145/2702123.2702142

[19] Mark Miller, Fernanda Herrera, Hanseul Jun, James Landay, and Jeremy Bailenson. 2020. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports* 10 (10 2020). https://doi.org/10.1038/s41598-020-74486-y

[20] Meredith Ringel Morris, Andreea Danielescu, Steven Mark Drucker, Danyel Fisher, Bongshin Lee, m. c. schraefel, and Jacob O. Wobbrock. 2014. Reducing legacy bias in gesture elicitation studies. *Interactions* 21, 3 (2014), 40–45. https://doi.org/10.1145/2591689

[21] Michael Nebeling and Katy Madier. 2019. 360proto: Making Interactive Virtual Reality & Augmented Reality Prototypes from Paper. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*. ACM, 596. https://doi.org/10.1145/3290605.3300826

[22] Michael Nebeling, Janet Nebeling, Ao Yu, and Rob Rumble. 2018. ProtoAR: Rapid Physical-Digital Prototyping of Mobile Augmented Reality Applications. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21-26, 2018*. ACM, 353. https://doi.org/10.1145/3173574.3173927

[23] Michael Nebeling, Shwetha Rajaram, Liwei Wu, Yifei Cheng, and Jaylin Herskovitz. 2021. XRStudio: A Virtual Production and Live Streaming System for Immersive Instructional Experiences. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*. ACM, 107:1–107:12. https://doi.org/10.1145/3411764.3445323

[24] Joon Sung Park, Joseph C. O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, UIST 2023, San Francisco, CA, USA, 29 October 2023- 1 November 2023*. ACM, 2:1–2:22. https://doi.org/10.1145/3586183.3606763

[25] Tomislav Pejsa, Julian Kantor, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2016. Room2Room: Enabling Life-Size Telepresence in a Projected Augmented Reality Environment. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, CSCW 2016, San Francisco, CA, USA, February 27 - March 2, 2016*. ACM, 1714–1723. https://doi.org/10.1145/2818048.2819965

[26] Shwetha Rajaram, Chen Chen, Franziska Roesner, and Michael Nebeling. 2023. Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI 2023, Hamburg, Germany, April 23-28, 2023*. ACM, 98:1–98:17. https://doi.org/10.1145/3544548.3581089

[27] Shwetha Rajaram and Michael Nebeling. 2022. Paper Trail: An Immersive Authoring System for Augmented Reality Instructional Experiences. In *CHI '22: CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April 2022 - 5 May 2022*. ACM, 382:1–382:16. https://doi.org/10.1145/3491102.3517486

[28] Shwetha Rajaram, Franziska Roesner, and Michael Nebeling. 2023. Reframe: An Augmented Reality Storyboarding Tool for Character-Driven Analysis of Security & Privacy Concerns. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, UIST 2023, San Francisco, CA, USA, 29 October 2023- 1 November 2023*. ACM, 117:1–117:15. https://doi.org/10.1145/3586183.3606750

[29] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J. Wang. 2014. World-Driven Access Control for Continuous Sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*. ACM, 1169–1181. https://doi.org/10.1145/2660267.2660319

[30] Adam Shostack. 2014. *Threat Modeling: Designing for Security* (1st ed.). Wiley Publishing.

[31] Sangho Suh, Sydney Lamorea, Edith Law, and Leah Zhang-Kennedy. 2022. PrivacyToon: Concept-driven Storytelling with Creativity Support for Privacy Concepts. In *DIS '22: Designing Interactive Systems Conference, Virtual Event, Australia, June 13 - 17, 2022*. ACM, 41–57. https://doi.org/10.1145/3532106.3533557

[32] Jacob O. Wobbrock, Meredith Ringel Morris, and Andrew D. Wilson. 2009. User-defined gestures for surface computing. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009, Boston, MA, USA, April 4-9, 2009*. ACM, 1083–1092. https://doi.org/10.1145/1518701.1518866

[33] Yixin Zou, Kevin A. Roundy, Acar Tamersoy, Saurabh Shintre, Johann Roturier, and Florian Schaub. 2020. Examining the Adoption and Abandonment of Security, Privacy, and Identity Theft Protection Practices. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*. ACM, 1–15. https://doi.org/10.1145/3313831.3376570

