

Designing Privacy-Informed Sharing Techniques for Multi-User AR Experiences

Shwetha Rajaram
University of Michigan

Franziska Roesner
University of Washington

Michael Nebeling
University of Michigan

Abstract

With recent augmented reality (AR) systems becoming more prevalent and collaborative, there is an increased need for AR interaction techniques to mitigate the unique privacy concerns with multi-user, always-on AR. We present a study design based on two multi-user AR scenarios which expands on user-driven elicitation as a popular method in HCI by pairing AR and privacy experts together to iteratively design interactions and analyze corresponding privacy threats. A pilot leads us to believe this paired study design is promising for interaction proposals to incorporate privacy concerns and create overall safer designs while shaping a more holistic design approach.

1 Introduction

As augmented reality (AR) applications are becoming increasingly widespread and adding support for collaboration, particularly in professional and educational settings [8], AR designers need to consider safety concerns which may arise in always-on usage scenarios. These threats include the collection of sensitive data from users and the physical environment [1, 5, 17], as well as challenges unique to multi-user AR, like users' agency over content placed in private spaces [16]. To mitigate these threats, privacy researchers have developed frameworks for safely sharing virtual content in multi-user AR experiences which function at the operating system or application level [10, 18]. However, this technical focus does not address how end-users prefer to use these techniques.

In parallel, the HCI research community has focused on developing usable AR systems and interaction techniques,

also with a recent focus on collaboration [4, 13, 15, 23]. Recent work also explores the feasibility of implementing these interactions through technical advancements, such as mixed reality telepresence [9, 21], object tracking techniques [19], and shared displays [6]. However, our review of the last five years of CHI and UIST literature finds relatively little work which prioritizes privacy when designing AR interaction techniques. In our current work, we therefore aim to develop systematic methods for AR technologists to integrate existing privacy guidelines into their interaction design processes.

As a first step, we considered how user-driven elicitation studies as popularized by Wobbrock et al. [22] could be extended to incorporate privacy considerations when proposing interaction techniques. Elicitation has been established as a powerful method for designing intuitive interaction techniques, and recent studies have contributed gesture sets for Kinect-based [2, 11] and mixed reality interfaces [3, 14]. The method often utilizes pairs of end-users to propose gestures to accomplish given system functions, and has been shown to improve the identification and memorability of interaction techniques [2, 3]. However, elicitation usually does not explicitly consider privacy threats, which can raise challenges for implementing the interaction techniques (e.g., ambiguous gestures, needing to instrument users with additional sensors) [12, 20]. To consider technical constraints which current devices and gesture recognizers may impose, prior work has extended the traditional elicitation process to incorporate functional system prototypes in the design process [12, 20].

Our work explores how elicitation studies can be further evolved to explicitly address privacy in the design of interaction techniques. We propose pairing AR and privacy experts together to iteratively produce interaction proposals while analyzing potential privacy threats with respect to two multi-user AR scenarios. In this position paper, we describe our current study design and multi-user AR scenarios, as well as share preliminary findings from two pilot studies conducted with pairs of graduate students. We reflect on the benefits and limitations of our approach, and outline our study plans with experts from academia and industry.

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

VR4Sec: 1st International Workshop on Security for XR and XR for Security 2021.

August 6, 2021, Vancouver, B.C., Canada.




Figure 1: In the *Future of Education* scenario, **Alice & Bob** are two high-school students collaborating on an AR physics lab involving a physical roller coaster setup and virtual windows to collect and display simulation data. **Eve**, a student in a different lab group, should not have access to Alice & Bob’s lab content due to class collaboration policies. **Charlie**, the instructor, may need to access lab groups’ content to provide help or feedback. After the class ends, Alice can take her headset home, but will need a way to access the lab content asynchronously. Bob will return his headset, which may be used by students in other class periods.

2 Multi-User AR Scenarios

We created two multi-user AR usage scenarios around the *Future of Education* and the *Future of Work* to use as a basis for designing interaction techniques in our elicitation study. We opted for scenario-driven elicitation rather than a more open-ended approach, in order to provide the participants with specific details about the collaboration context and physical environment which they could draw on during the privacy analysis to infer the users’ privacy needs and expectations.

Our goal in developing the scenarios was to provide coverage with respect to three design dimensions related to collaboration and privacy considerations: (1) the **time / space matrix** [7] which accounts for co-located vs. remote users and synchronous vs. asynchronous collaboration, (2) **public vs. private spaces**, and (3) a **threat model** developed by Ruth et al. [18] which describes privacy harms which other AR users can pose, such as accessing private virtual content or performing unwanted content manipulation.

The major differences between the scenarios involve co-located vs. remote collaboration and private vs. public usage settings. The *Future of Education* scenario involves students collaborating on an AR physics simulation in a public classroom, while the *Future of Work* scenario involves two co-workers collaborating remotely to design a virtual car engine,

one working from their private home and the other from a public coffee shop. We provide the digital scenario sketch for the *Future of Education* scenario in Figure 1.

