Leaks from Otherworlds: Discovering Protocols for Spontaneous Crossing Mixed Realities

BOTAO 'AMBER' HU*, University of Oxford, UK
YILAN ELAN TAO, Simon Fraser University, Canada
KAI LUKOFF, Santa Clara University, USA
YUE LI, Xi'an Jiaotong-Liverpool University, China

Mixed reality supports public art, play, and self-expression, yet people sharing a street may inhabit mutually invisible digital layers. Making a layer discoverable can reveal information or demand attention before anyone chooses to engage. We study **spontaneous crossing**: noticing an unfamiliar mixed reality and negotiating whether to take part. In an HCI classroom workshop with 40 participants we used **Protocol Futuring**, a speculative method in which teams propose the rules of an imagined encounter, challenge them, and revise. Comparing the questions this raised with existing literature, we map five facets: **leakage** and discoverability, which covers deliberate signals and incidental exposure alike; the crossing process; intentionality; Reality Access Control; and interoperability. Walkthroughs of a menu-based invitation and a handshake gesture show how one entry choice reshapes later permissions and responsibilities. This initial design space connects playful encounters to questions of who may see, join, refuse, and shape a shared reality, including those affected without taking part.

Additional Key Words and Phrases: mixed reality, spontaneous crossing, Protocol Futuring, design space, discoverability, consent, plural realities

ACM Reference Format:
Botao 'Amber' Hu, Yilan Elan Tao, Kai Lukoff, and Yue Li. 2026. Leaks from Otherworlds: Discovering Protocols for Spontaneous Crossing Mixed Realities. 1, 1 (October 2026), 53 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

It has long been anticipated that mixed reality (MR) glasses will become ubiquitous in the near future, worn by everyone in the street. Films and speculative design repeatedly picture a city covered by digital layers. In *Hyper-Reality*, the street is saturated with overlays of services and advertising [131]; in *Rainbows End*, a layer of reality may be seen by one person or by millions [209]. HCI has used design fiction and co-speculation to ask what ubiquitous augmented reality (AR) would bring [53, 100], including its implications for shared perception and perceptual rights [146]. Research on augmented clothing, public play, and place-based memorialization gives elements of this envisioned future concrete forms [65, 89, 127]. AR is anticipated as a medium through which people experience places and encounter one another.

*Corresponding author

Authors' Contact Information: Botao 'Amber' Hu, University of Oxford, Oxford, UK, botao.hu@cs.ox.ac.uk; Yilan Elan Tao, Simon Fraser University, Burnaby, Canada, elan_tao@sfu.ca; Kai Lukoff, Santa Clara University, Santa Clara, CA, USA, klukoff@scu.edu; Yue Li, Xi'an Jiaotong-Liverpool University, Suzhou, China, yue.li@xjtlu.edu.cn.

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.
© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
Manuscript submitted to ACM

Ubiquitous glasses do not ensure a shared digital city: people may inhabit different **actualities**, or currently experienced realities [12]. Current privacy protections and permission arrangements make sharing conditional. Content rendered on personal XR displays is ordinarily visible only to the wearer; systems such as Nearby SharePlay organize nearby sharing through explicit invitation and acceptance [7, 75]. Access also depends on applications and platforms, so co-location does not by itself establish a public layer that everyone enters [142, 196, 213]. Personalized augmentation is **permissionless**: a person can alter their own view without the agreement of others nearby [87]. The freedom to have one's own reality does not ensure that others can encounter it. Perceptual asymmetry, together with privacy controls and platform boundaries, can sustain private, fragmented, and siloed realities [34, 75, 156, 199].

Occasions to share mixed reality nevertheless keep arising. Friends who meet by chance want to revisit a memory attached to the place; a passer-by wants to join a street game in progress; an artist hopes that strangers will encounter a work. These encounters depend on **discoverability**. AR activism makes the tension especially clear: a creator can place an intervention without a site's owner's permission, yet still struggle to make it discoverable [190, 191]. Public art and heritage similarly depend on cues, demonstrations, and access routes that bring otherwise hidden content to attention [65, 222]. **Discoverability sits in tension with permissionless expression and privacy:** creators need ways for their realities to be encountered, while wearers and bystanders need control over what they reveal and receive. Permission to publish does not ensure that someone will notice, and public availability does not establish the awareness needed to consider participation.

