



Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality

Shwetha Rajaram
University of Michigan
Ann Arbor, MI, USA
shwetha@umich.edu

Franziska Roesner
University of Washington
Seattle, WA, USA
franzi@cs.washington.edu

Chen Chen
University of Michigan
Ann Arbor, MI, USA
noracc@umich.edu

Michael Nebeling
University of Michigan
Ann Arbor, MI, USA
nebeling@umich.edu

ABSTRACT

The HCI community has explored new interaction designs for collaborative AR interfaces in terms of usability and feasibility; however, security & privacy (S&P) are often not considered in the design process and left to S&P professionals. To produce interaction proposals with S&P in mind, we extend the user-driven elicitation method with a scenario-based approach that incorporates a threat model involving access control in multi-user AR. We conducted an elicitation study in two conditions, pairing AR/AR experts in one condition and AR/S&P experts in the other, to investigate the impact of each pairing. We contribute a set of expert-elicited interactions for sharing AR content enhanced with access control provisions, analyze the benefits and tradeoffs of pairing AR and S&P experts, and present recommendations for designing future multi-user AR interactions that better balance competing design goals of usability, feasibility, and S&P in collaborative AR.

CCS CONCEPTS

• **Human-centered computing** → **Scenario-based design**; **Mixed / augmented reality**; • **Security and privacy** → *Usability in security and privacy*.

KEYWORDS

elicitation studies, threat modeling

ACM Reference Format:

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 '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3544548.3581089>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9421-5/23/04...\$15.00
<https://doi.org/10.1145/3544548.3581089>

1 INTRODUCTION

The proliferation of new extended reality tools that enable collaboration between users in co-located and remote settings, such as Spatial, Microsoft Mesh, and Meta Horizon Workrooms¹, raises many new challenges. The HCI community is exploring multi-user augmented reality (AR) on the system infrastructure side; recent work addresses technical challenges such as bridging remote environments through spatial capture [4, 8, 23, 53], as well as improving usability through novel techniques to encourage awareness and communication between collaborators [17, 18, 41, 55].

A related stream of research in the security & privacy (S&P) domain explores novel threats involving access control of virtual content and physical spaces in multi-user AR. Prior work has studied threats related to unwanted access or manipulation of virtual content [49], norms for placing virtual content in personal or private physical spaces [27, 42, 49], and unwanted capture and sharing of environmental information involving other users or bystanders [1, 13, 48]. However, in reviewing the body of HCI research on AR systems and interaction techniques, we notice that security & privacy is often not a major consideration in design. This is particularly concerning given the accelerated adoption of AR, e.g., in educational contexts [43], and as AR form factors become more suitable for everyday, always-on usage. Our work seeks to bridge the gap between these two separate threads of research through exploring methods for integrating S&P considerations into the design process for multi-user AR interaction techniques.

As a step in this direction, we explore how *user-driven elicitation* [59] can be extended to incorporate S&P considerations in the design of interaction techniques. Elicitation studies have been widely-established in HCI as a method for working with end-users to propose intuitive interactions that accomplish a given system function or effect [33, 56, 59]. Benefits of this approach have been demonstrated with respect to *usability* design goals (e.g., increased memorability and identification of interactions [2, 3]). To consider constraints impacting *technical feasibility* (e.g., whether interaction techniques can be achieved with existing gesture recognizers or via on-device sensors [51] and the implementation effort required), prior work incorporated functional system prototypes in elicitation studies [36, 37, 51].

¹Spatial: <https://spatial.io>; Microsoft Mesh: <https://www.microsoft.com/en-us/mesh>; Meta Horizon Workrooms: <https://www.meta.com/work/workrooms/>

In addition to usability and technical feasibility, our work integrates consideration of *security & privacy* design goals, which are not explicitly incorporated in prior elicitation studies. Inspired by research in the usable privacy domain [35, 61], we adopted a scenario-based elicitation approach to provide a concrete basis for analyzing S&P considerations corresponding to specific interaction proposals. Drawing on established use cases for AR in education [43, 55], we created a scenario where students collaborate on an engineering lab using head-worn AR and designed a digital sketch to depict the students' different collaboration contexts (AR users vs. non-users, co-located vs. remote users, and public vs. private spaces). We focus on threats involving access control of virtual content and physical spaces in multi-user AR, adopting Ruth et al.'s threat model for multi-user AR interactions [49] as a framework for navigating threats from different people's perspectives. Our protocol adapts Morris' production, priming, and partners (PPP) method [33], pairing two designers together to produce interaction proposals for sharing AR content. In addition to production, we incorporate a revision phase, where partners critique each others' proposals with respect to the threat model and suggest ways to mitigate potential threats.

A key question for our research was whether our scenario-based elicitation approach provides sufficient guidance to mitigate threats in interaction proposals, or whether additional expertise in S&P is still required to elicit high quality proposals. To assess the impact that varying degrees of S&P expertise on the design team has on the set of elicited interactions, we conducted a between-subjects study with two conditions: pairing two AR experts together (*AR/AR* condition) and pairing an AR expert with a S&P expert (*AR/SP* condition). Overall, our elicitation study yielded sharing techniques enriched with access control provisions in *both* conditions. While the *AR/SP* pairs produced access control techniques earlier on in the elicitation session, pairs in both conditions produced similar types of these techniques, including both interactions adapted from legacy systems and more creative interactions tailored to the specific usage scenario. These results suggest that when it is not feasible for AR interaction designers to directly include S&P professionals in the design process, a scenario-based elicitation process as demonstrated by our approach – leveraging prompts that increasingly introduce threats along a threat model and facilitating turn-taking between designer and critic roles – can result in the design of S&P-aware yet creative design proposals.

This work contributes (1) an empirical study exploring the effects of pairing two AR experts versus AR and S&P experts for our scenario-based elicitation process, (2) the resulting set of expert-elicited interaction techniques for sharing AR content and providing access control, and (3) design recommendations for multi-user AR interactions which mitigate tradeoffs between design goals for collaboration and access control.

2 RELATED WORK

Our work builds upon prior research in elicitation studies, multi-user interaction techniques & systems, and security & privacy considerations for AR experiences related to access control of virtual and physical spaces.

2.1 Elicitation Studies

Our design approach is inspired by prior work in *user-driven elicitation*, which was popularized by Wobbrock et al. [59] and is widely used in HCI research as a method for designing interaction techniques with non-expert users [56]. These studies are often facilitated by using a Wizard-of-Oz prototype to present participants with a *referent* (the effect of a system function, e.g., advancing a slide or zooming in) and asking them to propose actions that could achieve that effect. Proponents of this approach note the advantages in designing interactions with non-technical users rather than software developers, who may prioritize the implementation constraints of the system over the mental models and capabilities of end-users [3, 59]. There are many demonstrated benefits of user-defined gesture sets, including that new end-users can more easily remember the symbols and identify their intended effects without having seen the gesture set before [2]. Related to our focus on interaction techniques for multi-user AR, prior elicitation studies contributed mid-air gestures for mixed reality systems [3, 40], Kinect-based interfaces [32], and virtual mirror displays [28].

However, user-defined gestures can be difficult to implement in interactive systems, e.g., requiring additional effort to train custom recognizers or instrumenting users with additional sensors to track areas of the body that are difficult to capture via AR HMDs' built-in cameras. Prior work addressed this limitation through utilizing functional systems prototypes that respond to user-defined actions [36, 37, 51] instead of Wizard-of-Oz, in order to understand end-users mental models and preferences towards interaction techniques that current gesture recognizers are capable of supporting. In addition to the design goals of usability and technical feasibility that prior literature explored, our work adds a third goal of security & privacy, which is not explicitly considered in previous elicitation studies but is increasingly important for multi-user AR systems.

We also draw on recent work studying users' privacy preferences with IoT devices, through jointly eliciting design proposals and conducting privacy analyses with various stakeholders. Yao et al. conducted a co-design workshop where non-expert users analyzed privacy concerns from the perspective of different users in a role play activity, then prototyped designs for privacy-friendly smart home devices [61]. Working with experts from a variety of privacy-related disciplines, Emani-Naeini et al. iteratively elicited factors to compose privacy and security "nutrition labels" for IoT devices, having the experts anonymously review an aggregated list of factors at each stage and provide their rationale for accepting or rejecting them [34]. We took inspiration from how these works considered different stakeholder perspectives to evaluate the privacy implications of the proposed designs. To more explicitly incorporate the dimension of implementation feasibility, we chose work with expert participants who could discuss and address technical tradeoffs that arise with the interaction proposals.

2.2 Multi-User AR Systems

There is a long trajectory of HCI research dedicated to developing system infrastructure for multi-user AR experiences. In this section, we review recent examples of collaborative AR systems from the CHI and UIST conferences and identify trends with respect to our three design goals of technical feasibility, usability, and S&P.

We observed an increased interest in enabling more flexible and ad-hoc distributed collaboration systems in AR, which requires overcoming technical challenges in environmental sensing to bridge the gap between remote environments. For remote meetings and training, techniques have been developed for capturing and conveying collaborators' environments to other AR users through first and third person video [8, 23], 360° video [41, 53], and 3D reconstructions of users' surroundings [4, 23, 24, 53] (as depth sensing technologies become increasingly available). Prior work has also leveraged fine-grained sensing mechanisms and algorithms for object-tracking [50] and people-tracking [8, 17, 21, 23], in order to enable more contextually-aware AR experiences which can operate in dynamic multi-user environments. A majority of these recent multi-user systems utilize hand-held and head-worn AR, although tabletop [6, 30, 50] and room-scale projective AR experiences [17, 21] continue to be explored.

In terms of usability design goals, we identified a trend towards interaction techniques for remote collaboration that provide a variety of visual and audio cues to aid remote users in completing tasks and increase their sense of presence. One example is Piumsomboon et al.'s technique utilizing a "Giant" worker who manipulates a 360° camera tracked with 6 DOF to help a "Miniature" collaborator navigate the giant's environment [41]. Common metrics used to validate the usability of system prototypes include subjective ease-of-use ratings via established scales such as SUS [4, 29, 52, 53], cognitive workload [4, 29, 41, 58], user preferences towards system features [8, 41], and task performance [8, 41, 53, 55, 58], which is consistent with Dey et al.'s findings from a review of 10 years of AR usability studies [14]. In more recent literature, we also identify an increased interest in measuring social presence [4, 8, 29, 41, 53].