A APPENDIX

A.1 Overview of REFRAME System; Characters & Analysis Prompts

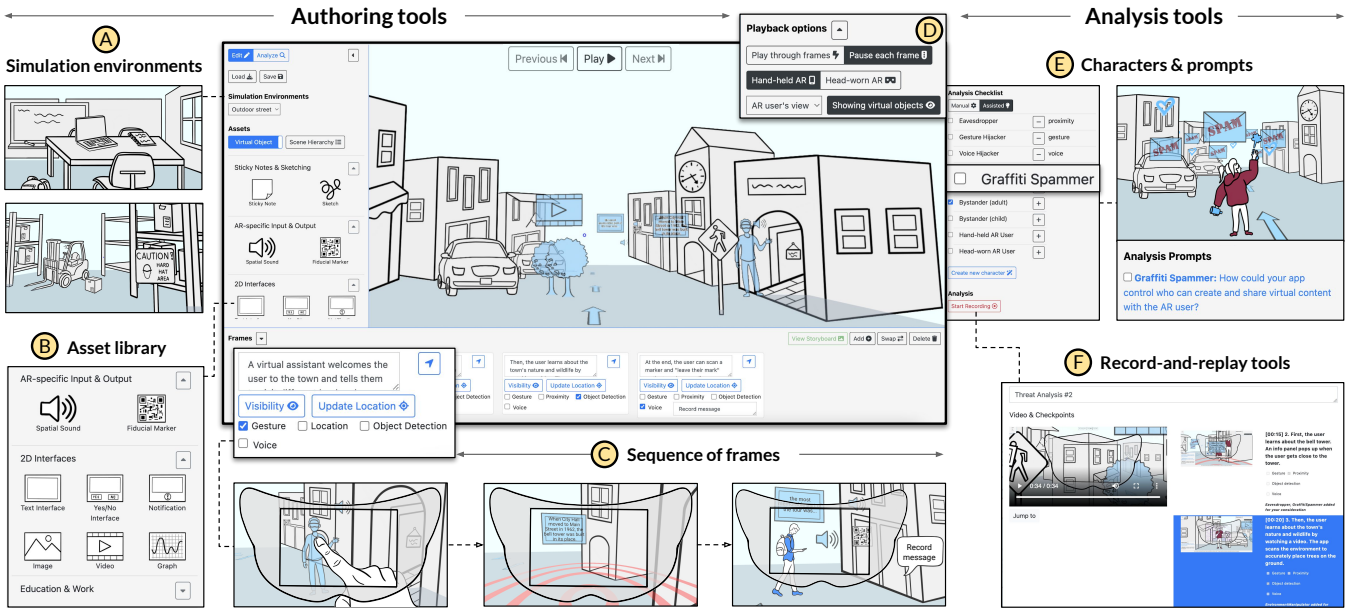

Figure 6: REFRAME User Interface. REFRAME enables designers to author storyboards as a sequence of of AR application states *(A, B, C)* and preview the sequence in the 3D scene, simulating either hand-held or head-worn AR *(D)*. REFRAME's character-driven analysis tools *(E)* insert personified representations of threats and analysis prompts to help designers brainstorm mitigation techniques. Designers can review their threat modeling via screen & audio recording tools *(F)*.

	Eavesdropper	Voice & Gesture Hijackers	Graffiti Spammer	Environment Manipulator	Bystanders (adult, child, hand-held & headworn AR users)
		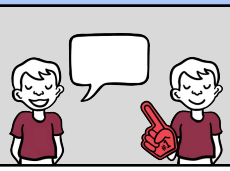		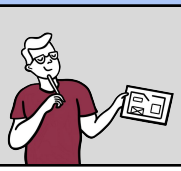	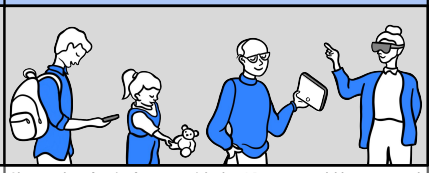
Description	Observes the AR user in the **physical space** to learn private info about the **virtual space**	Mimic the AR user's input in the **physical space** to make the AR app perform unwanted actions in the **virtual space**	Places unwanted **virtual content** in the **physical space**	Changes the **physical space** to manipulate app functionality in the **virtual space**	Shares the **physical space** with the AR user; could be captured by the AR device and represented in the **virtual space**
Analysis prompt	How could your app prevent others from learning sensitive info about the AR content or the AR user?	How could your app make sure that the AR content is only manipulated by authorized AR user(s)?	How could your app control who can create and share virtual content with the AR user?	How could your app still function with unforeseen changes to the physical environment?	How could your app minimize the info captured about bystanders and/or increase their awareness?