From a bystander's position, the problem begins before access: **how does someone first notice that another reality exists here at all?** We call the clues that cross this perceptual boundary **leakage**. Some clues are offered deliberately, such as a projected image, a sign, or a spoken invitation. Others arise incidentally: people casting virtual spells expose their bodily movements even when the spells remain invisible; sounds and device appearance similarly disclose that an experience is taking place [89, 171]. Videos of AR gameplay in public show bystanders seeing odd postures and gatherings rather than the game itself [123]; people working on virtual displays are likewise noticed by those nearby, even when their content remains unintelligible [155]. Leakage therefore includes both designed attraction and exposure that does not depend on the designer's or wearer's wish. It can disclose an experience's existence without revealing its content or granting access. It is in this gap between noticing and participating that a request to join, an invitation, or a decision to remain outside can arise. We keep one term across that range because both ends have the same effect on a bystander: another world makes itself felt across a perceptual boundary without thereby granting entry. The word does carry a connotation of escape beyond an intended limit, which fits incidental exposure more comfortably than it fits a creator openly drawing strangers toward a work. We therefore treat intention as a dimension along which leakage varies rather than a condition for it, and Section 5.3 maps that dimension as a facet in its own right.

Following such a leakage clue can initiate **spontaneous crossing into mixed reality**. Spontaneity describes how the encounter begins: someone initially unaware of a layer notices it and considers engaging. Subsequent requests, consent, participation, refusal, and exit can be deliberate. Collaborative MR, cross-reality, and transitional interfaces provide a substantial history of connecting experiences [13, 21, 188]. Much of this interaction research takes place in laboratory settings or begins with participants who already know the reality they will encounter [39, 157, 187]. Spectator research, participation frameworks, and MR street play examine how people approach and join activities [89, 130, 172, 218]. Bringing these traditions together directs attention to how awareness, access, and shared activity become connected when prior knowledge and an agreement to connect cannot be assumed.

Spontaneous crossing must also work **in the wild**, where it becomes a **protocol problem**. Two people meeting by chance may use different devices, applications, representations, and platform rules. For A to enter B's world, their

systems need a mutually intelligible way to discover an offer, negotiate access, establish what becomes shared, and end the relationship. Device association, AirDrop [6], and Nearby SharePlay [7] offer precedents for connecting nearby parties; cross-platform architectures and XR standards [203] address parts of interoperability [62, 207]. Such encounters may require a **crossing protocol**: shared rules through which parties and systems establish an agreement about the crossing. The agreement can be limited to an encounter, preserving different worlds. Who defines those rules, whose permissions count, and how decisions can be challenged also make this a question of governing plural realities [87, 145].

Early work on spontaneous-AR consent [216] and some sharing protocols provides starting points [87, 91], but the design space of spontaneous crossing connecting cues, receiving equipment, invitations, access rights, and governing authorities remains unmapped. Our overarching research question is: **What is the design space for enabling and governing spontaneous crossing between plural mixed realities?** We examine it through three connected dimensions:

- **Leakage and discoverability in spontaneous crossing:** How can leakage make another reality noticeable and intelligible to initially unaware bystanders while accommodating privacy and selective disclosure?
- **Spontaneous crossing protocols:** What protocol arrangements support discovery, entry, participation, and exit across different devices, applications, and platforms?
- **From spontaneous crossing to the governance of plural realities:** How can these crossings and the realities they connect be governed, including who sets, enforces, and can contest shared rules?

Protocol futuring [88] explores unknown protocol design spaces through a participatory, adversarial speculative process that iterates through blue-team proposals, red-team challenges, and revisions. We conducted a **protocol futuring** workshop with 40 participants to open the design space and develop a provisional **Protocol Design Space**. We then compared it with existing mechanisms and prior literature to extend and structure it. Our main contribution is this initial opening and structuring of the design space. We show what the second stage adds by returning to one workshop proposal: a team answered the question of who may enter a shared layer by reviewing uploaded code, and the map separates reviewing an artifact, admitting a person, and choosing what a receiving system renders (Section 5.6). We offer this as an exploratory approach whose usefulness in design practice remains to be tested.

