



REFRAME: An Augmented Reality Storyboarding Tool for Character-Driven Analysis of Security & Privacy Concerns

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

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

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

ABSTRACT

While current augmented reality (AR) authoring tools lower the technical barrier for novice AR designers, they lack explicit guidance to consider potentially harmful aspects of AR with respect to security & privacy (S&P). To address potential threats in the earliest stages of AR design, we developed REFRAME, a digital storyboarding tool for designers with no formal training to analyze S&P threats. We accomplish this through a *frame-based authoring approach*, which captures and enhances storyboard elements that are relevant for threat modeling, and *character-driven analysis tools*, which personify S&P threats from an underlying threat model to provide simple abstractions for novice AR designers. Based on evaluations with novice AR designers and S&P experts, we find that REFRAME enables designers to analyze threats and propose mitigation techniques that experts consider good quality. We discuss how REFRAME can facilitate collaboration between designers and S&P professionals and propose extensions to REFRAME to incorporate additional threat models.

CCS CONCEPTS

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

KEYWORDS

storyboarding, threat modeling

ACM Reference Format:

Shwetha Rajaram, Franziska Roesner, and Michael Nebeling. 2023. REFRAME: An Augmented Reality Storyboarding Tool for Character-Driven Analysis of Security & Privacy Concerns. In *The 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*, October 29–November 01, 2023, San Francisco, CA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3586183.3606750>

1 INTRODUCTION

There is increased interest among interaction designers to incorporate augmented reality (AR) elements in their applications [27, 35]. However, transitioning to AR design comes with a high technical

threshold and introduces new security & privacy (S&P) threats due to how AR devices capture and process the users' environments [41, 44]. To lower the barriers to entry for novice AR designers, prior work contributed authoring techniques for creating 3D content [34, 35], implementing AR interactions [27, 28, 53], and utilizing AR-specific sensing capabilities [13, 33, 51]. However, where these tools fall short is raising designers' awareness of potentially harmful aspects of their designs. In a separate research stream, the S&P community is exploring how known threats could manifest with more severity or in new ways with always-on AR devices. For example, environmental sensing techniques and subtler form factors (e.g., AR glasses) increase privacy risks for bystanders [10, 45]; sharing harmful AR content introduces new norms for content moderation and access control in virtual spaces [24, 40, 46].

As AR systems are increasingly developed and adopted, we envision **integrating threat modeling processes within AR prototyping tools** to accelerate Privacy by Design [6]. Threat modeling is a systematic method for identifying adverse impacts of technologies before they arise, based on a set of critical harms to protect against [49]. For novice AR designers who may lack formal training in S&P, we wanted to develop simpler abstractions to ease their understanding. We took inspiration from prior research on ideation cards [2, 29], design workbooks [7, 58], and privacy comics [52], which are strong examples of educational tools that promote holistic analysis of S&P threats, but are separate from authoring workflows. We aimed to closely couple prototyping and threat modeling processes by building new analysis capabilities directly into an early-stage AR authoring tool, enabling designers to rapidly iterate on safer interactions for end-users.

To demonstrate this, we developed REFRAME, an AR storyboarding tool that adopts a **character-driven analysis approach**, personifying S&P threats to make threat modeling more accessible for novice AR designers. Our key idea to align authoring and threat analysis is a **frame-based authoring technique**: we first focus designers' attention on prototyping AR interactions at a low-fidelity and sequencing their storyboards into a core set of application states. Then, in an iterative process driven by the tool, the designer refines storyboard elements that are subject to potential threats. REFRAME implements two main interventions to highlight S&P harms from our underlying threat model: (1) automatically inserting characters in the scene to demonstrate when and where threats could occur (e.g., based on the user's location and interaction modalities); (2) enhancing storyboards with visualizations of interaction techniques and AR devices, allowing designers to simulate different environments and consider bystanders' viewing perspectives for holistic threat analysis.

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.
UIST '23, October 29–November 01, 2023, San Francisco, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0132-0/23/10...\$15.00
<https://doi.org/10.1145/3586183.3606750>

We evaluated the benefits and limitations of REFRAME through a study with eight novice AR designers, where they recreated a previous storyboard in REFRAME and used the character-driven analysis tools to identify threats and prototype mitigation techniques. The designers found value in REFRAME's combination of authoring and analysis tools not only for supporting threat modeling, but also for encouraging experimentation in AR interaction design. To further assess the quality of designers' threat analyses, we conducted an expert review of the designers' work with four S&P researchers. While the S&P experts still wanted to iterate on the novices' proposals (e.g., to implement more granular controls for users and address more complex threats), they considered the threat modeling and mitigation techniques that designers produced through REFRAME to be of good quality.

Our key contributions are (1) the REFRAME system, which implements our character-driven approach to enable implicit threat modeling directly within a low-fidelity AR authoring tool; (2) two user evaluations with novice AR designers and S&P experts which demonstrate that REFRAME can assist novice AR designers in producing good quality threat analyses and mitigation techniques.

2 RELATED WORK

Our work contributes to a stream of HCI systems research studying non-technical AR/VR designers' workflows [4, 22] and authoring techniques that make it easier to create AR/VR experiences [13, 28, 35, 51]. We also review existing methods for educating designers about S&P considerations which were inspirational to us when developing REFRAME's character-driven analysis tools: ideation cards, card games, design workbooks, and privacy comics.

2.1 AR/VR Authoring Tools

In this section, we analyze key features of current AR/VR authoring tools that lower the technical barrier to entry for creating AR experiences. We draw on Nebeling & Speicher's categorization of authoring tools according to the level of skill required from designers and the level of fidelity that can be achieved [36].

We situate REFRAME in a class of low-fidelity authoring tools that adapt traditional techniques from web and mobile design for creating 3D content and implementing AR interactions [36]. Prior work contributed capture techniques to translate paper prototypes [35, 42] and 2D digital sketches [53] into 3D content which can be viewed immersively. We were particularly inspired by tools that enable prototyping spatial relationships in AR/VR through layered 2D or 360° images: Mental Canvas [12], 360proto [34], 360theater [50], and Henrikson et al.'s cross-device storyboarding tool [19]. To enable prototyping user interactions without programming, many tools incorporate performance animation [27, 28, 57] and visual scripting techniques [60, 61]. In REFRAME, we adopted timeline-based techniques to prototype interaction techniques and changes in AR application states, similar to Montage [26] and DART [30].

To further enable non-technical designers to prototype complex interactions, existing tools provide simpler interfaces for utilizing AR/VR devices' unique input and output capabilities. Commercial toolkits (e.g., Microsoft's MRTK¹) include built-in speech recognition and hand tracking models with input simulation tools for Unity

and Unreal, enabling designers to prototype voice and gesture-based interactions even without access to an AR/VR device. DepthLab [13] contributes depth-based models for mobile AR, including physics, occlusion rendering, and light estimation techniques.

In reviewing the landscape of AR authoring tools, we identify a lack of tool support for identifying and mitigating potentially harmful aspects of interaction designs with respect to S&P. To address this gap, REFRAME extends authoring techniques from prior work to provide new tools for analyzing potential harms in AR storyboards. To establish details about the AR user's physical surroundings for analyzing context-dependent S&P threats, we created simulation environments depicting different public and private locations, inspired by Unity MARS². We adapted prior work's timeline-based authoring techniques [26, 30] to identify when and where in the usage scenario threats may occur.

2.2 Security & Privacy Education for Designers

Our work also builds on prior methods for teaching designers about S&P considerations and integrating established threat analysis frameworks (e.g., threat modeling [49], privacy impact assessments [59]) into designers' workflows. In this section, we review four examples which were inspirational to us when developing REFRAME: ideation cards [2, 29], card games [1, 11], design workbooks [7, 58], and privacy comics [52].

Ideation cards and card games. Prior S&P educational tools build on *ideation cards*, which delineate different dimensions of a problem space through text and visual examples [16]. These card decks can be used to informally learn about different S&P considerations or employed in expert-facilitated workshops to educate designers. The Security Cards [2] stratify computer security considerations into dimensions of human impact and adversaries' motivations, methods, or resources. The Privacy Ideation Cards [29] make GDPR privacy and data protection regulations more accessible for designers. To raise awareness of different threats in a more engaging format, ideation cards have been adapted into games such as Control-Alt-Hack [11] and Elevation of Privilege [1].

Design workbooks and privacy comics. Prior work in *design workbooks* [17] and *privacy comics* [62] embeds S&P considerations in creativity support tools, enabling designers to learn about threats while actively designing. Wong et al.'s workbook of science fiction-inspired surveillance technologies [58] and Chen et al.'s concepts for preserving digital data after death [7] enable designers to directly compare vignettes of technology designs and reflect on different stakeholders' values to brainstorm more ethical solutions. PrivacyToon [52] enables designers to learn about digital privacy while creating their own comics, incorporating ideation cards to guide designers in crafting a story.

While these examples contribute strong educational techniques for S&P, we wanted to integrate scaffolding techniques directly within an AR authoring tool, in order to further reduce the skill level and effort for novice AR designers to operationalize S&P design guidelines in their prototypes. Inspired by the Control-Alt-Hack game [11], we developed character-driven analysis tools which display personified representations of threats and text prompts to guide designers in brainstorming mitigation techniques.

¹<https://docs.microsoft.com/en-us/windows/mixed-reality/mrtk-unity/mrtk2/>

²<https://unity.com/products/unity-mars>

3 THREAT MODEL

While prior work on S&P education addresses S&P issues broadly, we focused our investigation on a class of threats specific to AR: **threats involving proxemic interactions in AR**. Informed by a review of S&P literature, we structured our threat model into five classes of threats that could arise based on the AR user’s physical proximity to virtual objects, physical objects, and other people. Consider other AR users or non-users who may:

- (1) **Observe the AR user** to learn private information about the AR user or the AR app [46];
- (2) **Provide inauthentic input to the AR user’s device** to make the AR app perform unwanted actions [44], e.g., issuing a voice command or gesture to quit the AR user’s app;
- (3) **Create and share unwanted or harmful virtual content with the AR user**, e.g., to disrupt the user experience or obscure critical content in the physical environment [24, 40, 46];
- (4) **Manipulate the configuration or appearance of the physical space** to disrupt AR device tracking capabilities and degrade app functionality [15, 44];
- (5) **Experience privacy concerns due to the usage of the AR app** in the same physical space, e.g., having personal information captured by the AR user’s headset sensors [3, 10, 39]

Our motivation for focusing on proxemic interactions in AR is twofold. First, leveraging proximity between people, physical, and virtual spaces is a popular strategy for designing both implicit and explicit AR interactions [5, 31], but is less common in web and mobile design. Thus, we focus on a class of threats specific to AR interfaces that novice AR designers may be unfamiliar with.