However, a majority of the systems research we reviewed did not explicitly consider S&P in their design process. A few exceptions introduced interaction techniques to designate fully public vs. fully private AR content [18] and stop sharing camera data in remote collaboration scenarios [41], but do not support more fine-grained sharing controls or mitigate other threats with using AR in public environments, such as privacy harms involving bystanders. Our work seeks to address this gap in literature around multi-user AR interaction techniques which are designed with S&P in mind, in addition to the design goals of usability and feasibility. We opted for a more open-ended design approach based on multi-user AR usage scenarios, rather than eliciting interaction techniques through a functional prototype which could impose technical constraints (e.g., small field-of-view and prescribed input modalities).

2.3 Threats Involving Access Control in Multi-User AR

The widespread use of AR through always-on, personal computing devices could bring about novel S&P risks which have not yet been experienced with screen-based technologies, due to AR devices' unique sensing and immersive output capabilities [46, 47]. Prior work in AR/VR and lifelogging devices has explored a variety of social concerns which AR could enable, including surfacing sensitive information through the collection of biometric and environmental data [1, 13, 19, 22, 48], inserting undesirable or harmful

content [25, 26, 31, 42], and causing physical harm through manipulating users' perception of the real world [10, 26, 54].

Beyond these challenges, multi-user AR can pose new threats related to access control of virtual content and physical spaces, which is the focus of our work. Ruth et al. contribute a threat model involving untrusted other users in multi-user AR [49], including adversaries accessing confidential virtual content, performing unwanted content placement or manipulation, or learning private information about the user or their physical environment. Prior work has also studied users' perception of ownership within physical and virtual spaces and agency over when and where to see content [27, 42, 45]. To mitigate these threats, recent work from the computer security community proposes technical implementations of sharing frameworks and policies to mitigate S&P concerns. Ruth et al. developed a set of sharing principles to protect AR users from other untrustworthy users [49], including outbound controls to specify user permissions for accessing and editing virtual content and inbound controls to prevent receiving unwanted content. Roesner et al. proposed the world-driven access control framework [48], which utilizes the user's real-world context to control AR applications' access to sensor data (e.g., apps can be denied access to always-on sensors such as cameras and microphones when the user enters a bathroom). Lebeck et al. implemented a policy specification framework to mitigate malicious or accidental behavior from AR applications by automatically changing the appearance of virtual content that may pose safety concerns to users [26].

In our work, we study novel interaction techniques for sharing AR content, rather than platform-level policies that govern sharing. We adapt Ruth et al.'s [49] threat model in our elicitation process to analyze threats related to access control which could arise with specific interaction proposals.

3 MULTI-USER AR SCENARIO

To provide a concrete basis for designing sharing techniques and analyzing corresponding threats involving access control, we based our elicitation study on a specific multi-user scenario utilizing head-worn AR. Our *Future of Education* scenario was inspired from two established AR use cases from prior work: (1) enabling new educational experiences using AR [38, 43, 44, 55] and (2) supporting remote collaboration through immersive workspaces [4, 23, 41, 53]. We aimed to increase experimental control in our elicitation study by focusing on an established AR use case and familiar physical contexts (i.e., educational and personal home environments where people have clearly-defined expectations of security & privacy [39]), so that participants could concentrate on designing interaction techniques rather than debating the S&P needs for the scenario. We also covered a variety of design dimensions (e.g., co-located vs. remote collaboration across both public and private spaces) to increase the generalizability of the elicited sharing techniques to other scenarios that share these characteristics.

This section first presents the three design dimensions which informed the creation of the scenario and were later used in the elicitation study to support the participants' design and critique processes. Then, we walk through a digital sketch of the scenario, presenting our rationale for including specific visual elements in the sketch to convey different users and design dimensions.




Figure 1: Future of Education Scenario: Violet, Russell and Doug collaborate on an AR engineering lab involving a physical car engine and a room-scale AR experience. Violet and Russell are working from a co-located classroom along with Carly, a student in another lab group, and Jack, the instructor. Doug is working remotely from his private home while taking care of his daughter, Ellie.

3.1 Design Dimensions

To guide the development of the scenario, we defined three dimensions to consider when designing techniques for sharing AR content in multi-user experiences: (1) **the time/space matrix** [20] which maps collaboration dimensions in terms of co-located vs. remote users and synchronous vs. asynchronous sharing; (2) **considerations for public vs. private spaces**, e.g., people who are present in addition to the AR users and their ownership of objects within the space [42]; (3) **a threat model involving access control of virtual content and physical spaces in multi-user AR** developed by Ruth et al. [49], with four classes of threats: untrustworthy individuals accessing private virtual content, performing unwanted content placement or manipulation, or learning private information about the user or their physical space. In our scenario, we consider different granularities of access control, e.g., different permissions for viewing and editing AR content, allowing the AR device to capture varying degrees of detail about the physical environment. Following the scope of Ruth et al.'s threat model [49], we focus on ways that the AR users represented in our scenario could act as adversaries; we do not explicitly consider a higher level of adversarial entities, e.g., service providers, network hackers, or other applications. We further discuss considerations for access control in our scenario in Sec. 3.2.2.

3.2 Scenario: Future of Education

In Figure 1, **Violet, Russell** and **Doug** are three students collaborating on an AR engineering lab. They manipulate the configuration of a physical engine and use a room-scale, head-worn AR experience of a car to simulate its functionality. Violet and Russell are in a co-located classroom environment, and Doug collaborates from his personal home while taking care of his daughter, **Ellie**, who

frequently walks around the room to pick up various toys. **Carly** is a student working in a different lab group, but since the lab is a graded assignment and collaboration across lab groups is prohibited, she should not have access to Violet, Russell, and Doug's virtual simulation content. **Jack**, the instructor, oversees the class and provides feedback to students when necessary. The headsets may be used by students in different class periods, who are also prohibited from accessing Violet, Russell and Doug's lab content.

We also used the design dimensions to create a set of four prompts for eliciting interaction techniques (Table 1). We formulated these prompts to span the time/space matrix, first considering the co-located collaborators (Violet & Russell), then bringing in Doug who is collaborating remotely. The prompts also increasingly raise threats to access control by bringing in Doug's private home (Prompt 2) and considering the bystanders and non-users in the scenario (Prompt 3).

3.2.1 Scenario Design Rationale. We systematically designed the *Future of Education* scenario to provide coverage of the design dimensions (Sec 3.1), selecting story and visual elements to depict different considerations for multi-user AR in the digital scenario sketch. The public, co-located classroom is an open environment with few physical barriers between users; the equipment and tables are shared between many students in different class periods, so there is little notion of private spaces or objects within the environment. In contrast, Doug's personal home could contain private information (e.g., the mail on the table and the security system console on the left wall). We also incorporated bystanders and non-users to introduce threats to access control based on Ruth et al.'s threat model [49]. For example, Carly could attempt to gain access to Russell, Violet and Doug's virtual content. Doug's collaborators

Task 1 Elicitation prompts
1. What are possible interactions for Russell and Violet to create a shared view on the AR experiment?
2. What are possible interactions for Russell and Violet to create a shared view with Doug?
3. Would you change anything about these interactions when considering Jack, Ellie and Carly and if so, how?
4. What are possible interactions for Russell, Violet, and Doug to end the shared session?

Table 1: Elicitation design prompts: We provided the participants with step-by-step elicitation prompts to scaffold the design process in Task 1. Prompt 1 involves only the co-located collaborators (Russell & Violet), then Prompt 2 brings in the remote collaborator (Doug). Prompt 3 serves as a cue to consider privacy harms related to the interaction techniques by bringing in the Jack, Ellie, and Carly, who could be considered passive bystanders or potential adversaries. Prompt 4 involves ending the collaborative session, which may motivate discussions around accessing the AR content at a later time and preventing access for other students, who might use the same headsets in another class period.

could gain access to private information about Doug’s physical space including Ellie, whose identity Doug may want to protect.

While our scenario depicts the use of head-worn AR to enable a marker-less, room-scale experience, we intentionally left other system requirements open for the expert participants to define, based on the implementation needs of their interaction proposals (e.g., instrumenting the physical classroom with additional sensors to enable more robust body tracking). This enabled us to study a wide range of interaction techniques together with the experts’ assumptions and considerations for usability, feasibility and S&P.

3.2.2 Access Control Needs. Our scenario suggests the following access control needs according to Ruth et al.’s threat model [49], considering untrustworthy individuals who: (1) **Access private virtual content.** The AR application should preserve confidentiality of the AR lab content and prevent unauthorized individuals from gaining access, since only students assigned to a particular lab group should be able to access their own simulation. To enable their instructor to assist them when needed, students may need to grant viewing access to the AR lab on a per-user or per-role basis. (2, 3) **Perform unwanted content placement or content manipulation.** The AR app should also preserve the integrity of the AR lab content (e.g., prevent unauthorized parties from editing, prevent teammates from accidentally or purposely deleting others’ contributions) and give users agency over whether and where to place content in the physical environment (e.g., allow Doug to reposition content to avoid obscuring his daughter, prevent students from sharing spam content). (4) **Learn private information about the user or their physical space.** Students may want fine-grained control over their AR headset’s data collection and storage procedures, to prevent instructors or other students who use their headset at a later time from learning or inferring private information based on the device data (e.g., age, gender, grades). Additionally, students may feel uncomfortable with other students’ headsets passively capturing data about them, which raises a need for mechanisms to preserve bystanders’ privacy. This is likely a concern for Doug, whose application could capture and convey sensitive areas of his home to his collaborators, e.g., microphone data involving his daughter Ellie.

We encouraged our study participants to define additional access control needs within the scope of how AR users in the scenario could act as adversaries. We did not explicitly consider other adversarial entities, such as malicious service or application providers.

4 STUDY DESIGN

We designed an elicitation study adapting Morris et al.’s approach emphasizing production, priming, and partners (PPP) [33] to elicit interaction techniques for sharing AR content in multi-user AR experiences. We structured the elicitation process around our *Future of Education* scenario (Sec. 3) and incorporated an implicit threat modeling activity through a series of prompts related to access control, based on Ruth et al.’s threat model for multi-user AR. [49].

To investigate the effects of pairing AR experts with security & privacy experts when following our elicitation approach, we conducted a between-subjects study with a total of 16 participants. The study consisted of two conditions: **AR/AR**, where two AR experts were paired together and took turns proposing interactions and critiquing each others’ proposals, and **AR/SP**, where the AR expert proposed interactions and the S&P expert critiqued the proposals. We gave pairs in both conditions access to the design dimensions that we used to create the scenario and prompts (Sec.3.1), but they were not constrained to these and encouraged to consider additional design goals and S&P issues, given their expertise in AR interaction design and S&P. The study was IRB-approved and conducted remotely via Zoom to accommodate participants across different time zones and countries. Each study session lasted 1.5 hrs and participants were compensated with \$50 USD.