2 Background

2.1 Collaborative mixed reality, cross-reality, and transitional interfaces

Connecting realities has a substantial history in HCI. Early collaborative AR supported face-to-face interaction around shared objects, independent viewpoints, and differentiated information layers [21, 23, 64, 174, 201]. Mixed-reality boundaries and distributed architectures connected heterogeneous displays, scene representations, and physical and synthetic spaces [15, 30, 184, 185]. Reviews subsequently organized collaborative MR through awareness, communication, coordination, spatial distribution, and the symmetry of participants' roles and technologies [43, 54, 188]. Sharing a spatial reference need not entail identical views or capabilities. These classifications matter because the person outside an experience may differ from its wearer in knowledge, equipment, and permitted actions simultaneously. The technical possibility of connecting their displays addresses only part of that asymmetry.

Traversable and transitional interfaces explored movement between physical activity, augmented views, and immersive environments [22, 109, 114]. Cross-reality taxonomies situate these transitions alongside asymmetric participation and collaboration across the reality–virtuality continuum [13, 63, 98]. AR/VR and AR/desktop collaboration investigate differentiated viewpoints and awareness cues [56, 70, 82, 144, 157]; co-located systems connect headset users, nonusers,

and public spectators [50, 76, 165, 211]. Transition research examines continuity, bidirectionality, object transfer, and control across contexts [39, 67, 160, 186, 187, 210], supported by dedicated prototyping environments [73].

Much of this interaction research starts with recruited participants in laboratory tasks or otherwise prepared encounters: collaborators are equipped, environments configured, and the existence of another reality established before interaction begins [39, 67, 157, 160, 187]. Even asymmetric participation can assume that an outsider already knows what the headset user is doing. These starting conditions leave a different question: how does someone who has no prior awareness of another layer first discover it and spontaneously initiate a relationship with it?

2.2 Private, personalized, and platform-bounded realities

Wearable and pervasive AR bring this question into everyday settings, where people can occupy the same physical place while experiencing different mediated actualities [12, 57, 72, 75, 128]. Personalized-reality research considers how individual data and preferences reshape perceptual experience; accounts of fragmented, competing, and autonomous realities examine the consequences for shared worlds and authority [34, 87, 142, 156, 199, 206]. A common physical setting no longer guarantees a common perceptual or informational public. Here, a private view describes selective visibility, a personalized view describes adaptation, and a silo describes restricted movement across applications or providers. Distinguishing them explains how fragmentation can persist even when people wish to share, and why making everything public would not resolve every boundary.

Privacy is a central concern within this trajectory. Security research distinguishes sensing, application access, and rendered output, while multi-user AR investigates ownership, selective disclosure, and permissions over shared objects [44, 121, 167, 181, 182]. Bystander studies examine recording awareness, privacy mediation, sensitive contexts, and consent concerning information about others [1, 18, 46, 111, 112, 147, 180]. Social-acceptability and dyadic-interaction research further shows why wearers' preferences cannot stand in for those of nearby people [3, 4, 37, 122, 135, 139, 179]. Historically, the Google Glass backlash, including the derogatory label "Glassholes," made these conflicts publicly visible [116]. Criticism of Facebook's Project Aria extends them to corporate spatial-data collection and meaningful public refusal [11].

Together, privacy controls, personalization, and commercial platform architectures establish a trajectory toward AR that is private in view, customized in content, and bounded by applications and ecosystems [11, 69, 190, 199]. These forces can reinforce fragmentation, although privacy protection does not require proprietary silos, and a view hidden from bystanders need not be private from its provider. The resulting tension concerns how people retain selective control while making chosen experiences available beyond their existing audience or platform. Privacy scholarship provides essential resources for that control. The problem arises when an experience has no intelligible way to signal its availability beyond the people and systems already authorized to see it.

2.3 Public expression, permissionless publication, and discoverability

There are substantial reasons to cross those boundaries. Public art uses augmentation to reinterpret shared surfaces and places [20, 41, 222]; activism supports counter-narratives, community expression, and social change [99, 183, 190, 191]. Heritage and memorialization make histories encounterable where they matter [5, 65, 68, 79, 143]. Location-based play organizes sociability around movement and place, while augmented fashion makes self-presentation available to an audience [19, 25, 55, 96, 127, 154]. These practices constitute the demand for public encounters: expression, remembrance, play, or recognition depends on someone finding and interpreting a layer. They also imply different publics. A memorial may seek continuing community stewardship, a performance temporary spectators, and an outfit

recognition from willing passersby. Designing for their encounter therefore requires attention to the relationship being offered, rather than assuming that every publicly situated augmentation seeks unrestricted access.