Second, prior work has studied S&P considerations for proxemic interactions in personal computing and IoT [18, 56], but these threats could manifest in new ways due to AR devices’ spatial and multimodal interactions [43, 44]. Our threat model includes a combination of (1) known attacks studied through empirical studies (e.g., shoulder surfing [14, 47], recording bystanders without their consent [10, 20]); (2) threats that the S&P community anticipates will be more severe or novel in AR (e.g., sharing harmful AR content could be more distressing due to increased immersion [24, 25]).

Our threat model does not explicitly consider traditional S&P issues, e.g., network security threats for AR [54], or cases where service providers or other AR apps are intentionally malicious towards the AR user [24]. Instead, REFRAME enables designers to analyze threats where the AR user or other people act as adversaries.

4 FORMATIVE STUDY WITH NOVICE AR DESIGNERS

To inform how a traditional threat modeling process could be translated to our system, we conducted a formative study with 13 masters students taking an introductory course on AR/VR application design. We analyzed students’ assignment submissions where they directly applied a threat model from prior work [8] to identify threats in their own AR design concepts and brainstormed design revisions to mitigate the threats. This analysis informed key system components of REFRAME: simple controls for specifying interaction techniques and AR modalities, types of simulation environments to represent popular AR usage contexts, and a range of assets for authoring mitigation techniques.

4.1 Method

We studied with masters students from a graduate-level course at our university on introductory AR/VR application design, where they previously created multiple prototypes at different fidelities for an AR interaction design project (i.e., storyboards, physical and digital prototypes). First, we developed a lecture and accompanying exercise based on traditional threat modeling [49]. A member of our research team delivered the one-hour lecture in the AR/VR course. Then, the students completed the exercise by brainstorming 2-3 threats for each category in the threat model, specifying which stakeholders were harmed or caused harm and proposing ways to mitigate the threats through redesigning their AR interface. They also selected 2 threats which could be effectively mitigated through redesigning their AR interfaces and depicted the mitigation strategy through annotating images of their existing prototypes.

Then, we invited students from the course to participate in our study by contributing their threat modeling exercise. With the 13 submissions we received, we used an affinity diagramming approach [48] to summarize themes in the assignment write-up: common AR scenarios and usage environments, how students annotated their prototypes to depict threats, and types of mitigation techniques. This analysis contributed to the design of the simulation environments and asset library which we built into REFRAME.

4.2 Results

The designers explored a variety of topics in their AR projects: science education (P2, P4, P5), nutrition (P7, P8, P13), creativity support (P1, P9, P12), art and media (P6, P10, P11), and activism (P3). 7 out of 13 participants designed for hand-held AR and the remaining 6 designed for head-worn AR. We discuss two themes around how designers annotated their prototypes to depict threats and common mitigation technique proposals.

Designers extended their prototypes with details about the physical context to highlight how threats could occur, but seldom depicted AR interactions and AR modality. P2 drew areas of a lab environment to depict how the AR device could capture sensitive information displayed on computers; P3 and P10 added obstacles to the environment to illustrate physical safety issues when AR content blocks the AR user’s view of their physical surroundings. However, we only observed limited depiction of the AR modality and specific interaction techniques (e.g., using gestures in hand-held AR), which could provide important context for brainstorming threats related to AR input and output techniques [8].

Designers relied on familiar S&P interfaces to mitigate potential threats. Our designers typically resorted to established S&P techniques such as warnings (P1, P3, P6, P9-10), user authentication (P4-5, P8-9, P11), notice and consent (P3, P8), and interfaces to inform users about S&P issues (P4, P9-10). Most designers prototyped these interfaces following WIMP-style interaction design (windows, icons, mouse, pointer), e.g., a screen-based overlay for authenticating users on hand-held AR. They also proposed more AR-specific techniques, e.g., obfuscating bystanders to preserve their privacy (P2-3, P11-12) and placing fiducial markers in the physical environment to indicate sensitive areas where AR should not be used (P2).

5 REQUIREMENTS

Guided by our formative study and review of S&P concerns for AR, we identified four requirements to integrate an implicit threat modeling process within REFRAME:

R1: Simulating different AR usage environments. Designers relied on and extended their illustrations of the physical environment to brainstorm context-dependent S&P needs [37]. In REFRAME, we enable designers to easily simulate the AR usage context through authoring storyboards on top of 2.5D environments, based on popular public and private locations from our formative study. We also facilitate switching between environments, to encourage designers to brainstorm beyond the boundaries of traditional storyboards when analyzing potential threats.

R2: Supporting rapid prototyping of AR interaction techniques and modalities. Some classes of S&P threats require a strong understanding of AR interaction techniques and differ based on AR modality (e.g., more information can be discerned from mid-air gestures for head-worn AR as compared to subtler on-screen gestures for hand-held AR). In REFRAME, we provide simple controls for specifying interaction modalities and AR devices not only to facilitate rapid prototyping, but also to encourage designers to experiment with AR-specific and multimodal interactions. REFRAME generates corresponding visualizations to make the interaction techniques more visible to designers for threat analysis.

R3: Highlighting proxemic interactions between AR users, non-users, and the physical & virtual spaces. As established in our threat model, many AR S&P threats involve proxemic interactions between people and the physical/virtual spaces (e.g., shoulder surfing attacks, passive capture of bystanders). To help novice AR designers understand these threats, REFRAME implements a *frame-based authoring* approach where designers can specify a sequence of AR application states saved with the user's location in the scene. This enables "playing through" the frames to view how the AR user's interactions overlap with physical objects and people.

R4: Providing abstractions for understanding threat models and support for brainstorming solutions. Designers in our formative study were trained through a one-hour lecture on AR S&P considerations; however, we aimed to enable them to conduct threat modeling with no formal training. To provide simpler abstractions for understanding AR-specific threats, we implemented a *character-driven analysis* approach within REFRAME that displays personified representations of threats and analysis prompts to support designers in brainstorming mitigation techniques.

6 REFRAME SYSTEM

In this section, we introduce REFRAME, an AR storyboarding tool that promotes the design of safer end-user interactions by making threat modeling more accessible to novice AR designers. We first present a system walkthrough inspired by an example from our formative study with novice AR designers (Sec. 4). Then, we introduce two system concepts at the core of REFRAME: (1) a **frame-based authoring approach** that captures and enhances storyboard elements which are relevant for threat modeling; (2) **character-driven analysis tools** that provide abstractions for understanding S&P issues by personifying classes of threats in our underlying threat model (Sec. 3).

6.1 System Walkthrough

Figure 1 shows an overview of REFRAME's user interface. The system enables two main processes: (1) authoring storyboards of AR experiences in a 2.5D environment; (2) analyzing potential S&P threats in the design concept through character-driven tools. In our system walkthrough, a novice AR designer wants to create a head-worn AR app to educate users about historical landmarks as they walk through the city. We describe how REFRAME not only allows them to rapidly create an AR interaction design, but also to consider S&P concerns as they use more features of our system.

Using REFRAME's visual editor, the designer first selects a **2.5D simulation environment** of an outdoor street to integrate contextual details about the AR user's physical surroundings (Fig. 1A). They author the scene in 2.5D on top of the simulation environment by adding **physical and virtual content from the asset library** (Fig. 1B), including AR avatars to greet new visitors and text interfaces to list facts about the city.

Next, the designer **structures the storyboard into frames**, which represent key AR application states saved with a 3D location in the simulation environment (Fig. 1C). For each frame, they specify a caption, AR interaction techniques, and the visibility of scene assets. Using REFRAME's playback controls (Fig. 1D), they can "**play through**" the frames, simulating either hand-held or head-worn AR. REFRAME generates camera animations to show the AR user's motion through physical space and visualizations to demonstrate the interaction techniques (e.g., a pointing finger for gestures, speech bubble with customizable text for voice commands).

After authoring the storyboard, the designer uses REFRAME's assisted analysis mode to help them anticipate and address potential threats. REFRAME automatically inserts **bystander and adversary characters** into the scene and shows **analysis prompts** to guide the designer in brainstorming mitigation techniques (Fig. 1E). For example, REFRAME displays the *Graffiti Spammer* who inserts AR spam content, prompting the designer to consider techniques for access control and content moderation in AR. The designer also **jumps into characters' perspectives in the 2.5D scene** to observe the AR interactions from a third-person viewpoint.

While brainstorming solutions to address the threats illustrated by the characters and prompts, the designer captures their analysis process via REFRAME's **screen and audio recording tools** (Fig. 1F). They can review their analysis at a later time, using video checkpoints to revisit when each character was analyzed. To further explore design considerations and threats that could occur in other contexts, they can switch simulation environments and visualize the AR experience using a different AR modality.