4.1 Method

At the beginning of the study, we asked the two experts to provide a brief overview of their current roles and areas of expertise. As an introductory question, we also asked them to name design goals they consider important when designing AR interaction techniques. Then, we presented the *Future of Education* scenario (Sec. 3) and the three design dimensions (Sec. 3.1) to use as a reference throughout the study: the time/space matrix [20], considerations for public and private spaces, and a threat model considering access control of virtual content and physical spaces [49]. We provided the experts with a digital handout of the design dimensions (Sec. 3.1) to refer to throughout the study.

We structured each session into two tasks: (1) an iterative **production and design critique** cycle to elicit AR interaction techniques; (2) a **holistic review** of the interaction proposals with respect to our core design goals of usability, feasibility, and S&P. These tasks were followed by a debrief session to elicit experts’

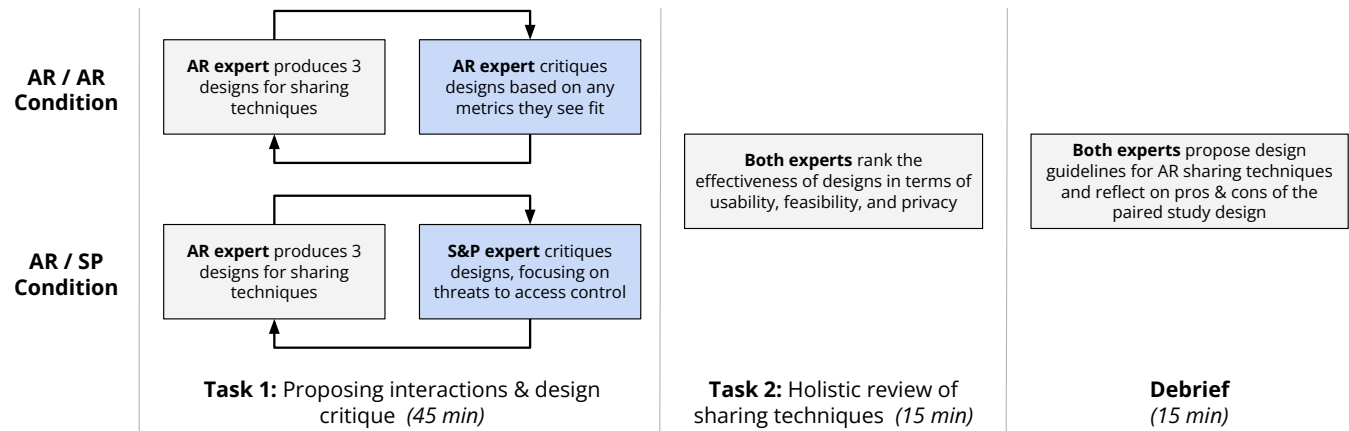


Figure 2: Our elicitation study is comprised of four stages where pairs of experts iteratively design and critique interaction techniques (Task 1), then assess the effectiveness of the techniques (Task 2) and overall design process (Debrief). The main difference between conditions is that in Task 1, the experts critique the techniques with respect to any metrics they see fit in the AR/AR condition, while in the AR/SP condition we specifically invoke the S&P expert to focus on privacy with their critique.

reflections on the pros and cons of each pairing. Figure 2 depicts the study design, which we detail below.

4.1.1 Task 1: Production and Design Critique of AR Sharing Techniques. In the first task, we adopted elements of the production, priming, and partners (PPP) method [33] to iteratively elicit and critique interaction proposals for sharing and controlling access to AR content. We facilitated turn-taking between the experts such that one expert first produced three interaction proposals, then the other expert provided a critique of the proposals. To systematically guide the experts through the *Future of Education* scenario, we posed four step-by-step prompts which each brought a new set of users into consideration (Table 1). This elicitation task was where a lot of creative energy was required and we wanted to allow exchange of ideas between the experts, so we allocated the majority of the time to this task (45 min).

For the production phase of each elicitation prompt, we instructed one expert to demonstrate their interaction proposals by thinking-aloud and annotating the scenario sketch using Google Jamboard². We named each technique (e.g., *Gaze & Point*) to create a point of reference throughout the study.

Then, we invoked the other expert to critique the interaction proposals, referring to the design dimensions handout as needed and defining additional design goals or quality metrics to assess the proposals as they saw fit. In the AR/SP condition, the AR expert always took on the role of designer and the S&P expert conducted the critique, whereas in the AR/AR condition, the AR experts swapped designing and critiquing roles after every prompt. While there was often an interesting debate between experts, rather than directly implementing revisions to the current interaction proposals based on the critiques at this stage, we encouraged participants to keep the feedback in mind as they continued working on the remaining elicitation prompts. For each subsequent prompt, we gave the experts the option of building on the existing interaction techniques or proposing new ones.

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

4.1.2 Task 2: Holistic Review of Interaction Techniques. The goal of the second task was to understand how the experts assessed the effectiveness of their interaction techniques with respect to the three design goals at the core of the investigation: usability, feasibility, and S&P. We asked them to compare three techniques elicited in Task 1 and rank them as more effective or less effective with respect to each design goal. We facilitated turn-taking between the experts, asking one expert to propose a ranking for each technique and asking the other expert to discuss whether they agree or disagree with the ranking. In the AR/SP condition, the AR expert proposed their ranking first, followed by a critique from the S&P expert.

4.1.3 Debrief. We concluded the study by asking the experts to reflect on the pros and cons of pairing two AR experts (AR/AR condition) or pairing AR and security & privacy experts (AR/SP condition) for an elicitation session.

4.2 Participants

We recruited 12 experts in AR design & development and 4 experts in security & privacy for our study (3 women, 11 men, average age of 29.3 years, 2 participants declined to answer). Our inclusion criteria were individuals with 2 or more first-authored publications in AR/VR or security & privacy or at least 2 years of industrial research experience in related fields. While some experts identified primarily as VR researchers (AR1, AR3), all experts reported having significant experience in interaction techniques common to both AR and VR interfaces (designing 3D spatial interactions, voice or gesture-based interfaces, etc.). We identified potential participants based on their recent publications (from venues including CHI, UIST, SOUPS, and IEEE Symposium on Security & Privacy) and invited them via email. Table 2 shows the experts' job roles and main areas of expertise, along with the condition and partner that they were assigned.

AR Experts				
Condition	Pair	ID	Job Role	Main Areas of Expertise
AR/AR	1	AR1	Assistant professor	shared and collaborative VR, including non-HMD users
AR/AR	1	AR2	Assistant professor	human perception, adaptive AR interfaces
AR/AR	2	AR3	Postdoctoral researcher	VR interaction techniques, camera networks
AR/AR	2	AR4	Assistant professor	XR workspaces, social acceptability
AR/AR	3	AR5	PhD student	mobile AR interactions
AR/AR	3	AR6	PhD student	asymmetric interactions for XR users & non-users
AR/AR	4	AR7	Research scientist	XR workspaces & productivity
AR/AR	4	AR8	Research scientist	XR interaction techniques
AR/SP	5	AR9	Assistant professor	AR sensing technologies
AR/SP	6	AR10	Research scientist	AR learning experiences, makerspaces
AR/SP	7	AR11	Assistant professor	human-AI systems, AR accessibility
AR/SP	8	AR12	PhD student	MR accessibility

Security & Privacy Experts				
Condition	Pair	ID	Job Role	Areas of Expertise
AR/SP	5	SP1	Postdoctoral researcher	usable security & privacy, security & privacy for IoT devices
AR/SP	6	SP2	Software engineer	AR security & privacy
AR/SP	7	SP3	PhD student	perceptions of security & privacy risks
AR/SP	8	SP4	Assistant professor	usable privacy & security, cybersecurity

Table 2: Participant information. We recruited 16 participants with expertise in a wide range of extended reality and security & privacy disciplines. Participants AR1-8 were paired with each other for the AR/AR condition, while AR9-12 were paired with SP1-4 for the AR/SP condition.

4.3 Data Collection & Analysis

We recorded and took notes on each study session, then followed a thematic analysis approach [7] to summarize the experts' design considerations and interaction techniques proposed throughout the study. To assemble an initial codebook, two of the authors analyzed one transcript from each study condition, identifying interaction techniques proposed during the elicitation task and paying attention to the experts' rationale for suggesting or critiquing particular interaction techniques. The two coders analyzed the remaining transcripts independently, then reviewed and aligned the codes. Then, they used the codes to categorize specific design considerations under overarching design goals (usability, feasibility, and S&P) and group the interaction techniques into similar modalities (e.g., gestures, proximity-based interactions) as shown in Table 3. Some design considerations had to be coded under multiple goals based on the experts' usage of the terms, e.g., scalability was mentioned both as a usability consideration for facilitating collaboration between more users and as an S&P consideration for translating legacy access control mechanisms to head-worn AR interfaces.

We used the codebook to generate a **timeline data visualization** for each expert pair (Fig. 3) to plot the design considerations mentioned for each elicitation prompt in chronological order. We note that experts mentioned design considerations with different levels of precision, which is also reflected in the timelines (e.g., security and confidentiality are shown as two separate design considerations under the overarching goal of S&P).

To determine the experts' underlying motivations in proposing techniques and how their consideration of design goals evolved throughout the study, we analyzed each interaction proposal alongside the timeline visualizations and identified design goals which they explicitly mentioned during the design and critique session. We further grouped techniques into two categories: (1) **techniques for sharing virtual content**, which were often motivated by usability

goals or to facilitate collaboration along the time/space matrix, and (2) **techniques for establishing access control policies**, which were suggested to mitigate threats raised in Task 1's back-and-forth critique sessions. Figure 3 shows a timeline for Group 8 from the AR/SP condition, which illustrates key interaction proposals that increasingly incorporated considerations for access control as the scenario evolved through the step-wise prompts.