AR activism makes one aspect particularly explicit. *With or Without Permission* describes site-specific interventions that can bypass a physical site's owner's approval [191]. This is a situated sense of permissionless publication; it does not abolish platform rules, viewer permissions, or the claims of represented people. *Understanding AR Activism* identifies discoverability, maintenance, and proprietary distribution as continuing barriers [190]. An augmentation can therefore be publicly published without becoming visible to a passerby. Application availability, access lists, privacy settings, and platform policies still condition who can encounter it [69, 142, 182]. Urban and heritage scholarship adds that technical availability does not settle who can authorize, contest, or sustain a representation [5, 24, 58, 60, 93, 208].

Discoverability bridges these conditions through markers, signage, demonstrations, application distribution, and community mediation [65, 96, 143, 222]. Public-AR and spectator research distinguishes unnoticed content, visible activity, intelligible effects, observation, participation, and departure [38, 74, 103, 130, 171, 172, 218]. Street play and public XR show how bodily conduct can signal otherwise inaccessible digital activity [89, 102, 153, 155, 164, 198, 204]. The design tension is to make an offer perceptible without making its content compulsory, exposing protected information, or treating observation as agreement to participate. A coarse cue could disclose that something is available while withholding the protected content, as selective-disclosure and sharing research explores [167, 182]. Someone may notice that cue yet choose to remain a spectator. Publicness consequently involves both a route toward participation and an intelligible position outside it.

2.4 Spontaneous linkage: nearby discovery, association, and sharing

Spontaneous connection is already an established concern in ubiquitous computing. Interest-signaling badges, profile exchange, local awareness, and serendipitous social computing explored encounters without a prearranged collaboration [26, 52, 115, 170]. Situated-information systems connected environmental identifiers and contextual annotations to portable interfaces [113, 176, 178]. Device association made nearby connection actionable through pick-and-drop, shaking or bumping devices together, and synchronized input [83, 84, 175, 177]. Authentication research and cross-device taxonomies distinguish endpoint selection, physical action, communication channels, and trust [14, 27, 36, 62]. These histories make spontaneous linkage a familiar design problem with substantial technical resources.

They also distinguish what a connection means socially. Public surfaces, phone sharing, and progressive interfaces separate awareness, showing, handing over, session participation, and keeping a copy [97, 108, 126, 129]. Proxemic critiques caution against treating approach or inferred attention as authorization [71]. MR pairing, shared anchors, and registration supply further components for establishing endpoints and spatial agreement [40, 95, 134, 138, 193, 194]. Their successful operation still leaves the purpose, scope, and duration of sharing to be negotiated.

AirDrop illustrates these distinctions in everyday infrastructure: nearby discovery, identity/contact checks, peer-to-peer transport, and receiving controls jointly organize an exchange [6]. Research on private discovery, historical wireless attacks, contact-authentication leakage, and unsolicited delivery examines different parts of that arrangement [81, 125, 200, 219]. Nearby SharePlay similarly combines invitations and shared context within an ecosystem [7]. These precedents shift the question toward how discoverability, authentication, willingness to receive, and continuing access work together when the object of connection is another person's experienced reality. A brief exchange of a file and admission to an ongoing spatial experience can entail different expectations about mutual visibility, editable objects, and what remains after departure. The history of spontaneous association thus supplies components and distinctions that a crossing protocol must connect.

2.5 Interoperability standards and the governance of shared realities

Cross-platform work addresses the infrastructures on which such encounters depend. Metaverse surveys consider representation, identity, security, and institutional incentives [48, 213]. MR architectures investigate heterogeneous applications, object sharing, federated positioning, and discoverable resources with user-controlled access [17, 45, 169, 207]; decentralized architectures and spatial registries explore shared state and persistent anchors [85, 94]. Interoperability already encompasses multiple technical and organizational dependencies.