6.2 Frame-based Authoring of AR Storyboards

Our goal with REFRAME was to closely align designers' authoring and threat modeling processes, to enable them to rapidly iterate on their storyboards and prototype safer interactions for their specific AR application scenario. This presented an interesting challenge of how to structure storyboards into building blocks that are meaningful both for interaction design and threat modeling (i.e., scaffolding designers' analysis by highlighting S&P threats at various application states). We wanted to make the threat modeling process a

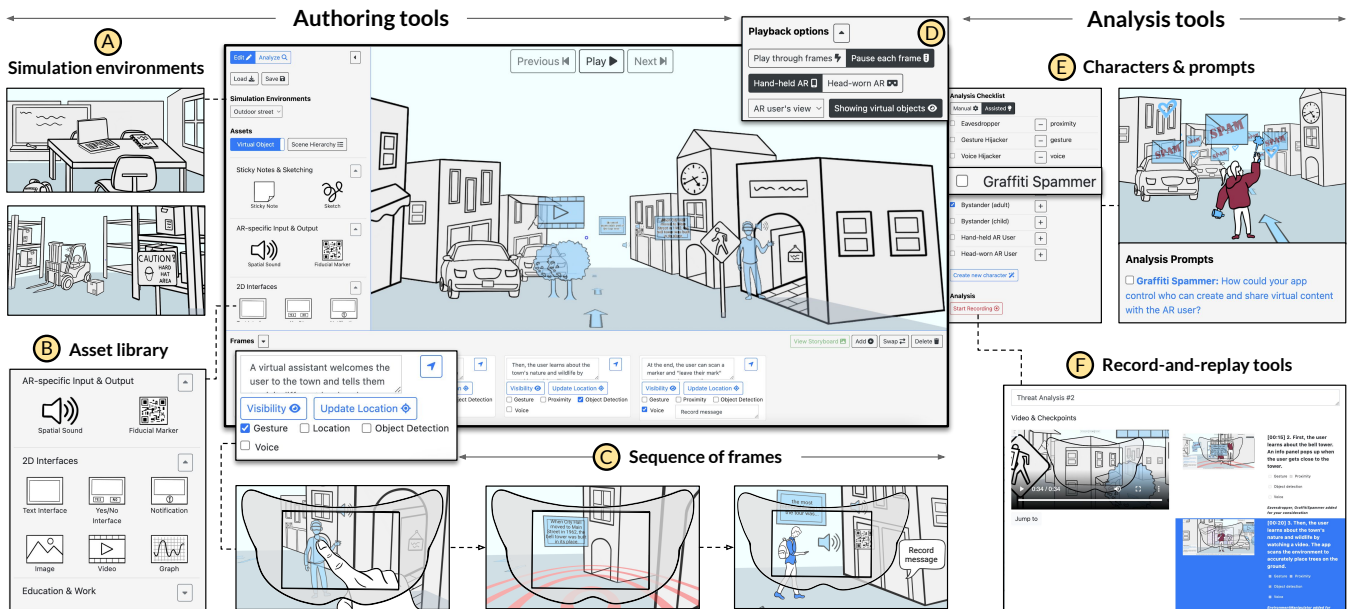


Figure 1: Overview of REFRAME system. REFRAME’s user interface consists of a scene view, where designers can both author and preview their storyboards, and side panels containing prototyping & analysis tools. To create storyboards, designers can (A) select an environment representing a public or private usage context, (B) author physical or digital content on top of the environment via the asset library, and (C) prototype AR application states through sequence of frames. Designers can preview frames through the playback interface (D); REFRAME generates camera animations to simulate states and visualizations of the interaction techniques specified for each frame. REFRAME’s character-driven analysis tools (E) support novice AR designers’ threat modeling process by inserting personified representations of S&P threats and analysis prompts to brainstorm mitigation techniques. Finally, designers can review their character-driven threat modeling via screen & audio recording tools (F).

“side effect” of storyboarding, putting it in the background while first focusing designers’ attention on prototyping interactions.

To enable this, we implemented a **frame-based authoring approach**. We extend principles from traditional 2D storyboards, which are typically composed as series of images that depict users’ key interactions with an application. Through the frames interface (Fig. 1C), we elicit the following information from designers and make it available to REFRAME in the frame’s metadata: a **text-based caption**, the **3D transform** representing the user’s location with respect to other elements in the scene, a **set of interaction techniques** available to the AR user, and the **visibility of scene elements**. Capturing this metadata serves a dual purpose: (1) forcing designers to make concrete choices by specifying the contextual situation and the employed interaction modalities; (2) providing key information for threat modeling, enabling REFRAME to confront designers with potential S&P threats (discussed in Sec. 6.3).

Next, we introduce REFRAME’s authoring features that build on the concept of frames: simulation environments, the asset library, and visualizations for AR devices and interaction techniques.

Simulating the AR usage context through different environments (R1). REFRAME allows designers to author storyboard content on top of five simulation environments (Fig. 1A). This enables easily integrating details about the AR user’s physical surroundings, which provides a basis for analyzing context-dependent

S&P issues. These environments depict a range of public and private locations (*outdoor street, home, park, classroom, and factory*) and were informed by examples from our initial study with 13 AR designers and inspired by Unity MARS³.

We implemented the environments in 2.5D using textured planes placed in 3D space (inspired by Mental Canvas [12], which simulates 3D scenes by placing 2D sketches at various depths and projection angles). This was to leverage novice AR designers’ familiarity with storyboarding in 2D [19, 35]. We designed the environments such that semantically similar elements (e.g., furniture and plants) remain at roughly the same location in each environment. This enables designers to test their storyboards in different environment without needing to significantly modify asset placement.

When the designer creates a new frame, the 3D transform of the scene camera is recorded in the frame’s metadata (Fig. 2b). This enables **“playing through” the sequence of frames** (Fig. 1C), using camera animations to simulate how the AR user moves throughout the physical environment and interacts with AR content, inspired by Mental Canvas [12]. REFRAME’s playthrough functionality further enables the character-driven tools for analyzing S&P threats (discussed later in Sec. 6.3), as it allows designers to revisit when and where potential threats could occur in the storyboard.

³<https://unity.com/products/unity-mars>

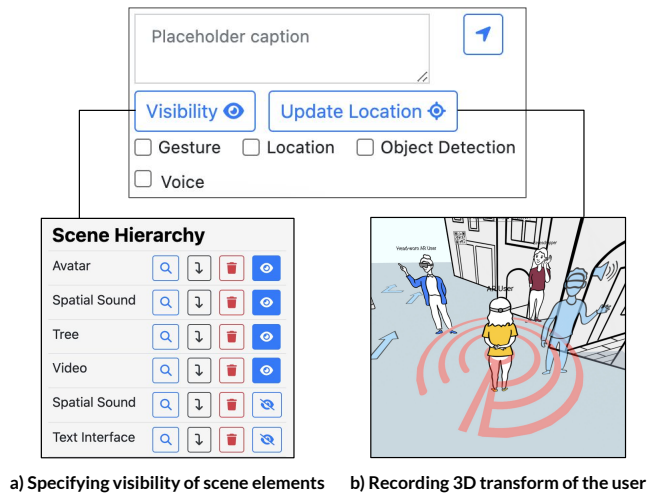


Figure 2: Frame-based authoring. REFRAME enables designers to prototype AR application states via frames, by specifying a text-based caption and set of interaction techniques available to the AR user. (a) Designers can toggle the visibility of scene elements to demonstrate changes in the physical or virtual spaces. (b) REFRAME implicitly records the 3D transform of the user for each frame, enabling “playing through” the frames to visualize how the AR user navigates through the environment while interacting with AR content.

Adding physical and virtual content through the asset library (R2). To facilitate rapid prototyping of AR experiences, REFRAME provides an asset library (Fig. 1B) informed by our formative study with novice AR designers (Sec. 4). The library includes visual representations of AR-specific input and output capabilities (e.g., fiducial markers and spatial sound), customizable text interfaces (e.g., notifications, yes/no interfaces), and categories of content related to the simulation environments (e.g., furniture, education and work). Designers can add assets to either the physical or virtual space; physical content is colored in light gray and AR content is colored in blue. REFRAME also enables sketching custom assets or placing sticky notes as placeholders for other assets.

Building on the concept of frames, REFRAME enables designers to specify the **visibility of scene elements** for each frame to illustrate the progression of AR application states (Fig. 2a), e.g., new AR content appearing when the user performs a gesture.

Specifying frame-based interactions and device simulation (R2, R3). Designers can prototype how users interact with AR content by **selecting a set of interaction techniques for each frame** (including gestures, voice commands, location-based interaction, and object detection). We provided a simple checklist interface to specify interactions (Fig. 2) for two main reasons. First, we could highlight a range of possible interaction techniques to support novice AR designers’ creativity (i.e., encouraging experimentation with AR-specific input modalities and multimodal interactions).

Second, we could use designers’ selections of interactions to enhance storyboards with **visual representations of interaction techniques and AR modality**, simulating either hand-held or

head-worn AR (Fig. 3). These visualizations are intended to help designers envision how the AR interaction space overlaps with the physical environment to brainstorm how threats from our threat model could manifest (e.g., how adversaries could learn sensitive information by observing the AR user’s interactions). Similarly, we aimed for the AR device visualizations to elicit different S&P considerations (e.g., on-screen gestures for hand-held AR are subtler than mid-air gestures for AR HMDs, which could impact bystanders’ awareness).

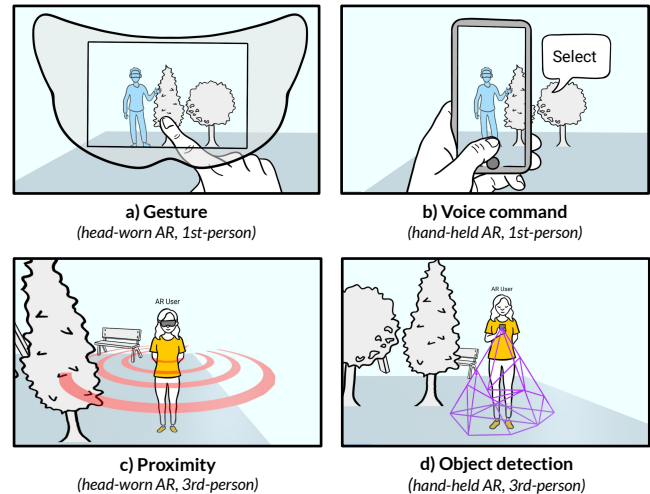


Figure 3: AR Devices and Interaction Visualizations. REFRAME generates visualizations to help designers envision how the AR interaction space overlaps with the physical environment. For the 1st-person view, we display overlays for head-worn AR (a) and hand-held AR (b). When viewing from a 3rd-person perspective, we show an avatar of the AR user wearing a headset (c) or holding a phone (d). We illustrate interactions with a pointing hand for gestures (a), speech bubble with customizable text for voice commands (b), radar visualization for proxemic interaction (c), and mesh visualization for object detection (d).