Figure 7: Characters and Analysis Prompts. REFRAME implements nine adversary and bystander characters depicting various privacy concerns based on a threat model for proxemic interactions in AR. We distinguish between adversaries (shown in red) and bystanders (shown in blue) to highlight threats posed *to* the AR user and threats posed *by* the AR user, respectively.

A.2 Privacy Adaptation Techniques: Access Control Techniques for Multi-User AR

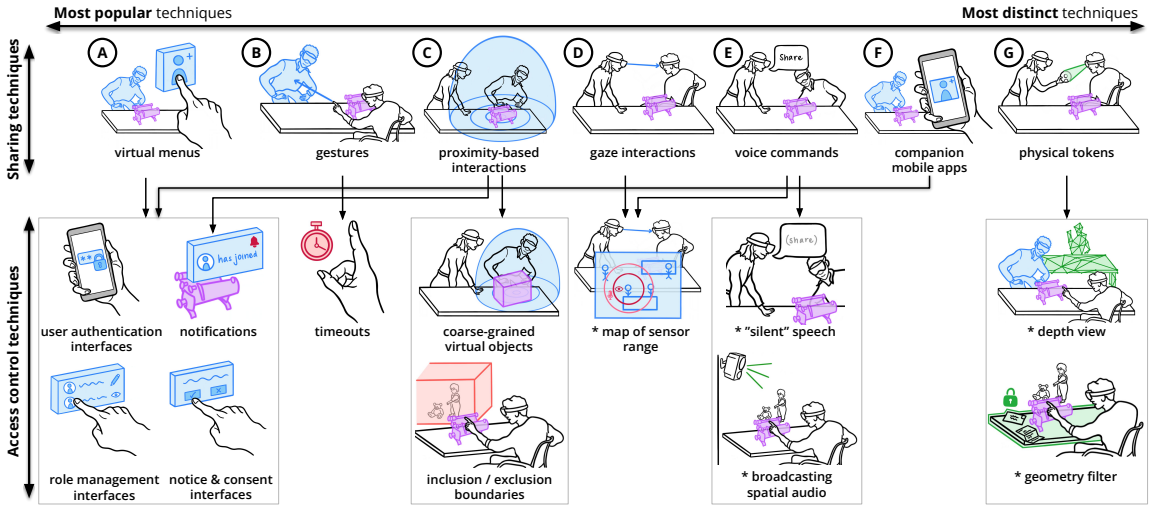

Figure 8: Design Catalog of Access Control Techniques for Multi-User AR: In our first elicitation study [26], we investigated how AR and privacy experts would design and adapt AR interaction techniques to mitigate threats to access control of physical and virtual content in a multi-user scenario. AR developers can use the resulting design space to augment base interaction techniques *(A-G)* with fine-grained access control mechanisms, e.g., allowing users to specify inclusion/exclusion boundaries to define areas of the physical environment that can be captured.