Standards address different layers of this work. Khronos OpenXR provides a common API across XR platforms and devices; W3C WebXR provides web applications with access to XR devices [203, 217]. glTF addresses transmission and loading of 3D assets, while OGC GeoPose specifies exchange of position and orientation [151, 202]. Open AR Cloud combines spatial-service discovery, content discovery, and positioning through its Open Spatial Computing Platform [150, 196]. The Metaverse Standards Forum coordinates requirements, resources, testbeds, and cooperation among standards organizations; it does not itself create standards [136].

Governance questions run through these infrastructures. WebXR explicitly addresses user intention, consent, and session duration [217]. AR research examines privacy management, normative tensions, perceptual agency, and expectations concerning bodies, possessions, and places [29, 101, 124, 145, 163]. Output management, access controls, and shared-state security enforce particular arrangements [119, 120, 182, 192]. A common format or API nevertheless leaves situated decisions about who may advertise a layer, admit others, withdraw access, or resolve competing claims. Cross-platform encounters require those decisions to remain intelligible across providers and affected parties. For example, a receiving system needs to distinguish permission to display an object from permission to copy or modify it; affected people need a way to understand and challenge the arrangement. This connects interoperability to the allocation and recognition of authority within an encounter.

2.6 Spontaneous crossing in the wild as a protocol problem

Spontaneous crossing takes place **in the wild**, where the people who encounter one another have not necessarily prepared to share an experience. Public XR encounters involve wearers, spectators, and passersby with different knowledge, interests, and opportunities to participate [89, 90, 155]. An unplanned encounter may occur within one application or across different devices and platforms; spontaneity does not itself imply incompatibility. Where a common application, account, spatial reference, or supported representation is absent, the parties must establish a supported form of exchange during the encounter. Where these resources already exist, noticing another reality and agreeing on access remain separate tasks.

This uncertainty makes spontaneous crossing a protocol problem. A protocol specifies the formats and procedures through which systems coordinate an association [189]. Standardizing those rules provides a basis for interoperability across independently developed devices, platforms, and applications. For crossing, we argue that this common specification needs to connect how an experience announces itself, how counterparts identify compatible capabilities, how an invitation and its permissions are expressed, and how access changes or ends. Existing standards supply parts of this foundation, including device access, asset exchange, and spatial references [151, 202, 203]. A crossing protocol would coordinate these components around the encounter, allowing implementations to offer a supported form of participation, such as a limited preview, or explain why crossing is unavailable. Interoperability here concerns both exchanging content and preserving the meaning of the agreement under which it is exchanged.

Direct work on these encounter-level arrangements is emerging. Windl et al. [216] develop scenarios and a consent-design framework for spontaneous interaction, but not crossing protocols. Hu et al. [91] propose a handshake-and-pull protocol linking nearby discovery, embodied initiation, alignment, and scoped sharing. Their proof of concept implements several components, but permission negotiation remains simulated and the complete encounter lacks participant or field evaluation. These contributions offer early design resources; their proposed arrangements still need examination in unplanned public encounters.

3 Methodological Overview: Futuring before Mapping

We use an **exploratory design inquiry: protocol futuring followed by literature-informed mapping**. Much collaborative and cross-reality research begins with participants already aware of an experience or prepared to share it (Section 2.1). We examine how an encounter becomes possible before that preparation. Protocol Futuring makes alternative arrangements concrete enough to question, without fixing an interface or platform in advance [35, 88].

Futuring elicits the problems to investigate. In one HCI classroom, blue teams propose an encounter, red teams challenge its assumptions or consequences, and blue teams revise. We examine what each challenge questions and what its proposed repair changes. Appendix B (Table B1) traces these arguments into questions for literature comparison, following design-space approaches that develop distinctions through examples [51, 78].

Mapping develops those questions through literature. We use a question-led, integrative comparison to find mechanisms, alternatives, and qualifications [106, 195]. Its vocabulary develops with the questions: making an outfit visible can lead to spectatorship, selective sharing, content admission, or infrastructure maintenance. A search restricted to "spontaneous cross-reality" would miss relevant work. Section 4.4 explains how comparison organizes these concerns into five facets.