6.3 Character-Driven Analysis of S&P Threats

A key challenge with integrating our threat model (Sec. 3) into REFRAME was providing simpler abstractions for novice AR designers to understand technical concepts and making subtler threats more visible. Inspired by prior work in S&P ideation cards and games [2, 11], we adopted a **character-driven analysis approach** to provide concrete, personified representations of each threat and depict adversaries’ abilities in a consistent way. We created nine characters (Fig. 4) that address the five classes of our threat model: adversaries (*Eavesdropper*, *Voice & Gesture Hijackers*, *Graffiti Spammer*, and *Environment Manipulator*) and bystanders (*adult*, *child*, *hand-held & head-worn AR users*).

Our character illustrations provide visual cues for novice AR designers to understand S&P threats; this personification is not meant to be interpreted literally, but rather to provide a concrete and provocative explanation of the threat. In designing these characters,


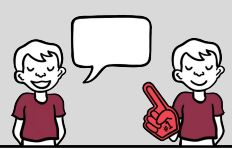


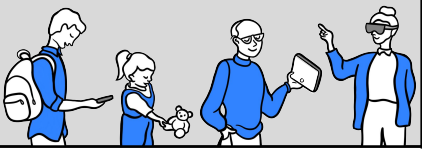
| | Eavesdropper | Voice & Gesture Hijackers | Graffiti Spammer | Environment Manipulator | Bystanders (adult, child, hand-held & headworn AR users) |
|-----------------|-----------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------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

Figure 4: Characters and Analysis Prompts. REFRAME implements nine adversary and bystander characters depicting various S&P concerns based on our threat model for proxemic interactions in AR (Sec. 3). In *assisted* analysis mode, adversary characters (shown in red) are automatically added to the scene and displayed during frames when the AR user is using a particular interaction technique (e.g., the *Environment Manipulator* is linked to object detection). Bystander characters (shown in blue) are persistent in the scene to encourage designers to consider risks to bystander privacy throughout the entire AR usage scenario. REFRAME displays text prompts for each character to summarize the threat and hint towards a potential solution.

we referred to prior work studying laypersons’ visual representations and mental models of S&P harms [38]. For example, we slightly exaggerated adversary characters’ features (e.g., eyes, hands) and incorporated familiar objects as metaphors for explaining threats (e.g., graffiti to represent unwanted virtual content).

We also created **analysis prompts corresponding to each character** that hint towards design solutions to mitigate the S&P threats (Fig. 4), e.g., designers are encouraged to think about access control and content moderation for the *Graffiti Spammer*.

Next, we detail REFRAME’s built-in threat modeling features that incorporate the characters and extend the frame-based authoring approach: threat analysis modes and recording tools.

Analyzing S&P threats through automatically-inserted characters; answering prompts hinting at mitigation strategies; switching into characters’ perspectives (R4). To operationalize the character-driven threat analysis approach in REFRAME, we developed two **analysis modes**. In *manual* mode, the designer decides when and where to add characters for analysis. In *assisted* mode, characters are automatically added to the scene and displayed based on the interaction techniques recorded in the current frame’s metadata (e.g., *Voice Hijacker* is shown for voice commands). REFRAME also displays corresponding analysis prompts to support designers in brainstorming mitigation techniques (Fig. 1D). Designers can indicate which characters and prompts they addressed in their analysis through the checklist interface. REFRAME also enables **creating custom characters** by specifying a name and choosing an avatar representation from our existing character set.

Designers can play through frames in 1st-person perspective from the AR user’s viewpoint or in **3rd-person perspective from other characters’ viewpoints** (Fig. 5), which we implemented by creating new cameras at each characters’ location in the scene. This allows designers to gain a broader view of the interaction area and brainstorm how the AR user’s interactions may be perceived or exploited by others who are co-located in the same space.




Figure 5: Character Perspectives. Designers can analyze AR interactions and related S&P threats (a) from the first-person perspective of the AR user or (b) from the third-person perspective of any bystander / adversary characters in the scene.

Capturing and reviewing the analysis process with screen and audio recording tools (R4). REFRAME offers record-and-replay tools to enable designers to review threats and mitigation strategies that they brainstormed during the analysis process (Fig. 1E). When designers start a recording session, REFRAME generates **video checkpoints for each frame**: (1) when each character is added to the scene (either manually or automatically), and (2) when prompts are checked off by the designer to indicate a completed analysis of that character. These checkpoints enable designers to jump to specific timestamps when replaying the video. REFRAME also supports generating a static storyboard with snapshots of each frame to capture the digital, 2.5D scene in a traditional storyboard format.

6.4 Implementation

We implemented REFRAME as a web-based editor using A-Frame⁴, a declarative framework for developing WebXR experiences on top of THREE.js (a WebGL library for rendering 3D content in

⁴<https://aframe.io/>

the browser). We used open-source A-Frame components for animation, wasd-controls and look-controls for navigation in the scene, and super-hands for dragging and dropping entities. We utilized the HTML5 Canvas API⁵ for sketching, WebMedia API⁶ for the analysis recording tools and Google Firebase⁷ for storing and loading storyboards and analysis recordings. More information on REFRAME and the source code is available at <https://www.mi2lab.com/research/reframe/>.

7 USER STUDY WITH NOVICE AR DESIGNERS

Figure 6 shows an overview of our two summative system evaluations. The first was a **user study with novice AR designers** who previously created a storyboard of an AR experience. Our goals were to investigate how designers' AR design concepts would evolve when using REFRAME to recreate and analyze their storyboards, and understand benefits and limitations of the system for threat modeling and prototyping mitigation techniques. The output of this user study (i.e., the designers' storyboards and analysis recordings) were used as input for our second study with S&P experts, who assessed the quality of the novice AR designers' threat modeling and mitigation techniques as enabled by REFRAME (Sec. 8).

7.1 Method

Using university mailing lists, we recruited current or recently graduated masters students who had previously created a storyboard for an AR interaction design. We selected eight designers (5 female, 3 male, average age = 26.4 years) who reported having one year or less experience designing AR interfaces. They also had limited training on S&P topics: 7 designers participated in one-time lectures or training sessions for internships (P1-5, P7-8) and 1 of these designers took a course in privacy in information technology (P7).

The designers participated in three tasks: (1) recreating their previous storyboard using REFRAME, (2) identifying potential threats in their design concepts using REFRAME's character-driven analysis tools, and (3) revising their AR storyboards to prototype mitigation techniques for the two threats. Each individual study session was conducted in 1 hr remotely over Zoom. Participants were compensated with a \$25 USD giftcard for their time.

Task 1: Recreating an AR storyboard using REFRAME's editing tools. A member of our research team introduced REFRAME's authoring features through a training video and guided the designer to recreate an existing storyboard which they brought to the study. We instructed the designers to think aloud to describe their prototyping process (i.e., selecting simulation environments, adding assets to the scene, and specifying frames with interaction techniques). After the task, designers discussed which parts of the task were easy and challenging to complete. We saved the storyboard for use in our following study with S&P experts (Sec. 8).

Task 2: Implicit threat modeling using REFRAME's analysis tools. Next, we introduced REFRAME's analysis features and guided the designers in playing through their storyboards in *assisted* mode, where characters are automatically inserted into the scene to encourage consideration of S&P threats from our threat

model (Sec. 3), guided by analysis prompts. We instructed the designers to think aloud to analyze threats which could arise with at least one bystander and one adversary character. Designers also created one custom character representing another individual who may be present in the usage scenario and jumped into characters' perspectives to experience the storyboard from their point of view. We collected screen and audio recordings of the analysis sessions using REFRAME's record-and-replay tools.

We did not compare designers' use of *manual* and *assisted* analysis mode in our study; this as an avenue for future work.

Task 3: Revising the storyboard to prototype mitigation techniques. In the final task, the designers described two threats which they felt were the most important to mitigate for their specific use case. We encouraged them to revisit their analysis using the recording checkpoints generated in the previous task. Then, they revised their storyboard to prototype mitigation techniques for the two threats they selected, using any of REFRAME's editing features (i.e. the asset library, sticky note and sketching features). We again recorded the prototyping session and saved the revised scene, for use in our follow-up expert evaluation.

Debrief. We concluded the study with a discussion around how the participants' design concepts evolved through using REFRAME and which system features, if any, were the most useful for brainstorming and mitigating threats, respectively. We also discussed pros and cons of using REFRAME for storyboarding AR experiences.

We used an affinity diagramming approach [48] to analyze themes in the designers' prototyping processes and post-task discussions. One author extracted observations for each designer by reviewing the video transcripts, then all three authors summarized and voted on themes across all participants. We categorized the types of mitigation techniques proposed for each character and summarized common prototyping strategies (e.g., introducing multimodal interactions or additional application states).

7.2 Results

All designers were able to recreate their original storyboards using REFRAME within the given timeframe of 30 minutes. The Appendix shows an overview of the designers' AR project topics, choice of AR modality and simulation environments. They explored a variety of interaction design topics: educational apps (D1-2), games and entertainment (D3-4, D8), a navigation system for lunar missions (D5), and a tutorial system for physical tasks (D7). Five designers designed for hand-held AR (D1, D3-4, D6, D8) and three designed for head-worn AR (D2, D5, D7).

In the rest of this section, we discuss five themes related to the designers' experience of using REFRAME to author storyboards, analyze potential threats through the character-driven analysis tools, and prototype mitigation techniques.