3 Study Design

We designed and piloted an elicitation study with two pairs of graduate students to gain preliminary insights into the effectiveness of pairing AR and privacy experts together to design safe interaction techniques. The study focused on the design of techniques to share virtual content in multi-user AR experiences, explicitly considering three design goals: 1) **usability**, 2) **technical feasibility**, and 3) **privacy**. The study was between-subjects with respect to the multi-user AR scenario: each pair only designed for either the *Future of Education* or *Future of Work* scenario due to time constraints. Our study design is outlined in Figure 2 and described in more detail in the rest of this section.

Participants: We recruited four Information Science graduate students through advertising in two courses offered at the University of Michigan specializing in AR/VR development and privacy. Most participants reported that these academic courses were their first formal experience studying either AR or privacy, so we asked each participant to watch two short videos from the Coursera Extended Reality for Everybody specialization¹ to ensure they had a baseline level of knowledge. The students acting in the AR expert role viewed videos on head-worn vs. hand-held AR, and those in the privacy expert role viewed videos on privacy and ethical considerations for AR devices.

Study introduction: First, we presented the participants with background information, including the study motivation and a digital handout describing the three design dimensions we used to create the scenarios (Sec. 2). To prompt participants to design novel interactions which go beyond existing techniques for sharing AR experiences (e.g. sending a URL or registering the same fiducial marker), we introduced a design concept where interactions with the physical environment can be utilized to share AR content.

Task 1 - Production and privacy analysis: To facilitate the iterative design of interaction techniques, we introduced the multi-user AR scenario using step-by-step prompts bringing in new user (e.g. for *Future of Education*, we first asked the AR expert to produce three interactions for Alice to give Bob access to the virtual lab content, then to revise the techniques considering Eve, who should not have access to Alice & Bob’s content). For each prompt, we facilitated turn-taking with the privacy expert, asking them to analyze threats which the sharing techniques could pose and suggest ways to prevent any privacy harms. To track the interaction proposals, we asked the experts to think aloud and collaboratively sketch

¹ Extended Reality for Everybody: <https://www.coursera.org/specializations/extended-reality-for-everybody>




Figure 2: We structured the pilot study in four phases: (1) an **introduction** to the 3 design dimensions, (2) a **back-and-forth design task** where the AR expert produces proposals for interaction techniques while the privacy expert performs threat modeling, (3) a **revision task** considering the design goals of usability and feasibility in addition to privacy, and (4) a **discussion** around the pros & cons of pairing AR & privacy experts and future design guidelines.

their ideas on top of the digital scenario sketch using Google Jamboard².

Task 2 - Holistic review: The primary goal of this task was to observe the interplay between the AR and privacy experts and infer which of the three design goals they were prioritizing when refining their sharing techniques. As the previous task was mainly focused on the design goals of usability and privacy, we first prompted the AR expert to explicitly consider technical feasibility by explaining how they would implement the sharing techniques they designed. Then, we facilitated another back-and-forth dialogue between the AR and privacy experts, this time asking either one to propose a revision to improve the sharing techniques and debate whether the revision should be accepted. We encouraged the experts to suggest holistic revisions addressing any of the three design goals and analyze the impact of the revision on all of the goals.

Discussion: To elicit feedback on our elicitation method, we conducted a discussion and debrief session with the participants. Based on their experience during the iterative design process in Tasks 1 & 2, we asked both experts to suggest guidelines for designing AR sharing techniques to offer to future designers and developers. Then, we reflected on the benefits and disadvantages of pairing two experts from different backgrounds together for elicitation sessions.

4 Preliminary Findings

In this section, we describe some of the initial findings from our pilot study with regards to similarities we observed in the AR sharing techniques designed by the experts and the effectiveness of the paired study design.

4.1 Commonalities Across the Sharing Techniques

We present two commonalities in the elicited AR sharing techniques from our analysis of the participants’ dialogue and Jamboard annotations. We were particularly interested in seeing whether there would be similarities involving the technical implementation and how much the techniques made use of the physical environment, given the differences in the scenarios with respect to our design dimensions (Sec 2).

Sharing virtual content through interacting with physical surfaces. Both pairs of participants explored similar sharing techniques where users can access virtual content through interacting with a designated physical surface, but envisioned different implementations. The *Future of Education* pair used a marker-based approach based on Alice & Bob’s physical proximity to a QR code fixed to the table. The *Future of Work* pair opted for a marker-less approach involving depth scanning to select physical planar surfaces, such as tables and walls; to share virtual content with their remote collaborators, the co-workers could anchor the content to the shared surface they designated in their physical environment.