Our purpose is to construct a design space across partly disconnected literatures while its questions and boundaries develop. We therefore use literature-informed mapping rather than report a systematic scoping review under PRISMA-ScR [205]. Uncertain terminology does not preclude scoping review; our choice concerns the inquiry's generative purpose. We report comparison procedures and source relationships without claiming exhaustive retrieval or saturation.

The map supports further Protocol Futuring and participatory deliberation by making a proposal's implications and omissions discussable, including where a repair shifts responsibility or authority.

4 Stage One: Futuring Spontaneous Crossings

4.1 Classroom setup and group prompts

We recruited 40 students from one HCI classroom. Participants worked in teams of four, comprising **two blue-team members and two red-team members**. Blue pursued a shared encounter goal, while red challenged the means and consequences of achieving it. Appendix D (Table D1) lists participant characteristics.

We prepared ten scenarios and use **G1–G10** to identify them. These labels index scenario tracks, not teams: teams chose freely, and several worked the same scenario independently. Teams recorded their exchanges on a shared pad, which holds twelve numbered sections; eleven contain recorded blue/red exchanges and one was left blank. The recorded material addresses six of the ten prepared tracks. **We worked from every record verbatim.**

We began with science-fiction inspiration about future mixed reality, inviting participants to imagine ordinary streets occupied by people experiencing different digital layers. Scenario prompts translated this broad future into situations involving fashion, public play, heritage, agent-mediated encounters, and different application ecosystems. An

AR-fashion worked example, video references, and a preparatory leakage-channel map helped participants consider how someone outside a layer might notice and approach it. These materials opened possibilities; they did not prescribe a common device, application, or sharing arrangement.

4.2 Ethics, consent, and researcher position

The study received a favourable ethics opinion from the last author's institution [Name redacted]. Research use of classroom material was voluntary and separable from teaching. Every student received the same two-hour activity, the same cards, the same responsibilities, and the same closing discussion, whether or not they agreed to research use. A researcher without grading authority distributed the information sheet in advance and collected decisions privately. Individual decisions were not disclosed to instructors or classmates, and declining or later withdrawing did not affect grades, attendance, course credit, feedback, or any relationship with teaching staff. Participants could pass on a prompt, take a break, or stop without giving a reason.

Consent was recorded as separate decisions for analysis, written quotation, and reproduction of team material, so agreeing to one did not imply the others. Because a team record carries contributions from several people, a record was eligible only where every contributor to it agreed; material linked to a declining or withdrawing contributor was excluded, including work made by students who had agreed. Names were replaced with codes, and the name-to-code key is held separately from the analyzed material and is not part of this submission. We did not record room audio or video, photograph faces, or quote spoken classroom discussion; the material we quote from teams is written team text. Participants were asked not to write names or identifying details about other people onto team material.

Two features of our position shape what this study can claim. The workshop took place inside a course connected to the author team, so students encountered the activity through an existing teaching relationship, and the measures above reduce but do not remove that asymmetry. An author also facilitated the AR-fashion discussion reported in Section 4.3, which makes that account closer to the research team's own framing than the records that teams wrote without facilitation. We therefore keep facilitated and unfacilitated material separate throughout, and treat neither as evidence of what people would do in a public setting.

4.3 Proposal, challenge, revision, and comic rendering

Each group began with a desired encounter, such as helping a willing passer-by discover an augmented outfit. **Blue proposed a way to achieve the goal; red challenged a concrete assumption or consequence; blue revised the proposal in response.** The next challenge addressed the revised arrangement, allowing the discussion to develop rather than repeatedly restart. Participants recorded the exchanges in writing. The worksheet requested at least five rounds and a ten-panel script retaining the questions and responses, including unresolved consequences.

The exercise invited challenges across four connected levels:

- **Interaction:** How does the encounter begin, and what must people and devices do?
- **Intention:** Who wants to reveal, notice, approach, participate, or remain outside?
- **Sharing and privacy:** What becomes visible or accessible, to whom, and with whose agreement?
- **Governance:** Who defines and enforces the rules, including those of a proposed public layer?

These prompts encouraged changes in perspective: a technical proposal could raise a social concern whose resolution introduced an institutional dependency. Groups could follow these consequences across levels without covering them in a fixed order.