A.3 Privacy Assistant: Annotation Task Interface

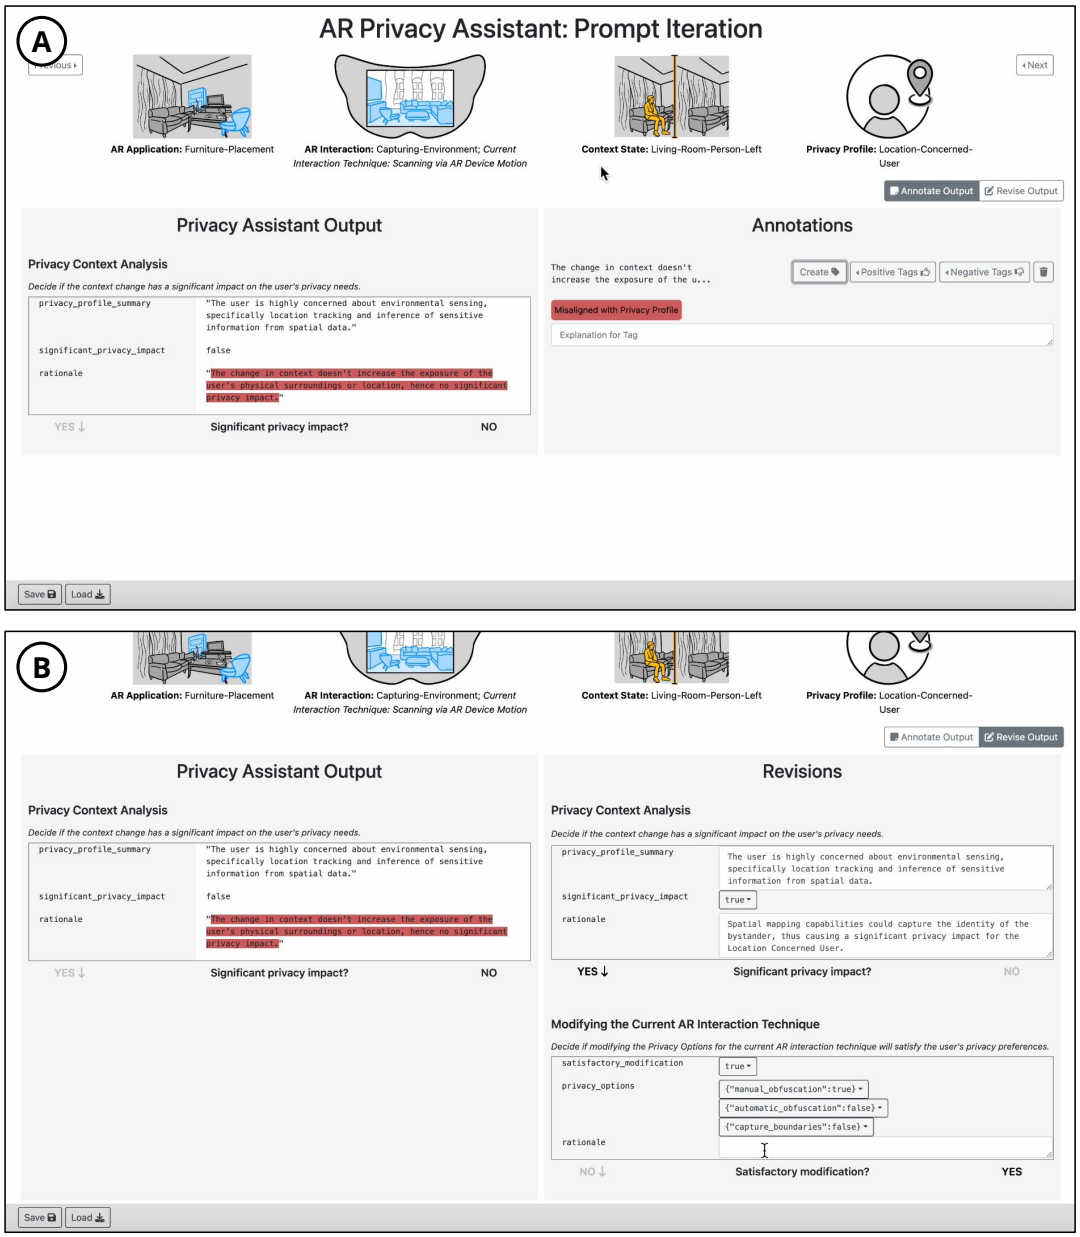

Figure 9: Privacy Assistant Annotation Interface. To assess the benefits and limitations of our initial LLM-enabled Privacy Assistant and work towards concrete design guidelines, we developed an annotation interface to facilitate a study with usable privacy experts. The interface supports reviewing the Privacy Assistant's output at each step of the recommendation workflow (Fig. 5) for different AR usage scenarios. To identify positive and negative aspects of the LLM's output, privacy experts can highlight text and assign custom tags (e.g., *"Statement is Misaligned with the User's Privacy Profile."*) (A). Experts can suggest revisions to the output by directly editing the text or variables (B). For example, the expert in (B) disagreed with the Privacy Assistant's assessment that the context event did not significantly impact the user's privacy needs; changing this "decision variable" allows the expert to view and edit the next phase of the recommendation workflow.