Using proximity-based techniques to mitigate confidentiality threats. Across both scenarios, the participants also adopted similar design concepts to specifically address threats of other users gaining access to private virtual content when working in public spaces. To prevent shoulder surfing attacks from Eve, the *Future of Education* participants suggested that the collaborators’ headsets should be paired with their assigned lab station by registering the QR code at a very close distance. The *Future of Work* pair also adopted a proximity-based technique which they described as an “AR force field” to prevent other people in the coffee shop from shoulder surfing or sharing unwanted content.

²Google Jamboard: <https://edu.google.com/products/jamboard/>

4.2 Pros & Cons of Paired Study Design

We reflect on the effectiveness of the study design, based on the participants' feedback during the discussion portion of our study as well as our observations on the interplay between experts. Overall, the participants responded positively to the pairing of AR and privacy roles, but noted the potential for the privacy analysis to impede the AR experts' design process.

Experts' diverse perspectives are useful for considering a variety of design goals. All participants expressed that a main benefit of pairing AR and privacy experts together was motivating each other to more effectively address all three design goals of usability, feasibility, and privacy. One AR expert explained that he learned a lot about novel privacy threats that AR technologies may pose through "understanding the privacy expert's perspective," and throughout the study, he increasingly felt the need to "bounce ideas" off of each other. However, both students serving as privacy experts noted the importance of establishing a "shared vocabulary" in order to better understand the AR experts' interaction technique designs, particularly with regards to technical concepts like marker-based vs. marker-less AR.

Privacy experts' analysis may limit the AR experts' creativity. One disadvantage raised by the privacy experts was their tendency to "shut down the creative process" of the AR experts through pointing out flaws in the sharing techniques. We observed this finding to some degree with one of the AR experts, who initially designed a relatively novel proximity-based technique, but quickly abandoned this design when the privacy expert pointed out the potential for other users to shoulder surf, instead opting for a legacy marker-scanning technique.

5 Reflection

We believe that our pilot studies showed promise for extending the elicitation method, as the privacy experts' analysis had the intended effect of encouraging the AR experts to design increasingly defensive interaction techniques as the study progressed. It was interesting to see that the privacy analysis seemingly made it more challenging for the AR expert to be creative in their designs, as they tended to opt for techniques which made use of fewer unique affordances of AR devices, but were more robust against possible privacy threats.

Limitations: We are aware of a few limitations of our study. Our two multi-user scenarios do not guarantee coverage of all potential privacy threats which can arise with always-on AR. As such, the sharing techniques elicited through our study may be limited in their generalizability to other use cases.

Additionally, there is still an open question of which metrics are important for assessing the quality of interaction techniques with respect to our three design goals. We are currently exploring metrics including how well the elicited interaction proposals align with existing AR usability guidelines, to what

extent the implementation could make use of existing devices and development toolkits, and coverage of potential privacy threats. We anticipate challenges in evaluating additional quality metrics without implementing a functional prototype of the interaction techniques, particularly usability and technical aspects, such as user performance and gesture ambiguity.

Future Work: To better establish these quality metrics, we plan to run a study with 12 professional experts in AR and privacy who have at least 3 years of experience working in related fields in industry or academia. Through analyzing the interplay and dialogue between the experts, we hope to identify their underlying mental models for what constitutes an effective interaction technique with respect to usability, technical feasibility, and privacy, and how they are implicitly prioritizing these design goals over one another. Through our analysis, we will also extract design guidelines to aid AR designers and developers in the creation of safe, multi-user interaction techniques in the future.

6 Contribution to the Workshop

We hope to share our insights from this project with other workshop participants, particularly around the challenges we experienced when designing the elicitation method and alternate designs we investigated. We will also contribute insights from piloting our approach, including the benefits and trade-offs we observed when pairing AR and privacy experts – two stakeholders with potentially conflicting goals and priorities – to design AR interaction techniques. Our study serves as one example of a systematic approach to consider privacy threats when designing AR systems. Through the workshop, we hope to engage in broader discussions on how to more effectively integrate existing privacy guidelines into XR designers' and developers' workflows.

7 Authors' Backgrounds

Shwetha Rajaram is a PhD student at the University of Michigan School of Information. She is interested in XR design & development and how to make XR systems more safe and privacy-friendly for users. **Franzi Roesner** is an Associate Professor in the Paul G. Allen School of Computer Science and Engineering at the University of Washington, where she co-directs the Security and Privacy Research Lab. She has studied security and privacy issues in XR systems since 2011. **Michael Nebeling** is an Assistant Professor at the University of Michigan School of Information, where his HCI research lab focuses on XR systems design. Since the Social HMD CHI 2019 workshop, he developed an interest in critical design and systematic approaches to safe XR systems.