To assess the extent to which the experts agreed upon the set of elicited interaction techniques, we utilized Morris et al.'s definition of **consensus-distinct ratio** [32] (the proportion of sharing techniques which were proposed by a minimum number of pairs), using a threshold of 50%. We calculated the consensus-distinct ratio separately for the AR/AR and AR/SP conditions (where techniques proposed by at least two out of four pairs achieved consensus), as well as for all eight pairs together (where techniques proposed by at least four out of eight pairs achieved consensus). The techniques which achieved the consensus threshold are indicated with a checkmark in Table 3. We selected the consensus-distinct ratio over other popular metrics (e.g., agreement score [59]) because prior work has shown the metric to handle cases when participants can suggest multiple symbols per referent in a more balanced manner [32].

5 RESULTS

We present our findings in four steps: (1) **characterizing the set of interaction proposals** across both conditions that were elicited for sharing AR content and were enriched with access control provisions; (2) **comparing trends in the design of interaction proposals** between AR/AR and AR/SP conditions; (3) **discussing design and revision strategies** adopted by the experts to incorporate considerations for access control in their interaction proposals; and (4) **reflecting on the effectiveness of the expert pairings** based on the debrief and comments from the experts.

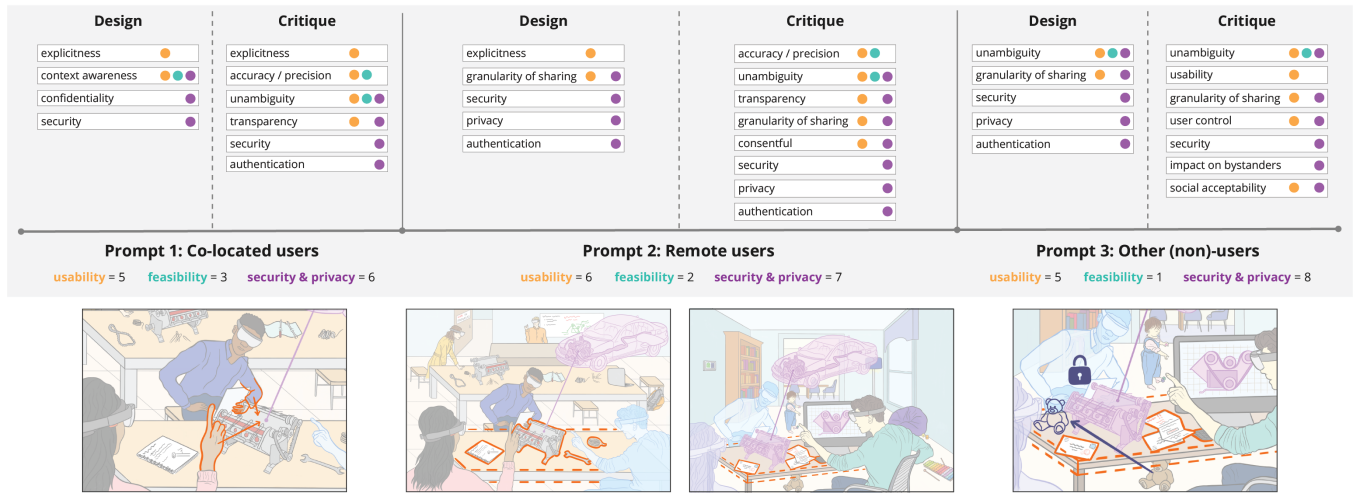


Figure 3: Example analysis for AR/SP condition. We show a recreation of Group #8’s interaction proposals alongside a portion of their timeline visualization, which depicts design considerations that the experts explicitly mentioned during the iterative design and critique task. The colored dots indicate to which design goal(s) each design consideration related (e.g., in the first row, explicitness referred to a usability consideration, accuracy / precision concerned both usability and feasibility). The number of distinct design considerations is shown below each prompt. For Prompt 1 (co-located users), AR12 designed a gesture for Russell and Violet to establish the engine as the virtual content anchor; we categorized this as a sharing technique because it was primarily based on the usability goal of explicitness. To share with Doug in Prompt 2 (remote users), AR12 proposed that collaborators draw a boundary where environmental data will be shared via depth capture. In Prompt 3 (other users / non-users), experts raised concerns that Ellie could interfere with the marker-less tracking by playing with the toys and moving around the room, which led to an interaction technique related to access control: “locking” the geometry of the table so the AR device disregards new geometry updates.

5.1 Popular Interactions across AR/AR and AR/SP Conditions

As a first step, we characterized the interaction techniques elicited in both the AR/AR and AR/SP conditions. Table 3 shows an overview of the techniques grouped by frequency for each condition, using consensus-distinct ratio [33] to indicate agreement among study pairs. We distinguish between techniques that were originally elicited for sharing virtual content in AR and those that were explicitly elicited to address threats related to access control.

First, we elicited a total of 73 interaction proposals for sharing virtual content in AR HMD-based environments, across all eight expert pairs and all four elicitation prompts. These sharing techniques ranged from more traditional or legacy [33] techniques, such as *virtual menus* (which was most the frequently proposed at 21 times), to more AR-specific interactions including *gestures* (17), *proxemic interactions* (16), *gaze* (14), and *voice* (11). AR-specific techniques were often motivated from prior work and directly cited by participants (e.g., zones from Slice of Light [57], collaborative authoring techniques from SpaceTime [60]). Proposals involving a *companion mobile app* or *physical tokens* were less frequent but shared between some study pairs. A smaller set of *distinct* techniques (i.e., only proposed by one pair [32]) included sweeping a virtual net to define a capture area (Pair #3), using notebooks as physical interfaces (#4), instrumenting AR headsets with touchpads for sharing (#5), and sharing all content placed on designated physical surfaces (#8).

Second, we identified 65 interaction proposals that were typically added onto the originally elicited sharing techniques to provide access control or balance requirements for access control with other design goals for collaboration. The most common techniques included *user authentication interfaces* (19 times), *notice and consent interfaces* (12), and *notifications* about collaborators’ activities (7). These more traditional techniques often resembled WIMP-style interfaces with minor adaptations for head-worn AR. A second group for enforcing access control policies, *role management interfaces*, *coarse-grained virtual objects*, and *timeouts* were also commonly proposed (5 times each). Among these were also AR-specific interactions, sometimes inspired from prior work, such as coarse-grained “ghost” objects [49] to improve other AR users’ awareness without fully revealing confidential AR content. Setting *inclusion/exclusion boundaries* was proposed by several pairs but was overall the least commonly proposed technique (3 times). There was also a group of *distinct* proposals (9 total) such as broadcasting audio to non-AR users as an awareness mechanism (Pair #2) and aggregating bystander data captured by AR devices (Pair #5).

Figure 4 shows the relationships between the original sharing techniques and the evolution into access control techniques, mapped on an axis from most popular to most distinct techniques (from left to right).

Virtual menus to share AR content (Fig. 4A) were proposed by all eight pairs and were often extended with *user authentication interfaces* for verifying users’ identities, *role management interfaces*

Interaction techniques for sharing virtual content					Consensus-Distinct (threshold=50%)		
Technique	AR/AR Frequency	AR/SP Frequency	Total	Pairs (#)	AR/AR=0.35	AR/SP=0.29	All=0.29
Virtual menus	9	12	21	1-8	✓	✓	✓
Gestures	7	10	17	1-5, 7, 8	✓	✓	✓
Proximity-based interactions	8	8	16	1-5, 7, 8	✓	✓	✓
Gaze	9	5	14	1-5, 8	✓	✓	✓
Voice	5	6	11	2-6	✓	✓	✓
Companion mobile app	2	2	4	3,4,6	✓		
Distinct (physical tokens, etc.)	7	3	10	1, 3-5, 8			

Interaction techniques for access control					Consensus-Distinct (threshold=50%)		
Technique	AR/AR Frequency	AR/SP Frequency	Total	Pairs (#)	AR/AR=0.31	AR/SP=0.31	All=0.31
User authentication interfaces	9	10	19	2-4, 6-8	✓	✓	✓
Notice & consent interfaces	4	8	12	2, 4-8	✓	✓	✓
Notifications	6	1	7	1,3,4,6	✓		✓
Role management interfaces	2	3	5	2,3,7,8	✓	✓	✓
Coarse-grained virtual objects	2	3	5	2, 4-7	✓	✓	✓
Timeouts	1	4	5	2,8			
Inclusion/exclusion boundaries	1	2	3	1,6,7		✓	
Distinct (silent speech, depth view, etc.)	3	6	9	2-6, 8			

Table 3: Interaction categories and modalities. We first grouped proposals into 2 categories (techniques for sharing virtual content and for addressing threats related to access control), then grouped techniques in each category by modality. The *Distinct* category includes techniques only suggested by 1 out of 8 pairs. Multimodal techniques are represented multiple times in the frequency column, e.g., a “look & point” technique is counted as both gaze and gesture. We calculated the *consensus-distinct* ratio, i.e., proportion of techniques which were proposed by 50% of pairs for each condition separately and together. For both conditions, 5 out of 17 sharing techniques (menus, gestures, proxemics, gaze, voice) and 5 out of 16 access control techniques (interfaces for user authentication, notice & consent, role management, notifications; coarse-grained virtual objects) achieved consensus (indicated by a checkmark).

for enforcing access control policies, *notifications* about collaborators’ activity, and *notice and consent interfaces* for accepting shared content from other AR users. Three pairs also proposed **companion mobile apps** (Fig. 4F) as an alternate modality to AR menus for sharing AR content or authentication. The second-most popular modality was **mid-air gestures** (Fig. 4B), which experts found beneficial for encouraging explicitness in collaborative scenarios and expressiveness (e.g., the same gesture could be used to express specific AR objects to share with specific collaborators). At the same time, gestures were considered less privacy-invasive than other AR-specific techniques (e.g., voice or gaze) if they were designed in a subtle way to disclose less detail about the AR application to bystanders or observers. To reduce the likelihood of hijacking attacks succeeding with gesture-based interactions (e.g., an adversary faking a gesture as if it were performed by the user), experts combined them with *timeouts* to control the time period when the AR app accepts continuous gestural input.

Proximity-based interactions (Fig. 4C), where AR content is automatically shared with AR users within the vicinity of the content’s owner, were considered intuitive and efficient for our collaboration scenario involving co-located students and the teacher. However, the experts expressed that proxemic interaction would not provide fine-grained sharing controls for both the physical and virtual space. To allow for more granular sharing of the physical environment, experts proposed *inclusion & exclusion boundaries* (e.g., Doug can draw an exclusion boundary around Ellie to prevent sharing information about her, even if she enters the defined zone).