Supporting efficiency and experimentation in storytelling. From an authoring standpoint, designers reported a main benefit of REFRAME is enabling "fast and consistent storytelling" (D2) through the interaction checklist and device simulation features which simplify repetitive actions from traditional storyboarding, e.g., hand-drawing the AR device (D1-3, D7-8). D1 expressed that there is "a lot more freedom to experiment with your storyboard,

⁵https://developer.mozilla.org/en-US/docs/Web/API/Canvas_API

⁶<https://developer.mozilla.org/en-US/docs/Web/Media>

⁷<https://firebase.google.com/>

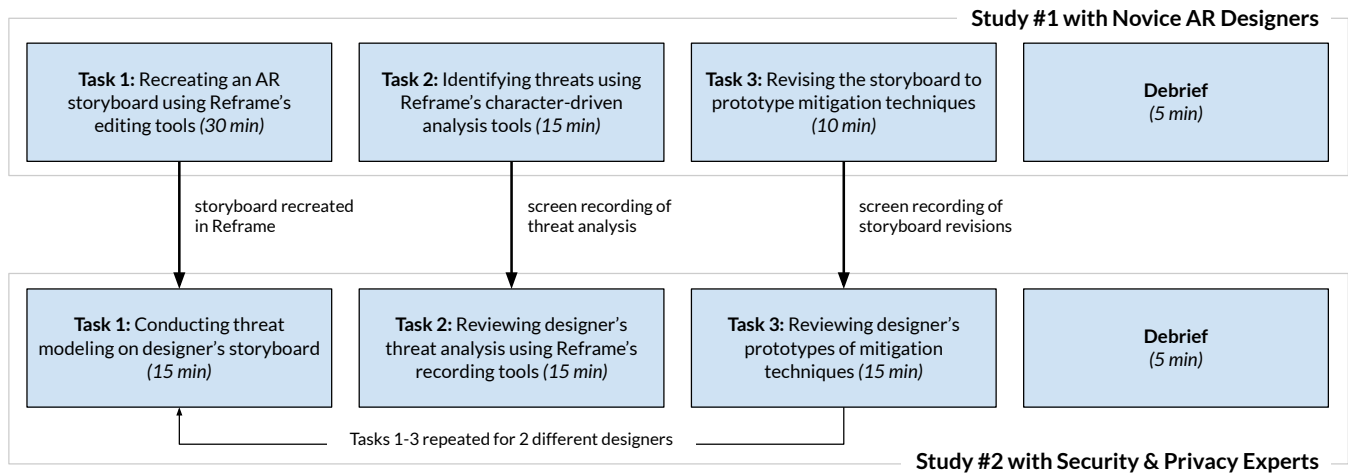


Figure 6: User studies with novice AR designers and S&P experts. In our first user study with 8 novice AR designers, we explored how their design concepts evolved through using REFRAME’s authoring tools to recreate a previous storyboard and analyzing their storyboards via the character-driven analysis tools. They also prototyped mitigation techniques for 2 threats that they considered the most important to address. Using the designers’ output as input to our study with 4 S&P experts, we investigated experts’ perception of the quality of the designers’ REFRAME-supported analyses and mitigation techniques.

and you don’t feel quite as locked in” compared to traditional storyboarding. “I was more willing to play with different options since I didn’t have to draw everything.” (D1).

Usability challenges compared to traditional storyboarding. Designers perceived the main drawbacks of REFRAME to be usability related, as they sometimes struggled to scale and place content accurately in the environment (D1-2, D4, D8); they suggested extending the system’s authoring features to allow for more precise transform editing similar to commercial tools like Unity. A few designers felt that digital tools like REFRAME could have a higher barrier to entry than traditional storyboarding (D3-5), “whereas anybody can use a pen and create a sketch” (D3).

Simulation environments enhanced threat analysis by demonstrating how the AR user interacts in the space. 7 out of 8 designers reported that their interaction design and consideration of threats improved as a result of prototyping in a 2.5D simulation environment as opposed to a static, 2D storyboard (D2-8). D2 expressed that ‘telling the story in a space’ helped them to “think of more ideas about privacy & security and how we can interact with the space.” In particular, we observed participants using the simulation environments and character perspectives to analyze threats to bystander privacy. When viewing the hand-held AR user from a third-person perspective, D1 commented that “it doesn’t look like any sort of data is being collected” and that bystanders might be more aware if the user was wearing a headset.

Characters and prompts provide effective scaffolding for brainstorming threats, but could pose challenges for brand-new design concepts. All designers reported that the characters and analysis prompts were valuable for brainstorming threats which are difficult to recognize when storyboarding with pen and paper. D6 expressed that the *assisted* analysis mode “forces [them] to think even beyond what’s given” in the storyboard to determine “what could distract the system from focusing on the user.” As D3 stated,

“characters that were behaving themselves, doing unexpected things really replicates the fact that when you’re using AR in the real world, there are things out of your control.”

D6 expressed that REFRAME is “definitely a good system for fleshed out ideas” where they had previously thought about different steps in the usage scenario. However, they anticipated challenges “coming in with a fresh idea” as it could be overwhelming to manage the threat analysis and interaction design simultaneously.

Designers’ mitigation techniques introduced multimodal interactions or additional application states. Figure 7 shows four examples of mitigation techniques which the designers prototyped to address the most important threats from their character-driven analysis. All designers prototyped new application states to mitigate S&P threats: startup procedures to authenticate the AR user (D6-7) or specify tracked areas of the physical environment (D2-3), notifying the user when bystanders or potential adversaries are nearby (D1, D5, D8). A common strategy was making input or output techniques increasingly multimodal to address S&P goals: D4 combined gesture and proximity to prevent gesture hijacking (Fig. 7c); D3 extended text instructions with spatial audio to notify bystanders about environmental sensing (Fig 7b).

8 EXPERT REVIEW OF DESIGNERS’ MITIGATION TECHNIQUES

To assess the quality of threat modeling and mitigation techniques enabled by REFRAME, we conducted a second study with security & privacy experts (Figure 6). In this study, we asked the experts to review the novice AR designers’ storyboards created with REFRAME; we aimed to understand the benefits and limitations of using our character-driven approach for analysis and our authoring tools to implement mitigation techniques from the experts’ perspectives.

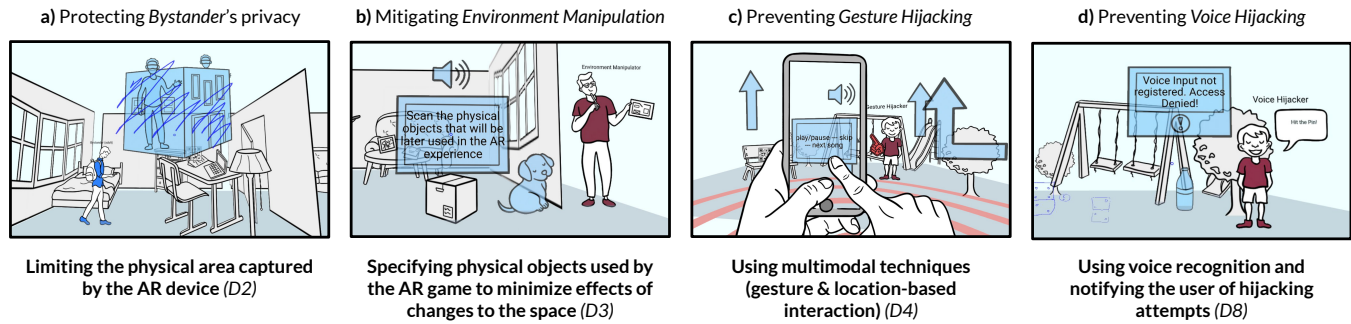


Figure 7: Examples of mitigation techniques. (a) D2 proposed allowing the AR user to specify areas of the physical environment where the AR device is allowed to capture, to minimize information captured about bystanders. (b) To prevent object detection from malfunctioning if an adversary introduces new objects in the physical environment, D3 proposed specifying objects which the AR app will track in the game. (c) D4 introduced a multimodal technique combining gestures with location-based interaction to prevent gesture hijacking. (d) To prevent voice hijacking, D8 envisioned using voice recognition and prototyped a notification interface shown to the user when an unauthorized user's voice is detected.

8.1 Method

We identified and emailed researchers active in S&P communities within the CHI, SOUPS, and USENIX conferences who had at least one first-authored publication or at least two years of PhD-level research on related topics. Four experts participated in our study (1 female, 3 male, average age of 27.5 years). Three out of four experts also had 1-2 years of AR development experience for both hand-held and head-worn AR (experts E2–E4). E1 had previously used AR apps and conducted research on VR authoring tools, but had not developed AR apps before.

We randomly assigned each expert the work of two AR designers to review. The study consisted of three tasks, which we repeated for both designers: (1) conducting an explicit threat modeling exercise based on the designers' REFRAME storyboards, (2) reviewing the quality of designers' implicit threat modeling supported by REFRAME's analysis tools, and (3) reviewing the quality of designers' mitigation techniques. We conducted 1.5 hr study sessions remotely over Zoom, compensating participants with a \$50 USD giftcard.

Task 1: Performing threat modeling on designers' REFRAME storyboards. Using REFRAME, the S&P experts first reviewed the storyboards from Task 1 of our AR designer study and summarized the key interactions. To establish a baseline threat modeling of the designers' storyboards, the experts brainstormed 1-2 threats for each class of our threat model. This enabled us to compare the experts' threat modeling to the novice AR designers' implicit analysis using REFRAME's character-driven analysis tools.

Task 2: Reviewing designers' threat analysis. The goal of the second task was to elicit the S&P expert's quality assessment of the designers' threat modeling. First, the experts watched a video to get an overview of REFRAME's character-driven analysis tools. Then, using REFRAME's record-and-replay tools, they reviewed the screen and audio recordings of the designers' analysis of two threats from the previous study. For each threat, the experts rated the quality of the analysis on a scale from 1-5 (1 = Poor, 5 = Excellent) and discussed the rationale for the rating. We did not impose a definition for the quality of threat modeling; instead, our goal was to elicit quality metrics from the experts.

Task 3: Reviewing designers' mitigation techniques. Following the same procedure as Task 2, the experts reviewed recordings of the designers' process of prototyping mitigation techniques (for the same two threats discussed in the previous task). They rated the quality of the mitigation techniques on a scale from 1-5 (1 = Poor, 5 = Excellent) and explained their rationale.

Debrief. After completing Tasks 1-3 for two different designers, we concluded with a semi-structured interview portion to discuss the pros and cons of using REFRAME to brainstorm threats and prototype mitigation techniques.

We used the same affinity diagramming analysis approach [48] as in the previous study (Sec 7.1). We summarized themes in the experts' threat modeling as compared to the designers' analysis using REFRAME (Task 1), the experts' rationale for assigning specific quality ratings (Tasks 2 & 3) and the pros and cons of REFRAME.