References

- [1] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 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*, pages 427–442. USENIX Association, 2018.
- [2] Abdullah X. Ali, Meredith Ringel Morris, and Jacob O. Wobbrock. Crowdlit: A system for conducting distributed end-user elicitation and identification studies. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*, page 255. ACM, 2019.
- [3] Abdullah X. Ali, Meredith Ringel Morris, and Jacob O. Wobbrock. "i am iron man": Priming improves the learnability and memorability of user-elicited gestures. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*, pages 359:1–359:14. ACM, 2021.
- [4] Huidong Bai, Prasanth Sasikumar, Jing Yang, and Mark Billingham. A user study on mixed reality remote collaboration with eye gaze and hand gesture sharing. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*, pages 1–13. ACM, 2020.
- [5] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 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*, pages 2377–2386. ACM, 2014.
- [6] Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. Sharevr: Enabling co-located experiences for virtual reality between HMD and non-hmd users. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, May 06-11, 2017*, pages 4021–4033. ACM, 2017.
- [7] R. Johansen. Groupware: Computer support for business teams. 1988.
- [8] Kangsoo Kim, Mark Billingham, Gerd Bruder, Henry Been-Lirn Duh, and Gregory F. Welch. Revisiting trends in augmented reality research: A review of the 2nd decade of ISMAR (2008-2017). *IEEE Trans. Vis. Comput. Graph.*, 24(11):2947–2962, 2018.
- [9] Balasaravanan Thoravi Kumaravel, Fraser Anderson, George W. Fitzmaurice, Bjoern Hartmann, and Tovi Grossman. Loki: Facilitating remote instruction of physical tasks using bi-directional mixed-reality telepresence. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST 2019, New Orleans, LA, USA, October 20-23, 2019*, pages 161–174. ACM, 2019.
- [10] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. Securing augmented reality output. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 320–337. IEEE Computer Society, 2017.
- [11] Meredith Ringel Morris. Web on the wall: insights from a multimodal interaction elicitation study. In *Interactive Tabletops and Surfaces, ITS'12, Cambridge/Boston, MA, USA, November 11-14, 2012*, pages 95–104. ACM, 2012.
- [12] Michael Nebeling, Alexander Huber, David Ott, and Moira C. Norrie. Web on the wall reloaded: Implementation, replication and refinement of user-defined interaction sets. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces, ITS 2014, Dresden, Germany, November 16 - 19, 2014*, pages 15–24. ACM, 2014.
- [13] Michael Nebeling, Katy Lewis, Yu-Cheng Chang, Lihan Zhu, Michelle Chung, Piaoyang Wang, and Janet Nebeling. Xrdirector: A role-based collaborative immersive authoring system. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*, pages 1–12. ACM, 2020.
- [14] Thammathip Piumsomboon, Adrian J. Clark, Mark Billingham, and Andy Cockburn. User-defined gestures for augmented reality. In *Human-Computer Interaction - INTERACT 2013 - 14th IFIP TC 13 International Conference, Cape Town, South Africa, September 2-6, 2013, Proceedings, Part II*, volume 8118 of *Lecture Notes in Computer Science*, pages 282–299. Springer, 2013.
- [15] Thammathip Piumsomboon, Gun A. Lee, Andrew Irlitti, Barrett Ens, Bruce H. Thomas, and Mark Billingham. On the shoulder of the giant: A multi-scale mixed reality collaboration with 360 video sharing and tangible interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*, page 228. ACM, 2019.
- [16] Franziska Roesner, Tadayoshi Kohno, and David Molnar. Security and privacy for augmented reality systems. *Commun. ACM*, 57(4):88–96, 2014.
- [17] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J. Wang. 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*, pages 1169–1181. ACM, 2014.
- [18] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. Secure multi-user content sharing for augmented reality applications. In *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 141–158. USENIX Association, 2019.
 - [19] Kihoon Son, Hwiwon Chun, Sojin Park, and Kyung Hoon Hyun. C-space: An interactive prototyping platform for collaborative spatial design exploration. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*, pages 1–13. ACM, 2020.
 - [20] Maximilian Speicher and Michael Nebeling. Gesturewiz: A human-powered gesture design environment for user interface prototypes. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21-26, 2018*, page 107. ACM, 2018.
 - [21] Theophilus Teo, Louise M. Lawrence, Gun A. Lee, Mark Billingham, and Matt Adcock. Mixed reality remote collaboration combining 360 video and 3d reconstruction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*, page 201. ACM, 2019.
 - [22] Jacob O. Wobbrock, Meredith Ringel Morris, and Andrew D. Wilson. 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*, pages 1083–1092. ACM, 2009.
 - [23] Haijun Xia, Sebastian Herscher, Ken Perlin, and Daniel Wigdor. Spacetime: Enabling fluid individual and collaborative editing in virtual reality. In *The 31st Annual ACM Symposium on User Interface Software and Technology, UIST 2018, Berlin, Germany, October 14-17, 2018*, pages 853–866. ACM, 2018.