Experts also proposed replacing AR content with *coarse-grained virtual objects* that show more basic representations (e.g., showing a virtual cube instead of the car engine model) to maintain confidentiality of the content while increasing other users’ awareness of virtual content placement in the physical space.

Gaze and voice interactions (Fig. 4D, E) were proposed by a majority of pairs to achieve natural and intuitive interactions while communicating with collaborators. However, gaze and voice were overall less popular than virtual menus, gestures, and proxemic interactions due to privacy concerns. For example, experts argued that eye tracking requires access to more sensitive biometric information as compared to gesture recognition; microphones have a higher likelihood of capturing bystanders due to a larger sensor range). To increase users’ awareness of physical areas captured by the AR device, Pair #5 combined gaze and voice with a *virtual map* depicting the spatial range that AR device sensors can capture. They further combined voice commands with the distinct technique of *silent speech*, where users breathe in while speaking voice commands to make them quieter (e.g., to prevent adversaries from overhearing voice commands and using them to gain access to the AR content). Pair #2 proposed a technique similar to coarse-grained virtual objects, namely *broadcasting spatial audio* from the AR app into the physical space, to increase awareness for non-users without revealing full details about confidential content (e.g., to make Ellie aware that her dad is communicating with remote collaborators in our scenario).




Figure 4: Evolution of AR sharing techniques into access control techniques. We illustrate the expert-elicited techniques for sharing AR content in HMD-based environments (mapped from most popular to most distinct, from left to right) and corresponding access control techniques that were added on to address threats. We indicate distinct techniques with (*). Virtual menus (A) were often paired with additional menus for user authentication, notifications, role management, and notice & consent. Gestures (B) were extended with timeouts, which limit the time period that AR devices accept gestural input. Proximity-based interactions (C) were extended through coarse-grained virtual objects (to maintain confidentiality of AR content while promoting awareness for other AR users) and inclusion/exclusion boundaries to define areas of the physical environment that can be captured. Gaze and voice commands (D, E) were combined with a virtual map of the sensor range to promote awareness of input hijacking attacks. For voice, experts also proposed silent speech to disclose less detail and broadcasting spatial audio, to promote awareness for bystanders while retaining some degree of confidentiality. Companion mobile apps (F) were paired with similar access control interfaces as virtual menus. Finally, distinct techniques involving physical tokens (G) led to distinct access control proposals for sharing depth views rather than RGB camera feeds and “locking” updates to AR devices’ spatial maps to prevent losing access to app content, in the event that marker-less tracking is disrupted.

Distinct interaction proposals included sharing through **physical tokens** (Fig. 4G), e.g., placing virtual content on designated physical “shared surfaces” or using virtual buttons embedded in students’ lab notebooks. Experts perceived these techniques to be more robust to threats related to access control, since users could physically protect or hide their personal tokens to prevent unwanted access. However, experts also noted potential implementation challenges with scaling physical tokens to large classroom settings while depicting shared surfaces with a high level of detail, which could require complex object recognition setups. These distinct sharing techniques also led to distinct proposals addressing threats related to access control. Pair #5 proposed sharing the AR device’s *depth view* instead of camera data, as this coarser-grained, diminished representation would disclose less detail about the physical environment. Pair #8 designed a *geometry filter* technique, where the AR device disregards changes in the physical environment to prevent the AR user from losing access to app content, e.g., if Ellie were to interfere with the marker-less tracking by playing with the toys and moving around the room.

5.2 Comparing Evolution & Creativity of Interaction Proposals between AR/AR and AR/SP Conditions

Next, we compared the elicited interaction techniques from each condition. As shown in Table 3, we observed that gaze and distinct techniques involving environmental tracking were proposed more frequently in the AR/AR condition. Virtual menus and gestures were more popular in the AR/SP condition. The experts agreed these techniques can be more secure and privacy-friendly since they only require gesture detection, as compared to other sharing techniques that require more invasive forms of biometric sensing or environmental mapping. This is indicative of different design strategies, e.g., prioritizing S&P design goals earlier on by discarding proposals with S&P weaknesses (discussed more in Sec. 5.3). The AR/SP pairs produced a greater frequency of access control techniques (e.g., user authentication and notice & consent interfaces); this is because the AR experts tended to address threats earlier on in the elicitation task given the S&P experts’ critiques, and carried their access control proposals over to later iterations. Notifications were more frequent among AR/AR pairs because they had the tendency

to first establish a base sharing technique such as virtual menus and proximity-based interactions, then add notifications to address threats related to access control.

We observed two major trends: *AR/AR* and *AR/SP* pairs proposed similar access control techniques though *AR/SP* pairs proposed them earlier, and the presence of S&P experts did not limit the creativity of AR experts in the *AR/SP* condition.

Similar access control techniques in both conditions. The *AR/SP* pairs often generated proposals to mitigate threats during Prompts 1 & 2, which led them to produce a higher frequency of access control techniques (37, as compared to 28 from the *AR/AR* condition). Examples include notice & consent interfaces and time-outs, which were proposed in both conditions but achieved higher frequencies among *AR/SP* pairs due to being proposed earlier and carried over to successive prompts (Table 3). Based on our timeline analysis, *AR/AR* pairs proposed a majority of their threat mitigation strategies later on during Prompt 3, which considered the bystanders and non-users in the scenario.

However, pairs in both conditions produced similar types of access control techniques. As shown in Table 3, four techniques achieved a consensus threshold in both conditions, i.e. proposed by 50% of pairs in each condition. Techniques which achieved consensus include legacy interfaces for user authentication (pairs #2-4, 6-8), role management (#2,3,7,8) and notice & consent (#2, 4-8), as well as a more creative approach of sharing coarse representations of virtual objects (#2, 4-7) to preserve the confidentiality of AR content while promoting awareness for other AR users.

AR designers' creativity not limited by the presence of S&P experts. In our debrief, S&P experts raised concerns that they could be “killjoys” (SP4), causing AR experts to “not deliver the same amount of creativity” as they would when designing interaction techniques on their own (SP1). Thus, one might expect that *AR/SP* pairs' interaction techniques were skewed towards legacy proposals (e.g., WIMP-style menus for authentication and role management) which are familiar from 2D interfaces and perceived as secure (discussed more in Sec. 5.3). However, our analysis showed that this was not the case: pairs in both conditions produced similar numbers of legacy, WIMP proposals (30 for *AR/AR*, 33 for *AR/SP*) and post-WIMP proposals which were more specific to AR interfaces, e.g., gesture and proximity-based inclusion boundaries (45 for *AR/AR*, 49 for *AR/SP*).

5.3 Design and Revision Strategies

Next, we analyzed the experts' higher level approaches to revising their techniques to incorporate considerations for access control, in order to better understand the similarities and differences that we observed in the type and frequency of interaction proposals. By using the design goal timeline visualizations to analyze the evolution of interaction proposals (Fig. 3), we identified two design strategies employed by the AR experts: a *usability-first strategy* and *security & privacy-first strategy*.

Usability-first strategy: enhancing initial sharing technique proposals by adding interaction modalities with access control techniques. While we gave the experts the option of building on their existing techniques or proposing new ones for each successive prompt, many first optimized their interaction proposals

for usability goals (e.g., explicitness and flexibility) and gradually added on access control techniques to mitigate threats (Fig. 4). Eight out of twelve AR experts adopted this design approach: six from the *AR/AR* condition (AR1-4, AR6-7) and two from the *AR/SP* condition (AR9, AR12). In describing the benefits of this approach, AR9 argued that prioritizing usability before S&P is most practical when designing interactions techniques; while “it's hard to optimize for multiple parameters at the same time,” they could “always find an alternative way to make [the existing techniques] more privacy-preserving.”

As a result of this design strategy, a significant number of sharing techniques were multimodal (40 out of 73). The most agreed-upon multimodal sharing techniques involved menus and proximity, gaze and proximity, and combining gaze, menus, and gestures (each proposed by two pairs). While the experts gravitated towards this subset of common AR modalities, the specific mechanics of these interactions were often unique to each pair. Experts noted the potential disadvantage with multimodal techniques imposing additional software requirements; some distinct techniques may require new hardware components as well (e.g., AR9's proposal to add new physical buttons to the AR HMD).

Security & privacy-first strategy: discarding techniques with S&P weaknesses early on. The remaining four experts from the *AR/AR* (AR5, AR8) and *AR/SP* conditions (AR10-11) designed defensively from the first elicitation prompt to mitigate threats related to access control. They often weighed multiple ideas and anticipated critiques that their partners may raise before proceeding with an interaction proposal, discarding techniques which had security weaknesses or were privacy-invasive. For example, AR10 initially considered voice commands, but discarded the proposal due to concerns of adversaries using speech interjection to gain access to the AR app; AR11 commented that “[SP3] will probably critique” their authentication proposal based on facial recognition.

A potential disadvantage we observed with the S&P-first design approach was a delayed prioritization of usability goals. The final interaction proposals tended to include legacy techniques (e.g., authentication and role management interfaces), which are standard approaches for access control in 2D form factors. However, during the holistic ranking task, some expert pairs ranked these proposals lower in terms of usability due to AR HMDs' limitations of pointing accuracy and small field-of view [5].

The usability-first and S&P-first design strategies were utilized by experts in both conditions. While we might expect the S&P-first strategy to be tied to the *AR/SP* condition, we found this was not the case. We observed instances where the *AR/AR* pairs first optimized for threats related to access control (AR5, AR8) and discarded techniques they perceived to be privacy-invasive. We also observed cases in the *AR/SP* condition where AR experts initially optimized for usability and gradually incorporated access control techniques, despite receiving a thorough threat analysis from the start of the elicitation task (AR9, AR12). This suggests that the presence of a S&P expert was not always required to encourage defensive design; the scenario and step-wise elicitation prompts that included bystanders and non-users provided sufficient cues for S&P-aware design from the very start of the elicitation task for some *AR/AR* pairs.

Summary of Results (Section 5)	
5.1	Across all 8 pairs, we elicited 73 proposals for sharing techniques and 65 proposals for enhancing these techniques with access control provisions. Legacy techniques were the most popular (e.g., virtual menus); distinct techniques were often scenario-specific and made use of the physical environment.
5.2	The <i>AR/SP</i> pairs produced access control techniques earlier on, but pairs in both conditions produced similar types of access control techniques. The AR designers' creativity was not limited by the presence of S&P experts.
5.3	Pairs in both conditions utilized the usability-first strategy (first optimizing interaction proposals for usability goals, gradually adding on access control techniques) and the S&P-first strategy (discarding techniques with S&P weaknesses early on).
5.4	The <i>AR/AR</i> pairing was beneficial for nuanced AR interaction design, but poses risks of "tunnel vision" without more diverse expertise. The <i>AR/SP</i> pairing encouraged security & privacy as top-level design goals, but could be time and resource-demanding for design teams.