8.2 Results

We identified three themes around the experts' perception of the quality of designers' threat analysis and mitigation strategies. We also discuss the benefits and limitations of REFRAME that the experts identified for supporting threat modeling.

Experts found novice designer's threat modeling with REFRAME to be of good quality. The experts rated the quality of the designers' threat modeling favorably (average of 3.91/5, range of 2-5) based on the similarity to their own analysis (i.e., assigning a higher score if the designer suggested a threat which the expert did not consider and vice versa). They assigned higher ratings when designers "fully thought about the type of information that can be captured" (E1) and brainstormed holistically about how the AR app is perceived from different stakeholders' viewpoints (E1, E4). REFRAME supports designers in this regard by providing "interactive checkboxes which list out all the factors to consider" for interaction modalities and enabling designers to view the scene from different characters' perspectives (E1).

Characters and environments are valuable for concretizing threat modeling, but require making assumptions about real-world scenarios. Overall, the experts reported that REFRAME's

character-driven approach is promising for supporting novice AR designers in identifying harms from our threat model. E1 expressed that “when you’re so focused on building,” potential threats can “fall under the cracks” and it can be “so easy to miss these threats even for people who consider themselves experts.” Similarly to the AR designers, E3 suggested that the 2.5D environments and frame-based authoring enhanced the character-driven analysis: “you have a timeframe and you have a physical path... this provides more detail than just brainstorming with sticky notes.”

However, some experts anticipated that with lower-fidelity depictions of the physical environment and AR interactions, some “security needs might be under assumption” (E4), e.g., assuming all adversaries are located within the scene. E3 and E4 suggested that analyzing threats with respect to more realistic representations of real-world use cases may elicit different S&P properties to protect.

Designers proposed base mitigation techniques of good quality; S&P experts considered more ways a threat can manifest. The experts also rated the designers’ mitigation techniques as good quality (average of 3.97/5, range of 3-5). They favored mitigation techniques that provided granular controls with “different degrees of restriction” to preserve privacy (E1), e.g., multiple levels of access control for viewing, editing, and sharing AR content (D4). Consequently, the experts assigned lower ratings for techniques that provided only “black and white” controls (E2). For example, E2 critiqued D2’s proposal of muting the AR user’s microphone when bystanders are nearby, since this privacy-preserving measure completely limits the use of voice-based interaction.

All experts wanted the designers to push mitigation techniques further to cover complex variations of the threats, e.g., considering the next generation of AR devices, multiple types of harmful input and output (E1-4). For example, E4 argued “there are a lot of different ways one could put ‘graffiti’ in your virtual world” beyond adding new AR spam content e.g., manipulating the scale of existing virtual content to block critical real-world objects [24] or tricking users into thinking virtual objects are physical, assuming that future AR devices will have hyper-realistic rendering capabilities.

9 DISCUSSION

Our studies around REFRAME were encouraging, as they demonstrated that AR designers with no formal training in S&P can produce good quality threat analyses and mitigation techniques using the system, validated by the S&P experts. Reflecting on this research, we identified three opportunities to extend our approach: (1) facilitating co-design between designers and S&P experts; (2) bringing character-driven threat modeling closer to real-world use cases; and (3) supporting additional threat models.

9.1 Facilitating co-design between novice AR designers and S&P experts

A major goal in designing REFRAME was enabling novice AR designers to analyze potential S&P threats on their own. It was encouraging that the novice designers often had the right intuition when analyzing threats and prototyped good quality mitigation techniques, in the opinion of the S&P experts.

However, our studies demonstrate that there is still a strong need for S&P experts in the design pipeline to iterate on designers’

threat analysis and mitigation technique proposals. Reflecting on the role of S&P experts in AR designers’ workflows, E2 expressed that REFRAME “is already, to a specific level, good enough to give designers insights on what could go wrong.” However, we have not yet reached the “point we can consider what the designers design with authoring tools are secure & privacy aware.” Where all experts saw the most potential for improvement is providing granular, rather than binary, S&P controls and further addressing complex threats (e.g., considering multiple classes of threats in combination, different types of inauthentic input and output).

We find it promising that the experts saw value in REFRAME for facilitating their collaboration with designers. E2 expressed that REFRAME’s record-and-replay tools are valuable for asynchronously iterating on designers’ initial mitigation techniques, stating that “it’s amazing to be able to think about threats in such a collaborative way.” To broaden REFRAME’s support for collaboration, we could extend our recording facilities to entire prototyping sessions and enable custom recording checkpoints for collaborators to draw each other’s attention to specific design aspects.

9.2 Aligning character-driven threat modeling with real-world use cases

We designed REFRAME as an early-stage prototyping tool with built-in threat analysis tools, in order to enable rapid and flexible exploration of interaction designs and S&P threats in different contexts. While the experts were positive about the quality of designers’ threat analyses based on the low-fidelity storyboards, they anticipated that more realistic depictions of usage environments may yield more authentic threats. For example, our current implementation uses 2.5D environments and does not accurately model proxemic interactions between people and physical & virtual content. Studying threats related to fine-grained proxemic interactions (e.g., sensing F-Formations of AR users and bystanders to infer their social relationships [5]) may require transitioning to higher-fidelity tools that more precisely simulate AR devices’ sensor ranges, users’ position and orientation (e.g., Unity MARS).

Conversely, it is possible for REFRAME to depict threats that are not probable to occur in real-world usage environments (e.g., participants perceived threats of *Environment Manipulation* to be less realistic, as they require adversaries to have advanced knowledge of AR tracking capabilities). This actually enables an important step in traditional threat modeling [49] of ranking the severity of all threats (including both realistic and improbable threats) to prioritize which mitigation techniques to implement. As such, we encourage future work extending REFRAME’s approach to continue supporting more unconventional threats while extending support for other threat models (discussed further in Sec 9.3).

To enable brainstorming threats more aligned with authentic usage environments, we can imagine extending REFRAME to support immersive authoring and analysis, similar to prior tools including Pronto [28] and RealitySketch [53]. Future work should explore how to balance character illustrations to match the realism of the environment without losing the room for interpretation which our current implementation affords. We also note that an immersive authoring mode that requires access to the designer’s environment raises its own S&P risks that future work should address.

9.3 Extending support for other threat models

REFRAME's current character-driven analysis tools are designed to personify our threat model involving proxemic interactions in AR (Sec. 3). Our assisted analysis mode places a focus on adversaries who are co-located with the AR user and attack synchronously during the usage scenario. As such, the S&P concerns and mitigation techniques elicited through REFRAME may be limited in their generalizability to other threat models. In this section, we discuss three examples of alternate threat models and propose how REFRAME could be extended to support them.

Threats carried out remotely or asynchronously. Considering different dimensions of the time/space matrix [21], prior work has also studied S&P threats where adversaries operate remotely (e.g., remote collaborators learning private information about the AR user's surrounding environment [46]) and asynchronously compared to when the AR usage scenario occurs (e.g., adversaries obtaining access to application data after the user ends the AR session [8]). REFRAME could be extended to support multiple side-by-side environments and illustrate data flows between remote users.

Threats involving service providers or other AR applications. Our threat model does not explicitly consider threats where the adversary is a service provider (e.g., app developer inferring personal information from app data to curate advertisements [32]) or other AR apps running in the background (e.g., another AR app inserting unwanted content in the active app's view [23, 25]). These threats may require a more technical understanding of the broader ecosystem of AR applications, such as network security and operating system permissions across multiple applications. However, future work should investigate how to best represent these concepts to designers without narrowing or oversimplifying the technical concepts and threats (e.g., how to illustrate a cloud layer).

Threats involving future AR devices. We based REFRAME's authoring tools and threat model on capabilities of current AR devices with set of common interaction techniques (i.e., gesture, voice, location-based interaction, and object detection). To an extent, REFRAME already supports consideration of new sensing capabilities and corresponding threats by specifying new characters, analysis prompts, and visualizations for AR modalities and interaction techniques (cf. Fig. 3). This may require additional exploration on how to effectively represent interaction techniques which are not performed "in-front" of the AR device. For example, exploring threats to physical safety with haptic redirection [55]) may need a full-body representation of the AR user in first-person viewing mode.

9.4 Limitations

We discuss limitations of our work with regards to (1) generalizability of our study insights to different populations of designers and S&P experts; (2) how REFRAME's threat analysis tools compare and could be applied to existing AR authoring tools.

Limitations with the study sample. We studied with masters students who had 1 year or less of AR design experience to control for expertise. While the opinions and challenges they expressed may not be generalizable outside of an academic context, there is agreement between our studies and prior work studying with professional designers in terms of design challenges [4, 22]. Despite studying with only 4 S&P experts, we are encouraged by the fact

that they agreed on the benefits our system may provide to novice AR designers and that it can enable novice AR designers to produce threat analyses of good quality. We reduced potential participant response bias [9] by having the S&P experts review the novice's usage of the tool against a baseline of their own threat modeling and we explicitly asked both groups to discuss limitations of REFRAME.

Generalizability and effectiveness of REFRAME's threat analysis features. We developed REFRAME as new AR authoring tool in order to study how a combination of authoring and threat analysis features could promote S&P-minded design. We were concerned that participants might have varying experience with or certain opinions of the existing tool. This could pose threats to validity when trying to isolate the effects of REFRAME's threat analysis tools from the existing authoring environment. That said, we see potential for REFRAME's approach to be adapted in existing tools since some key authoring features are shared (e.g., MRTK provides a set of profiles to specify interactions, Unity MARS supports authoring on top of simulation environments).

The lack of a baseline condition is a common limitation of AR tools research, including our work on REFRAME. We considered conducting a comparative evaluation with prior S&P educational methods (e.g., ideation cards [2, 29] or design workbooks [7, 58]). However, given key differences in the threat models (AR-specific threats in REFRAME vs. S&P issues broadly) and modality (character-driven tools embedded within a digital prototyping system vs. physical instructional materials), this was not a fair comparison. Ultimately, we worked with PhD-level S&P researchers to represent the state-of-the-art, asking them to conduct threat modeling on novice AR designers' storyboards without REFRAME, then compare their work with the designers' character-driven analysis. While this is not a perfect baseline, we find it encouraging that in the S&P experts' opinions, the designers produced good-quality threat modeling and mitigation techniques, without having formal training or introduction to the underlying threat model.

10 CONCLUSION

This paper presented REFRAME, an AR storyboarding tool which enables designers with limited background in security & privacy to analyze potential threats in their design concepts through an implicit, character-driven threat modeling approach. Through studies with novice AR designers and S&P experts, we demonstrated that REFRAME can assist designers in producing good quality threat analyses and mitigation techniques. Designers' feedback was encouraging in that they found value in REFRAME's authoring and analysis tools, not only to improve their storyboards with respect to S&P, but also to encourage experimentation with a range of interaction techniques. The experts also saw potential in REFRAME's record-and-replay tools to facilitate collaboration with AR designers to iterate upon mitigation techniques. Future work could continue to explore how to effectively incorporate consideration of different threat models in the AR prototyping process, enable more realistic threat analysis through immersive authoring, and explore techniques to enhance more standardized AR authoring tools (e.g., Unity and Unreal) with S&P guidance. As a final thought, we are excited about incorporating tools like REFRAME into residential and online curricula teaching AR design with S&P in mind.

REFERENCES

- [1] 2010. Elevation of Privilege (EOP) Threat Modeling Card Game. <https://agilestationery.com/products/elevation-of-privilege-game>
- [2] 2013. The Security Cards: A Security Threat Brainstorming Kit. <https://securitycards.cs.washington.edu/>
- [3] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security, SOUPS 2018, Baltimore, MD, USA, August 12–14, 2018*. USENIX Association, 427–442. <https://www.usenix.org/conference/soups2018/presentation/adams>
- [4] Narges Ashtari, Andrea Bunt, Joanna McGrenere, Michael Nebeling, and Parmit K. Chilana. 2020. Creating Augmented and Virtual Reality Applications: Current Practices, Challenges, and Opportunities. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25–30, 2020*. ACM, 1–13. <https://doi.org/10.1145/3313831.3376722>
- [5] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. 2010. Proxemic interaction: designing for a proximity and orientation-aware environment. In *ACM International Conference on Interactive Tabletops and Surfaces, ITS 2010, Saarbrücken, Germany, November 7–10, 2010*. ACM, 121–130. <https://doi.org/10.1145/1936652.1936676>
- [6] Ann Cavoukian. 2010. Privacy by Design: The 7 Foundational Principles. Revised: October 2010.
- [7] Janet X. Chen, Francesco Vitale, and Joanna McGrenere. 2021. What Happens After Death? Using a Design Workbook to Understand User Expectations for Preparing their Data. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8–13, 2021*. ACM, 169:1–169:13. <https://doi.org/10.1145/3411764.3445359>
- [8] 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>
- [9] Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. 2012. "Yours is better!": participant response bias in HCI. In *CHI Conference on Human Factors in Computing Systems, CHI '12, Austin, TX, USA - May 05 - 10, 2012*. ACM, 1321–1330. <https://doi.org/10.1145/2207676.2208589>
- [10] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies. In *CHI Conference on Human Factors in Computing Systems, CHI '14, Toronto, ON, Canada - April 26 - May 01, 2014*. ACM, 2377–2386. <https://doi.org/10.1145/2556288.2557352>
- [11] Tamara Denning, Adam Lerner, Adam Shostack, and Tadayoshi Kohno. 2013. Control-Alt-Hack: The Design and Evaluation of a Card Game for Computer Security Awareness and Education. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (Berlin, Germany) (CCS '13)*. Association for Computing Machinery, New York, NY, USA, 915–928. <https://doi.org/10.1145/2508859.2516753>
- [12] Julie Dorsey, Songhua Xu, Gabe Smedresman, Holly E. Rushmeier, and Leonard McMillan. 2007. The Mental Canvas: A Tool for Conceptual Architectural Design and Analysis. In *15th Pacific Conference on Computer Graphics and Applications, PG 2007, Maui, HI, USA, October 29 - November 2, 2007*, Marc Alexa, Steven J. Gortler, and Tao Ju (Eds.). IEEE Computer Society, 201–210. <https://doi.org/10.1109/PG.2007.64>
- [13] Ruofei Du, Eric Turner, Maksym Dzitsiuk, Luca Prasso, Ivo Duarte, Jason Dourgarian, João Afonso, Jose Pascoal, Josh Gladstone, Nuno Cruces, Shahram Izadi, Adarsh Kowdle, Konstantine Tsotsos, and David Kim. 2020. DepthLab: Real-time 3D Interaction with Depth Maps for Mobile Augmented Reality. In *UIST '20: The 33rd Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 20–23, 2020*. ACM, 829–843. <https://doi.org/10.1145/3379337.3415881>
- [14] Malin Eiband, Mohamed Khamis, Emanuel von Zezschwitz, 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, CO, USA, May 06–11, 2017*. ACM, 4254–4265. <https://doi.org/10.1145/3025453.3025636>
- [15] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. 2018. Robust Physical-World Attacks on Deep Learning Visual Classification. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. Computer Vision Foundation / IEEE Computer Society, 1625–1634. <https://doi.org/10.1109/CVPR.2018.00175>
- [16] Batya Friedman and David G. Hendry. 2012. The envisioning cards: a toolkit for catalyzing humanistic and technical imaginations. In *CHI Conference on Human Factors in Computing Systems, CHI '12, Austin, TX, USA - May 05 - 10, 2012*. ACM, 1145–1148. <https://doi.org/10.1145/2207676.2208562>
- [17] William W. Gaver. 2011. Making spaces: how design workbooks work. In *Proceedings of the International Conference on Human Factors in Computing Systems, CHI 2011, Vancouver, BC, Canada, May 7–12, 2011*, Desney S. Tan, Saleema Amer-shi, Bo Begole, Wendy A. Kellogg, and Manas Tungare (Eds.). ACM, 1551–1560. <https://doi.org/10.1145/1978942.1979169>
- [18] Saul Greenberg, Sebastian Boring, Jo Vermeulen, and Jakob Dostal. 2014. Dark patterns in proxemic interactions: a critical perspective. In *Designing Interactive Systems Conference 2014, DIS '14, Vancouver, BC, Canada, June 21–25, 2014*. ACM, 523–532. <https://doi.org/10.1145/2598510.2598541>
- [19] Rorik Henrikson, Bruno Rodrigues De Araújo, Fanny Chevalier, Karan Singh, and Ravin Balakrishnan. 2016. Multi-Device Storyboards for Cinematic Narratives in VR. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology, UIST 2016, Tokyo, Japan, October 16–19, 2016*. ACM, 787–796. <https://doi.org/10.1145/2984511.2984539>
- [20] Roberto Hoyle, Robert Templeman, Steven Armes, Denise L. Anthony, David J. Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. In *The 2014 ACM Conference on Ubiquitous Computing, UbiComp '14, Seattle, WA, USA, September 13–17, 2014*, A. J. Brush, Adrian Friday, Julie A. Kientz, James Scott, and Junehwa Song (Eds.). ACM, 571–582. <https://doi.org/10.1145/2632048.2632079>
- [21] R. Johansen. 1988. Groupware: Computer Support for Business Teams.
- [22] Veronika Krauß, Alexander Boden, Leif Oppermann, and René Reiners. 2021. Current Practices, Challenges, and Design Implications for Collaborative AR/VR Application Development. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8–13, 2021*. ACM, 454:1–454:15. <https://doi.org/10.1145/3411764.3445335>
- [23] Kiron Lebeck, Tadayoshi Kohno, and Franziska Roesner. 2019. Enabling Multiple Applications to Simultaneously Augment Reality: Challenges and Directions. In *Proceedings of the 20th International Workshop on Mobile Computing Systems and Applications, HotMobile 2019, Santa Cruz, CA, USA, February 27–28, 2019*, Alec Wolman and Lin Zhong (Eds.). ACM, 81–86. <https://doi.org/10.1145/3301293.3302362>
- [24] 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, 320–337. <https://doi.org/10.1109/SP.2017.13>
- [25] 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, 392–408. <https://doi.org/10.1109/SP.2018.00051>
- [26] Germán Leiva and Michel Beaudouin-Lafon. 2018. Montage: A Video Prototyping System to Reduce Re-Shooting and Increase Re-Usability. In *The 31st Annual ACM Symposium on User Interface Software and Technology, UIST 2018, Berlin, Germany, October 14–17, 2018*. ACM, 675–682. <https://doi.org/10.1145/3242587.3242613>
- [27] Germán Leiva, Jens Emil Gronbaek, Clemens Nylandstedt Klokmose, Cuong Nguyen, Rubaiat Habib Kazi, and Paul Asente. 2021. Rapido: Prototyping Interactive AR Experiences through Programming by Demonstration. In *UIST '21: The 34th Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 10–14, 2021*. ACM, 626–637. <https://doi.org/10.1145/3472749.3474774>
- [28] Germán Leiva, Cuong Nguyen, Rubaiat Habib Kazi, and Paul Asente. 2020. Pronto: Rapid Augmented Reality Video Prototyping Using Sketches and Enaction. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25–30, 2020*. ACM, 1–13. <https://doi.org/10.1145/3313831.3376160>
- [29] Ewa Luger, Lachlan Urquhart, Tom Rodden, and Michael Golembewski. 2015. Playing the Legal Card: Using Ideation Cards to Raise Data Protection Issues within the Design Process. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI 2015, Seoul, Republic of Korea, April 18–23, 2015*, Bo Begole, Jinwoo Kim, Kori Inkpen, and Woontack Woo (Eds.). ACM, 457–466. <https://doi.org/10.1145/2702123.2702142>
- [30] Blair MacIntyre, Maribeth Gandy, Steven Dow, and Jay David Bolter. 2004. DART: a toolkit for rapid design exploration of augmented reality experiences. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology, Santa Fe, NM, USA, October 24–27, 2004*. ACM, 197–206. <https://doi.org/10.1145/1029632.1029669>
- [31] Nicolai Marquardt, Ken Hinckley, and Saul Greenberg. 2012. Cross-device interaction via micro-mobility and f-formations. In *The 25th Annual ACM Symposium on User Interface Software and Technology, UIST '12, Cambridge, MA, USA, October 7–10, 2012*. ACM, 13–22. <https://doi.org/10.1145/2380116.2380121>
- [32] Abraham Hani Mhaidli and Florian Schaub. 2021. Identifying Manipulative Advertising Techniques in XR Through Scenario Construction. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8–13, 2021*. ACM, 296:1–296:18. <https://doi.org/10.1145/3411764.3445253>
- [33] Leon Müller, Ken Pfeuffer, Jan Gugenheimer, Bastian Pflöging, Sarah Prange, and Florian Alt. 2021. SpatialProto: Exploring Real-World Motion Captures for Rapid Prototyping of Interactive Mixed Reality. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8–13, 2021*. ACM, 363:1–363:13. <https://doi.org/10.1145/3411764.3445560>
- [34] Michael Nebeling and Katy Madier. 2019. 360proto: Making Interactive Virtual Reality & Augmented Reality Prototypes from Paper. In *Proceedings of the 2019*

- CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04–09, 2019*. ACM, 596. <https://doi.org/10.1145/3290605.3300826>
- [35] Michael Nebeling, Janet Nebeling, Ao Yu, and Rob Rumble. 2018. ProtoAR: Rapid Physical-Digital Prototyping of Mobile Augmented Reality Applications. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21–26, 2018*. ACM, 353. <https://doi.org/10.1145/3173574.3173927>
- [36] Michael Nebeling and Maximilian Speicher. 2018. The Trouble with Augmented Reality/Virtual Reality Authoring Tools. In *IEEE International Symposium on Mixed and Augmented Reality, ISMAR 2018 Adjunct, Munich, Germany, October 16–20, 2018*. IEEE, 333–337. <https://doi.org/10.1109/ISMAR-Adjunct.2018.00098>
- [37] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79, 1 (Feb. 2004), 119–157.
- [38] 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>
- [39] Joseph O'Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. 2022. Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 4 (2022), 177:1–177:35. <https://doi.org/10.1145/3569501>
- [40] Lev Poretzki, Joel Lanir, and Ofer Arazy. 2018. Normative Tensions in Shared Augmented Reality. *Proc. ACM Hum. Comput. Interact.* 2, CSCW (2018), 142:1–142:22. <https://doi.org/10.1145/3274411>
- [41] Shwetha Rajaram, Chen Chen, Franziska Roesner, and Michael Nebeling. 2023. Eliciting Security & Privacy Informed Sharing Techniques for Multi-User Augmented Reality. In *CHI '23: CHI Conference on Human Factors in Computing Systems, Hamburg, Germany, 23 April 2023 - 28 April 2023*. ACM. <https://doi.org/10.1145/3544548.3581089>
- [42] Shwetha Rajaram and Michael Nebeling. 2022. Paper Trail: An Immersive Authoring System for Augmented Reality Instructional Experiences. In *CHI '22: CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April 2022 - 5 May 2022*. ACM, 382:1–382:16. <https://doi.org/10.1145/3491102.3517486>
- [43] 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*.
- [44] 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>
- [45] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J. Wang. 2014. World-Driven Access Control for Continuous Sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3–7, 2014*. ACM, 1169–1181. <https://doi.org/10.1145/2660267.2660319>
- [46] 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, 141–158. <https://www.usenix.org/conference/usenixsecurity19/presentation/ruth>
- [47] Florian Schaub, Ruben Deyhle, and Michael Weber. 2012. Password entry usability and shoulder surfing susceptibility on different smartphone platforms. In *11th International Conference on Mobile and Ubiquitous Multimedia, MUM '12, Ulm, Germany, December 4–6, 2012*. ACM, 13. <https://doi.org/10.1145/2406367.2406384>
- [48] RAYMOND SCUPIN. 1997. The KJ Method: A Technique for Analyzing Data Derived from Japanese Ethnology. *Human Organization* 56, 2 (1997), 233–237.
- [49] Adam Shostack. 2014. *Threat Modeling: Designing for Security* (1st ed.). Wiley Publishing.
- [50] Maximilian Speicher, Katy Lewis, and Michael Nebeling. 2021. Designers, the Stage Is Yours! Medium-Fidelity Prototyping of Augmented & Virtual Reality Interfaces with 360theater. *Proc. ACM Hum. Comput. Interact.* 5, EICS (2021), 1–25. <https://doi.org/10.1145/3461727>
- [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, CHI 2018, Montreal, QC, Canada, April 21–26, 2018*. ACM, 107. <https://doi.org/10.1145/3173574.3173681>
- [52] Sangho Suh, Sydney Lamorea, Edith Law, and Leah Zhang-Kennedy. 2022. PrivacyToon: Concept-driven Storytelling with Creativity Support for Privacy Concepts. In *DIS '22: Designing Interactive Systems Conference, Virtual Event, Australia, June 13 - 17, 2022*. ACM, 41–57. <https://doi.org/10.1145/3532106.3533557>
- [53] Ryo Suzuki, Rubaiat Habib Kazi, Li-Yi Wei, Stephen DiVerdi, Wilmot Li, and Daniel Leithinger. 2020. RealitySketch: Embedding Responsive Graphics and Visualizations in AR through Dynamic Sketching. In *UIST '20: The 33rd Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 20–23, 2020*. ACM, 166–181. <https://doi.org/10.1145/3379337.3415892>
- [54] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 3789–3806. <https://www.usenix.org/conference/usenixsecurity22/presentation/trimananda>
- [55] 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 *CHI '22: CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April 2022 - 5 May 2022*. ACM, 612:1–612:15. <https://doi.org/10.1145/3491102.3517728>
- [56] Daniel Vogel and Ravin Balakrishnan. 2004. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology, Santa Fe, NM, USA, October 24–27, 2004*. ACM, 137–146. <https://doi.org/10.1145/1029632.1029656>
- [57] Tianyi Wang, Xun Qian, Fengming He, Xiyun Hu, Ke Huo, Yuanzhi Cao, and Karthik Ramani. 2020. CAPturAR: An Augmented Reality Tool for Authoring Human-Involved Context-Aware Applications. In *UIST '20: The 33rd Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 20–23, 2020*. ACM, 328–341. <https://doi.org/10.1145/3379337.3415815>
- [58] Richmond Y. Wong, Deirdre K. Mulligan, Ellen Van Wyk, James Pierce, and John Chuang. 2017. Eliciting Values Reflections by Engaging Privacy Futures Using Design Workbooks. *Proc. ACM Hum. Comput. Interact.* 1, CSCW (2017), 111:1–111:26. <https://doi.org/10.1145/3134746>
- [59] David Wright. 2013. Making Privacy Impact Assessment More Effective. *The Information Society* 29, 5 (2013), 307–315. <https://doi.org/10.1080/01972243.2013.825687>
- [60] Enes Yigitbas, Jonas Klauke, Sebastian Gottschalk, and Gregor Engels. 2021. VREUD - An End-User Development Tool to Simplify the Creation of Interactive VR Scenes. In *IEEE Symposium on Visual Languages and Human-Centric Computing, VL/HCC 2021, St Louis, MO, USA, October 10–13, 2021*. IEEE, 1–10. <https://doi.org/10.1109/VL/HCC51201.2021.9576372>
- [61] Lei Zhang and Steve Oney. 2020. FlowMatic: An Immersive Authoring Tool for Creating Interactive Scenes in Virtual Reality. In *UIST '20: The 33rd Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 20–23, 2020*. ACM, 342–353. <https://doi.org/10.1145/3379337.3415824>
- [62] Leah Zhang-Kennedy, Sonia Chiasson, and Robert Biddle. 2016. The Role of Instructional Design in Persuasion: A Comics Approach for Improving Cybersecurity. *Int. J. Hum. Comput. Interact.* 32, 3 (2016), 215–257. <https://doi.org/10.1080/10447318.2016.1136177>

A APPENDIX: SUMMARY OF DESIGNERS' AR PROJECTS

| | AR Application Scenario | AR Modality | Environment | Characters Selected for Prototyping Mitigation Techniques |
|----|-------------------------------------------|--------------|-------------|-----------------------------------------------------------|
| D1 | American sign language education for kids | Hand-held AR | Home | Bystander, Custom |
| D2 | Connecting remote learners | Head-worn AR | Home | Bystander, Voice Hijacker |
| D3 | Game for elephant poaching awareness | Hand-held AR | Home | Bystander, Environment Manipulator |
| D4 | Immersive music player | Hand-held AR | Park | Gesture Hijacker, Graffiti Spammer |
| D5 | Navigation system for lunar missions | Head-worn AR | Park | Bystander, Environment Manipulator |
| D6 | Kids' game for connecting and exercising | Hand-held AR | Home | Gesture Hijacker, Voice Hijacker |
| D7 | Tutorial system for experts & students | Head-worn AR | Factory | Gesture Hijacker, Graffiti Spammer |
| D8 | Domino game for people waiting in queues | Hand-held AR | Park | Gesture Hijacker, Voice Hijacker |

Table 1: Novice AR Designers' Storyboard Topics. We summarize the AR interaction design projects that novice AR designers explored in our user study (Sec. 7), where we assessed the benefits and limitations of REFRAME for supporting an implicit threat modeling process. The designers explored a variety of AR usage scenarios involving education (D1-2), games and entertainment (D3-4, D8), a navigation system for lunar missions (D5), and a tutorial system for physical tasks (D7). Five out of eight participants designed for hand-held AR (D1, D3-4, D6, D8) and three designed for head-worn AR (D2, D5, D7). A majority of designers utilized the home environment (D1-3, D6), as many projects were aimed at connecting individuals during remote activities during the COVID-19 pandemic. In Task 3 of the user study (i.e., prototype mitigation techniques for two REFRAME characters), the designers selected the following characters (ordered by frequency): *Bystander* (4), *Gesture Hijacker*, (4), *Voice Hijacker* (3), *Environment Manipulator* (2), *Graffiti Spammer* (2).