Table 4: Summary of Results: In Section 5, we presented our analysis of the characterization of the interaction proposals (5.1), trends in the design of interaction techniques between conditions (5.2), design and revision strategies adopted by the experts (5.3), and experts' reflections on the effectiveness of each pairing (5.4).

5.4 Effectiveness of the Expert Pairings

In our debrief, we also asked the participants to share their opinions on the pros and cons of their pairings.

AR/AR pairing beneficial for nuanced AR interaction design; risk of "tunnel vision" without more diverse expertise. Experts in the *AR/AR* condition found value in building on each others' ideas (AR3, AR5-6) and challenging each others' opinions from an AR interaction design perspective (AR1-3, AR7-8). AR2 thought that even with two AR experts, they approached the design process "from multiple angles" due to nuances in their research interests (e.g., balancing explicitness and social acceptability in the case of AR3 and AR4). However, AR3 raised the issue of a "shared tunnel vision" as two AR experts who prioritize similar goals "may not gain much from each other," while experts from different fields would "have the opportunity to leave the tunnel for a moment." AR6 expressed that having overly similar expertise and agreeing on interaction proposals could result in their "confidence getting artificially reinforced," potentially preventing pairs from analyzing their proposals with a critical lens.

AR/SP pairing encourages security & privacy as top-level design goals, but could be resource-demanding. The *AR/SP* pairs mentioned benefits of diverse perspectives (AR10-12, SP1-3) in encouraging "privacy and security as top level concepts of design" rather than as a "last mile approach" (SP4). However, they also expressed that involving a S&P expert could potentially be resource-demanding by incurring a time cost in communicating critiques and iterating upon interaction proposals (AR9-12, SP1, SP3). AR9 discussed a potential "opportunity cost at the beginning" of the interaction designers' workflows, as they might abandon proposals primarily on the basis of S&P weaknesses and without weighing the benefits from an interaction design perspective. However, AR9 (who we categorized as a usability-first designer) argued they could usually find ways to extend their proposals to make them more privacy-preserving, so the S&P analysis "did not add too much overhead" to their design process.

6 DISCUSSION

Overall, all participants found it valuable to work with expert partners to iteratively design and critique interaction techniques through our scenario-based elicitation approach. It was encouraging that pairs in both the *AR/AR* and *AR/SP* conditions produced similar

types of interaction techniques for sharing and access control, with both legacy proposals and more creative techniques specific to the *Future of Education* scenario. A summary of our results is shown in Table 4. In this section, we (1) reflect on and extract key elements of our study design that enabled AR experts to design techniques which incorporate S&P considerations for access control, as one example, (2) present design recommendations for AR sharing techniques that balance competing design goals for collaboration and access control, and (3) discuss limitations of our approach.

6.1 Incorporating S&P Considerations in Elicitation Studies

In this work, our goal was to encourage AR interaction designers to consider potential threats by making S&P design goals explicit within the elicitation process. We still aimed to enable designers to understand tradeoffs with usability and technical feasibility, which have been the focus of prior elicitation studies. It was encouraging that we observed no major differences between the *AR/AR* and *AR/SP* conditions in terms of the types of access control techniques and creativity of proposals. This suggests that the scenario-based elicitation process, facilitated through guided prompts which traverse the time/space matrix and raise threats to access control involving different AR users and bystanders, provided sufficient structure and cues for AR experts alone to design more thoughtfully. The key design decisions that we made to achieve this result were grounding the elicitation process within a multi-user AR scenario, crafting prompts which increasingly introduce potential threats along a threat model [49], and structuring the iterative design and critique of interaction proposals through assigning specific roles to partners and facilitating turn-taking between them.

Scenario-based design: We previously piloted a more open-ended approach similar to prior gesture elicitation studies, where participants designed sharing techniques solely based on generic system operations [32, 40, 59]. However, we observed that our pilot testers struggled to brainstorm interaction techniques without concrete details on the collaboration context; they eventually developed their own scenarios involving specific physical environments and characters to ground their interaction proposals. This posed issues for experimental control and analysis of results. Inspired by prior work in the usable privacy domain [39, 61], we made our scenario-based elicitation approach more explicit by increasing specificity

and applying stricter constraints, e.g., by depicting sensitive objects in the environment, incorporating both AR users and non-user and establishing their roles as collaborators or adversaries.

Facilitation strategy: We note our potential impact on the results given the stricter facilitation that was required by our study design, to understand how participants addressed the prompts and considered S&P issues as the scenario unfolded. While this is a possible limitation, facilitation was critical to establish a step-by-step process and help “jumpstart” the initial elicitation prompts, where participants often required clarification on what assumptions they could make about the scenario (e.g., whether Violet and Russell are already in a shared session). After establishing a back-and-forth cadence between the experts, we transitioned to facilitating less and observed how their collaboration strategy evolved throughout the rest of the design process. We would thus say that our facilitation did not require S&P expertise and is in fact common to scenario-based design [9].

The need for S&P expertise: While our elicitation approach yielded access control techniques of a similar type and quality in both conditions, we do not mean to say that our elicitation approach can replace the expertise of S&P professionals. Most of the expert participants recommended the pairing of AR and S&P experts and emphasized them working together “from the beginning to get used to each others’ opinions” (SP1). However, when it is not feasible to directly include S&P professionals, our work provides an example of how to adopt an S&P-minded design process that incorporates a structured framework for identifying and mitigating threats in line with threat models from prior work.

In practice, AR design teams can take the following steps to organize and conduct an elicitation study like ours: (1) **Create a scenario visualization** that depicts the typical usage context and potential AR users and non-users. While this visualization can build on top of existing design artifacts (e.g., storyboards and user personas), designers may need to brainstorm additional bystander personas whose S&P needs are important to consider [15]. (2) **Assign roles of designer, critic, and facilitator.** Similar to our elicitation study, the designer and critic roles could be filled by participants who prioritize different or competing design goals (e.g., interaction designers for usability, developers for feasibility) in addition to S&P. The facilitator should have an overview of these different design goals and work to balance the participants’ discussions when necessary. (3) **Choose an appropriate threat model.** For multi-user scenarios especially, design teams could adopt the same threat model involving AR access control that we used in our study; Ruth et al.’s paper [49] demonstrated the generalizability of the threat model to many collaborative scenarios along the time / space matrix with opt-in or opt-out sharing policies. However, if access control is not a top-level design goal for the particular AR use case, design teams could draw on other threat models from prior work (e.g., performing sensory manipulations to inflict harm on end-users in extended reality [54], privacy considerations involving data flows in AR [11]). While establishing alternate threat models representing critical threats for AR use cases may require consulting others with expertise in S&P and threat modeling, these experts would not be required to participate in the elicitation process.

6.2 Design Recommendations for Access Control in Multi-User AR Interactions

In the experts’ analysis of sharing techniques (Task 2), there was no commonly agreed-upon set of “best” techniques, i.e., ones that expert pairs rated as highly effective for all three design goals. Taking our holistic review approach, the experts often identified tradeoffs between design goals for facilitating collaboration (e.g., explicit vs. opportunistic interaction) and requirements for access control. In this section, we provide recommendations for interaction techniques to achieve a balance between these potentially conflicting design goals, based on the experts’ discussions.

Making interactions more obvious yet secure for collaborators. A majority of AR experts emphasized the need for *explicit* interactions which “communicate the state” of the application (AR2) and provide awareness for collaborators in multi-user scenarios (AR1-4, AR7-8, AR12), e.g., “obvious gestures” to indicate when group members are leaving or joining the shared session (AR4). However, a disadvantage of explicit interactions is enabling new threats from co-located adversaries who are directly observing AR users. Possible attacks include shoulder surfing [16] to learn private information about the AR users’ virtual content or learn how to perform key input techniques to manipulate virtual content (e.g., specific voice commands to delete AR content).

Experts suggested three strategies to allow for a greater degree of explicitness while preventing adversaries from learning private information or gaining unwanted access to AR content: (1) choosing interactions which “disclose less detail” about what the AR user is trying to accomplish (SP1), e.g., choosing mid-air gestures over voice commands; (2) making sharing techniques increasingly multimodal in order to make an attack more difficult e.g., Pair #2 proposed a “point to the stars” technique combining gesture, voice, and gaze; (3) combining sharing techniques with *timeouts* to prevent input hijacking (e.g., giving the user 3 seconds to perform a gesture, so that adversaries cannot mimic the AR user’s input at a later time to trigger an unwanted action). While these multimodal techniques were ranked as more secure by our experts in the holistic review task, we note it is possible for them to add complexity from an interaction design perspective and potentially have an adverse effect on usability.

Balancing access control with opportunistic interaction. In our *Future of Education* scenario, we clearly defined whether each character should have access to the AR users’ virtual simulation content. However, some expert pairs extended the scenario to discuss *opportunistic interactions* which could be desirable for educational contexts (#1, 2, 4), e.g., encouraging collaboration between different student teams to support learning or promoting bystanders’ awareness of AR users’ actions. These pairs expressed concerns that access control techniques could be “exclusionary” and result in “removing the happenstance” in classroom interactions (AR4).

To still enable opportunistic interaction while protecting access to virtual content, the experts proposed: (1) *proximity-based sharing* which allows non-collaborators to view more details about the AR user’s virtual space when they walk nearby, while notifying the current collaborators with visual or audio feedback (Pair #1); (2) displaying *coarse-grained representations of AR objects* to other AR users to promote their awareness and open opportunities for them

to request to join the AR user's experience (Pairs #2, 4-7). Since these interaction techniques involve a tradeoff between confidentiality and awareness (e.g. revealing the basic shape and placement of AR content within a physical space), the experts argued they could be ideal for "lower stakes" collaborative scenarios (e.g., educational contexts where "there's a greater sense of safety" due to being familiar with your teacher and classmates, as SP2 reported).

6.3 Limitations

We discuss limitations of our work with respect to the study sample, lack of baseline comparison to assess the quality of interaction proposals, and generalizability of the sharing techniques and overall elicitation approach. We also note that our research team includes individuals with expertise in AR and S&P; authors from other academic backgrounds might have envisioned the study design and interpreted its results differently.

Limitations of the study sample: Different from traditional elicitation, we studied with individuals with expertise in AR and S&P topics, rather than with non-technical designers or end-users as in traditional elicitation studies [2, 32, 37]. We made this decision to balance the conditions by ensuring that participants have a comparable level of experience with AR interaction design or threat modeling. However, studying with designers or end-users could yield different interaction sets and new requirements for the scenario-based elicitation process, which future work could explore.

Our between-subjects design required fewer participants with expertise in S&P (4) and more AR experts (12). This may be perceived as an imbalance; however, we did not reuse the S&P experts and the results did not appear to be skewed in the direction of the AR experts. We also considered that participants' positive feedback on our elicitation method could be a sign of participant response bias [12]. However, to lessen the effects, we explicitly asked them to comment on limitations of the process and the pairing.

We studied with a smaller sample than prior elicitation studies adopting the production, priming, and partners (PPP) approach [33] (i.e., 16 participants compared to 25 for Web on the Wall [32]). This tradeoff was partly due to working with AR and S&P experts, who were a harder population to access than end-users due to their specialized backgrounds. However, working with the 16 experts still allowed us to elicit a rich set of interaction techniques, both in terms of quantity (138 total) and range (legacy vs. post-WIMP). Comparing the experts' interaction proposals and discussions with the novice AR designers' responses in our pilot study, we believe working with experts resulted in higher quality proposals and nuanced discussions, especially around technical feasibility.

Lack of comparison to a baseline: Similar to prior work that derived new interaction techniques for novel application scenarios through user-driven elicitation [3, 32], there is no baseline that could serve as a common benchmark or universal point of comparison to assess the quality of the suggested interactions. However, we adopted three main measures to improve the quality of interaction proposals based on prior elicitation studies [33, 59]. (1) We worked with experts who have significant experience in S&P or AR development, rather than end-users. Being familiar with the state-of-the-art AR technologies, the experts frequently picked known techniques from prior work and commercial toolkits as a starting point, trying

to strike a balance between usability, feasibility, and S&P for our scenario. (2) We further assessed the quality of interaction proposals through the consensus-distinct ratio established by Morris et al. [32], in addition to asking the experts to assess the effectiveness of the sharing techniques in the holistic review (Task 3). We find it encouraging that a majority of both sharing and access control techniques achieved the consensus threshold (i.e. proposed by 50% of expert pairs in each condition). (3) We assessed the proportion of legacy vs. post-WIMP interactions to add context to the experts' discussions and understand their rationale for emphasizing security over novelty in some cases.

Generalizability of sharing techniques and elicitation approach: Due to our scenario-based approach, some interactions designed around the *Future of Education* scenario may be limited in their generalizability to other AR use cases, particularly those requiring specific environmental features (e.g., using tables as shared surfaces, passing the wrench to add collaborators). However, abstracting from the specifics of our scenario, our study yielded techniques similar to those in the research literature (e.g., collaborative authoring techniques from SpaceTime [60] and Slice of Light [57], ghost objects from Ruth et al. [49]) and those already available in some commercial AR HMD-based interfaces (e.g., virtual menus, gestures, proximity-based interactions), which suggests some generalizability to other AR use cases. To further improve generalizability and creativity of techniques, future work could study alternate scenarios and add a priming task to focus on new design goals that enable a wider exploration of the design space.

We do not claim generalizability of our scenario-based elicitation approach beyond content sharing in multi-user, head-worn AR environments. We focused on AR HMDs to strike a balance between realistic and future-facing interaction techniques, as they represent a class of current devices (e.g., HoloLens 2, Meta Quest with Passthrough) that have not yet achieved mass adoption due to high cost, but are widely used in the HCI research community. This enabled us to elicit both legacy and creative interaction proposals that the experts agreed were feasible to implement, based on their prior experiences developing AR apps for these devices. While some sharing techniques we elicited are already common to mobile and projective AR (e.g., proxemic interactions), we expect that different techniques are required to promote S&P goals in these other AR form factors. For example, interaction techniques in mobile AR (e.g., multi-touch gestures) are often subtler as compared to present-day AR HMDs, potentially making it difficult for bystanders to recognize mobile AR users in-the-wild; this warrants more techniques to promote bystander awareness and privacy (e.g., enforcing policies through Bluetooth signals that limit mobile AR users' access to capture sensitive physical areas [48]). Since AR content can be viewed by any co-located people in projective AR experiences, displaying private AR content that only specific users can access requires custom solutions (e.g., using a projector with an AR HMD to display shared and private views in AR [18]).

We believe that in principle, our scenario-based approach could be extended to single-user AR use cases (e.g., by considering the primary AR user and non-users one at a time). This may require adjusting or extending the threat model and developing new prompts (e.g., with a focus on harms involving bystanders), but would not require changing our overall approach.

7 CONCLUSION

In this paper, we demonstrate our approach extending user-driven elicitation to design techniques for sharing AR content, while considering potential security & privacy threats related to access control of the virtual and physical spaces. We contribute a set of multi-user AR sharing techniques enhanced with access control provisions and insights from our comparative elicitation study with 16 participants, where we explore the effects of pairing two AR experts with each other versus pairing an AR expert with a S&P expert. The studies were promising in that our scenario-based approach, with elicitation prompts that increasingly incorporated threats involving access control, encouraged pairs from both conditions to design interaction techniques with a critical lens and consider the S&P implications of using AR from users' and non-users' perspectives.

While we studied with experts rather than end-users or novice designers, we are exploring how to adapt a similar elicitation method to teach designers about S&P considerations for AR in future work. We also encourage future work to extend our approach and explore how other threat models, e.g., Guzman et al.'s categorization of mixed reality data flows [11], could guide the elicitation process.

ACKNOWLEDGMENTS

We thank the AR and security & privacy experts who participated in our study for engaging in creative and thoughtful elicitation sessions. We thank Florian Schaub and our anonymous reviewers for their valuable feedback to improve our paper. This work was supported in part by the National Science Foundation under Award CNS-1651230.

REFERENCES

- [1] 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, Berkeley, CA, USA, 427–442. <https://www.usenix.org/conference/soups2018/presentation/adams>
- [2] Abdullah X. Ali, Meredith Ringel Morris, and Jacob O. Wobbrock. 2019. Crowdlicit: A System for Conducting Distributed End-User Elicitation and Identification Studies. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, Article 255, 12 pages. <https://doi.org/10.1145/3290605.3300485>
- [3] Abdullah X. Ali, Meredith Ringel Morris, and Jacob O. Wobbrock. 2021. "I Am Iron Man": Priming Improves the Learnability and Memorability of User-Elicited Gestures. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 359, 14 pages. <https://doi.org/10.1145/3411764.3445758>
- [4] Huidong Bai, Prasanth Sasikumar, Jing Yang, and Mark Billinghurst. 2020. A User Study on Mixed Reality Remote Collaboration with Eye Gaze and Hand Gesture Sharing. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 423, 13 pages. <https://doi.org/10.1145/3313831.3376550>
- [5] Anil Ufuk Batmaz and Wolfgang Stuerzlinger. 2019. Effects of 3D Rotational Jitter and Selection Methods on 3D Pointing Tasks. *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* 2019 (2019), 1687–1692.
- [6] Hrvoje Benko, Ricardo Jota, and Andrew Wilson. 2012. MirageTable: Freehand Interaction on a Projected Augmented Reality Tabletop. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 199–208. <https://doi.org/10.1145/2207676.2207704>
- [7] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101.
- [8] Yuanzhi Cao, Xun Qian, Tianyi Wang, Rachel Lee, Ke Huo, and Karthik Ramani. 2020. An Exploratory Study of Augmented Reality Presence for Tutoring Machine Tasks. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 559, 13 pages. <https://doi.org/10.1145/3313831.3376688>
- [9] John M. Carroll. 1999. Five Reasons for Scenario-Based Design. In *Proceedings of the Thirty-Second Annual Hawaii International Conference on System Sciences - Volume 3 - Volume 3 (HICSS '99)*. IEEE Computer Society, USA, 3051.
- [10] Kaiming Cheng, Jeffery Tian, Tadayoshi Kohno, and Franziska Roesner. 2023. Exploring User Reactions and Mental Models Towards Perceptual Manipulation Attacks in Mixed Reality. In *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9–11, 2023*. USENIX Association, Berkeley, CA, USA, 18. <https://www.usenix.org/conference/usenixsecurity23/presentation/cheng>
- [11] 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>
- [12] Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. 2012. "Yours is Better!": Participant Response Bias in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 1321–1330. <https://doi.org/10.1145/2207676.2208589>
- [13] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In Situ with Bystanders of Augmented Reality Glasses: Perspectives on Recording and Privacy-Mediating Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 2377–2386. <https://doi.org/10.1145/2556288.2557352>
- [14] Arindam Dey, Mark Billinghurst, Robert W. Lindeman, and J. Edward Swan II. 2016. A Systematic Review of Usability Studies in Augmented Reality between 2005 and 2014. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*. IEEE, New York, NY, USA, 49–50. <https://doi.org/10.1109/ISMAR-Adjunct.2016.0036>
- [15] 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, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5228–5239. <https://doi.org/10.1145/2858036.2858214>
- [16] Malin Eiband, Mohamed Khamis, Emanuel von Zeszczwitz, Heinrich Hussmann, and Florian Alt. 2017. Understanding Shoulder Surfing in the Wild: Stories from Users and Observers. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 4254–4265. <https://doi.org/10.1145/3025453.3025636>
- [17] Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. 2017. 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, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 4021–4033. <https://doi.org/10.1145/3025453.3025683>
- [18] Jeremy Hartmann, Yen-Ting Yeh, and Daniel Vogel. 2020. AAR: Augmenting a Wearable Augmented Reality Display with an Actuated Head-Mounted Projector. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 445–458. <https://doi.org/10.1145/3379337.3415849>
- [19] Roberto Hoyle, Robert Templeman, Steven Armes, Denise L. Anthony, David J. Crandall, and Apu Kapadia. 2014. Privacy Behaviors of Lifeloggers Using Wearable Cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Seattle, Washington) (UbiComp '14). Association for Computing Machinery, New York, NY, USA, 571–582. <https://doi.org/10.1145/2632048.2632079>
- [20] Robert Johansen. 1988. *GroupWare: Computer Support for Business Teams*. The Free Press, USA.
- [21] Brett R. Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira. 2014. RoomAlive: Magical Experiences Enabled by Scalable, Adaptive Projector-Camera Units. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 637–644. <https://doi.org/10.1145/2642918.2647383>
- [22] Marion Koelle, Matthias Kranz, and Andreas Möller. 2015. Don't Look at Me That Way! Understanding User Attitudes Towards Data Glasses Usage. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Copenhagen, Denmark) (MobileHCI '15). Association for Computing Machinery, New York, NY, USA, 362–372. <https://doi.org/10.1145/2785830.2785842>
- [23] Balasaravanan Thoravi Kumaravel, Fraser Anderson, George W. Fitzmaurice, Björn Hartmann, and Tovi Grossman. 2019. 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 (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 161–174. <https://doi.org/10.1145/3332165.3347872>
- [24] Balasaravanan Thoravi Kumaravel and Andrew D. Wilson. 2022. DreamStream: Immersive and Interactive Spectating in VR. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 636, 17 pages. <https://doi.org/10.1145/3491102.3517508>
- [25] Lutz Lammerding, Tim Hilken, Dominik Mahr, and Jonas Heller. 2021. Too Real for Comfort: Measuring Consumers' Augmented Reality Information Privacy Concerns. In *Augmented Reality and Virtual Reality*, M. Claudia tom Dieck, Timothy H. Jung, and Sandra M. C. Loureiro (Eds.). Springer International Publishing, Cham, 95–108.
- [26] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2017. Securing Augmented Reality Output. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22–26, 2017*. IEEE Computer Society, New York, NY, USA, 320–337. <https://doi.org/10.1109/SP.2017.13>
- [27] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2018. Towards Security and Privacy for Multi-user Augmented Reality: Foundations with End Users. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21–23 May 2018, San Francisco, California, USA*. IEEE Computer Society, New York, NY, USA, 392–408. <https://doi.org/10.1109/SP.2018.00051>
- [28] Gun A. Lee, Jonathan Wong, Hye Sun Park, Jin Sung Choi, Chang-Joon Park, and Mark Billinghurst. 2015. User Defined Gestures for Augmented Virtual Mirrors: A Guessability Study. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI EA '15). Association for Computing Machinery, New York, NY, USA, 959–964. <https://doi.org/10.1145/2702613.2732747>
- [29] Chuan-en Lin, Ta Ying Cheng, and Xiaojuan Ma. 2020. ARchitect: Building Interactive Virtual Experiences from Physical Affordances by Bringing Human-in-the-Loop. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 487, 13 pages. <https://doi.org/10.1145/3313831.3376614>
- [30] David Lindbauer, Jens Emil Grønbaek, Morten Birk, Kim Halskov, Marc Alexa, and Jörg Müller. 2016. Combining Shape-Changing Interfaces and Spatial Augmented Reality Enables Extended Object Appearance. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 791–802. <https://doi.org/10.1145/2858036.2858457>
- [31] Abraham Hani Mhaidli and Florian Schaub. 2021. Identifying Manipulative Advertising Techniques in XR Through Scenario Construction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 296, 18 pages. <https://doi.org/10.1145/3411764.3445253>
- [32] Meredith Ringel Morris. 2012. Web on the Wall: Insights from a Multimodal Interaction Elicitation Study. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces* (Cambridge, Massachusetts, USA) (ITS '12). Association for Computing Machinery, New York, NY, USA, 95–104. <https://doi.org/10.1145/2396636.2396651>
- [33] Meredith Ringel Morris, Andreea Danieleescu, 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>
- [34] Pardis Emami Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. 2020. Ask the Experts: What Should Be on an IoT Privacy and Security Label?. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18–21, 2020*. IEEE, New York, NY, USA, 447–464. <https://doi.org/10.1109/SP40000.2020.00043>
- [35] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujio Bauer, Lorrie Faith Cranor, and Norman M. Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Thirteenth Symposium on Usable Privacy and Security, SOUPS 2017, Santa Clara, CA, USA, July 12–14, 2017*. USENIX Association, Berkeley, CA, USA, 399–412. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini>
- [36] Michael Nebeling and Anind K. Dey. 2016. XDBrowser: User-Defined Cross-Device Web Page Designs. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5494–5505. <https://doi.org/10.1145/2858036.2858048>
- [37] Michael Nebeling, Alexander Huber, David Ott, and Moira C. Norrie. 2014. 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* (Dresden, Germany) (ITS '14). Association for Computing Machinery, New York, NY, USA, 15–24. <https://doi.org/10.1145/2669485.2669497>
- [38] 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 *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 107, 12 pages. <https://doi.org/10.1145/3411764.3445323>
- [39] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. 2018. Turtles, Locks, and Bathrooms: Understanding Mental Models of Privacy Through Illustration. *Proc. Priv. Enhancing Technol.* 2018, 4 (2018), 5–32. <https://doi.org/10.1515/popets-2018-0029>
- [40] Thammathip Piumsomboon, Adrian J. Clark, Mark Billinghurst, and Andy Cockburn. 2013. 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 (Lecture Notes in Computer Science, Vol. 8118)*. Springer, New York, NY, USA, 282–299. https://doi.org/10.1007/978-3-642-40480-1_18
- [41] Thammathip Piumsomboon, Gun A. Lee, Andrew Irlitti, Barrett Ens, Bruce H. Thomas, and Mark Billinghurst. 2019. 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* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–17. <https://doi.org/10.1145/3290605.3300458>
- [42] Lev Poretzki, Joel Lanir, and Ofer Arazy. 2018. Normative Tensions in Shared Augmented Reality. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 142 (nov 2018), 22 pages. <https://doi.org/10.1145/3274411>
- [43] Iulian Radu and Bertrand Schneider. 2019. What Can We Learn from Augmented Reality (AR)? Benefits and Drawbacks of AR for Inquiry-Based Learning of Physics. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, Article 544, 12 pages. <https://doi.org/10.1145/3290605.3300774>
- [44] Shwetha Rajaram and Michael Nebeling. 2022. Paper Trail: An Immersive Authoring System for Augmented Reality Instructional Experiences. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 382, 16 pages. <https://doi.org/10.1145/3491102.3517486>
- [45] Derek F. Reilly, Mohamad H. Salimian, Bonnie MacKay, Niels Mathiasen, W. Keith Edwards, and Juliano Franz. 2014. SecSpace: Prototyping Usable Privacy and Security for Mixed Reality Collaborative Environments. In *Proceedings of the 2014 ACM SIGCHI Symposium on Engineering Interactive Computing Systems* (Rome, Italy) (EICS '14). Association for Computing Machinery, New York, NY, USA, 273–282. <https://doi.org/10.1145/2607023.2607039>
- [46] Franziska Roesner and Tadayoshi Kohno. 2021. Security and Privacy for Augmented Reality: Our 10-Year Retrospective. In *VR4Sec: 1st International Workshop on Security for XR and XR for Security (at CHI 2021)*. Association for Computing Machinery, New York, NY, USA.
- [47] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (2014), 88–96. <https://doi.org/10.1145/2580723.2580730>
- [48] 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, New York, NY, USA, 1169–1181. <https://doi.org/10.1145/2660267.2660319>
- [49] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2019. 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*. USENIX Association, Berkeley, CA, USA, 141–158. <https://www.usenix.org/conference/usenixsecurity19/presentation/ruth>
- [50] Kihoon Son, Hwiwon Chun, Sojin Park, and Kyung Hoon Hyun. 2020. C-Space: An Interactive Prototyping Platform for Collaborative Spatial Design Exploration. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 325, 13 pages. <https://doi.org/10.1145/3313831.3376452>
- [51] Maximilian Speicher and Michael Nebeling. 2018. GestureWiz: A Human-Powered Gesture Design Environment for User Interface Prototypes. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, Article 107, 11 pages. <https://doi.org/10.1145/3173574.3173681>
- [52] Hariharan Subramonyam, Steven Mark Drucker, and Eytan Adar. 2019. Affinity Lens: Data-Assisted Affinity Diagramming with Augmented Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, Article 398, 13 pages. <https://doi.org/10.1145/3290605.3300628>
- [53] Theophilus Teo, Louise M. Lawrence, Gun A. Lee, Mark Billinghurst, and Matt Adcock. 2019. Mixed Reality Remote Collaboration Combining 360 Video and 3D Reconstruction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, Article 201, 14 pages. <https://doi.org/10.1145/3290605.3300431>

- [54] Wen-Jie Tseng, Elise Bonnal, Mark McGill, Mohamed Khamis, Eric Lecolinet, Samuel Huron, and Jan Gugenheimer. 2022. The Dark Side of Perceptual Manipulations in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 612, 15 pages. <https://doi.org/10.1145/3491102.3517728>
- [55] Ana M. Villanueva, Zhengzhe Zhu, Ziyi Liu, Kylie Pepler, Thomas Redick, and Karthik Ramani. 2020. Meta-AR-App: An Authoring Platform for Collaborative Augmented Reality in STEM Classrooms. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 19, 14 pages. <https://doi.org/10.1145/3313831.3376146>
- [56] Santiago Villarreal-Narvaez, Jean Vanderdonckt, Radu-Daniel Vatavu, and Jacob O. Wobbrock. 2020. A Systematic Review of Gesture Elicitation Studies: What Can We Learn from 216 Studies?. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) (DIS '20). Association for Computing Machinery, New York, NY, USA, 855–872. <https://doi.org/10.1145/3357236.3395511>
- [57] Chiu-Hsuan Wang, Chia-En Tsai, Seraphina Yong, and Liwei Chan. 2020. Slice of Light: Transparent and Integrative Transition Among Realities in a Multi-HMD-User Environment. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '20). Association for Computing Machinery, New York, NY, USA, 805–817. <https://doi.org/10.1145/3379337.3415868>
- [58] Thomas Wells and Steven Houben. 2020. CollabAR - Investigating the Mediating Role of Mobile AR Interfaces on Co-Located Group Collaboration. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 414, 13 pages. <https://doi.org/10.1145/3313831.3376541>
- [59] Jacob O. Wobbrock, Meredith Ringel Morris, and Andrew D. Wilson. 2009. User-Defined Gestures for Surface Computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 1083–1092. <https://doi.org/10.1145/1518701.1518866>
- [60] Haijun Xia, Sebastian Herscher, Ken Perlin, and Daniel Wigdor. 2018. Space-time: Enabling Fluid Individual and Collaborative Editing in Virtual Reality. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 853–866. <https://doi.org/10.1145/3242587.3242597>
- [61] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 2019. Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, Article 198, 12 pages. <https://doi.org/10.1145/3290605.3300428>