What the analysis uses. Our analysis uses the written exchanges that teams recorded during the session. Six scenario tracks (G1–G6) carry recorded team material; G7–G10 were prepared but not taken up by any team, and we use those four names only as illustrations, never as evidence. Teams wrote in Chinese, in English, or in both; we worked from the original wording, and where a team supplied its own English rendering we used that alongside.

Two AR-fashion sources are kept separate. The five-round sequence presented in Table 1 and Figure 2 is an author-facilitated account of the AR-fashion scenario, developed by an author with students and reconstructed afterwards. The classroom team record for the same scenario is a different discussion. We do not merge them, and we do not count them as two independent cases.

What generative AI did and did not do. Generative AI rendered the comic panels in Figure 2 and expanded the Appendix A portfolio dialogues. Those dialogues illustrate what each prepared scenario invites; they are not reports of what any team said, and no analysis in this paper rests on them. Generative AI did not produce, edit, or summarize the team records, and it did not perform the analysis. Appendix C states the full disclosure.

4.4 Protocol Futuring portfolio

The Protocol Futuring portfolio comprises ten scenarios, identified as G1–G10: **G1, AI Fashion; G2, Street Magic; G3, Shop with Agents; G4, AR Qupid; G5, City & City; G6, Live Heritage; G7, Uninvited Gallery; G8, Reality Alert; G9, Invisible Target; and G10, Public Canvas.** Together, they situate spontaneous crossing in self-expression, public play, agent-mediated services and encounters, plural realities, cultural interpretation, emergency interruption, and shared public space. Figure 1 shows the illustrated situation prepared for each one. Appendix A (Table A1) presents these situations with their team goals; its worked dialogues are generative-AI expansions of the prepared scenarios, not records of what any team said. Here, we use the AR-fashion scenario to show how successive proposals and challenges open protocol questions.



Fig. 1. The ten prepared Protocol Futuring scenarios, (a)–(j) corresponding to G1–G10. Each situates a spontaneous crossing in a public setting: purple denotes virtual content visible only inside a layer, while grey shows what everyone shares. Appendix A gives each scenario's inspiration, team goal, and worked dialogue. Teams recorded exchanges for six of them, G1–G6.

AR Fashion illustrates the method's progression (Figure 2). Blue begins with the wish to make an augmented outfit visible, and red asks how an outsider would encounter it. A public layer addresses availability but raises the burden and safety of receiving unfamiliar garment content. Code review then prompts questions about who performs it, who controls the service, and what incentives support that role. Proposing a private operator brings exchange across operators into question; a common format subsequently raises compatibility across device generations. The goal remains recognizable, while its requirements expand from visibility to reception, review, authority, and continuity. Table 1 and Figure 2 follow the G1 discussion.

Table 1. Five proposal–challenge pairs from the author-facilitated AR-fashion account, corresponding to Figure 2. This is the facilitated account described in Section 4.2, not the classroom team record for the same scenario, which is a separate discussion. The final column identifies the protocol questions opened by each exchange.

Account and round	Blue proposal	Red challenge	Protocol questions opened
G1, B1 / R1	Make the wearer's virtual outfit visible to passersby.	Where would they see it, when their ordinary view does not show the augmentation?	**Discoverability and perceptual asymmetry:** what cue makes an unseen layer encounterable? Whose view must change, and what equipment does that require?
G1, B2 / R2	Create a public layer for AR fashion.	Must everyone download the garment content, and who consented to receiving it? What if it is large or contains malicious code?	**Reception, permission, resource burden, and safety:** does announcing a layer authorize receiving its contents? What can be previewed or checked before loading it?
G1, B3 / R3	Review garment code before upload.	Who controls the layer and performs the review, and what incentives do they have?	**Governance and admission:** who sets review criteria, operates the service, and has authority to admit content? What sustains that work?
G1, B4 / R4	A private operator could govern the layer, with running the AR cloud providing an incentive.	Can the same outfit work across different companies' public AR clouds?	**Operator dependence and interoperability:** what must providers share or recognize for an outfit to cross their boundaries? Which decisions remain provider-specific?

G1, B5 / R5	Define a common AR-fashion format so the outfit looks the same across platforms.	Will the format change as devices improve? What if older devices cannot support it?	**Temporal compatibility and continuity:** how are versions, device capabilities, and fallback handled? A shared format leaves its evolution and support obligations open.

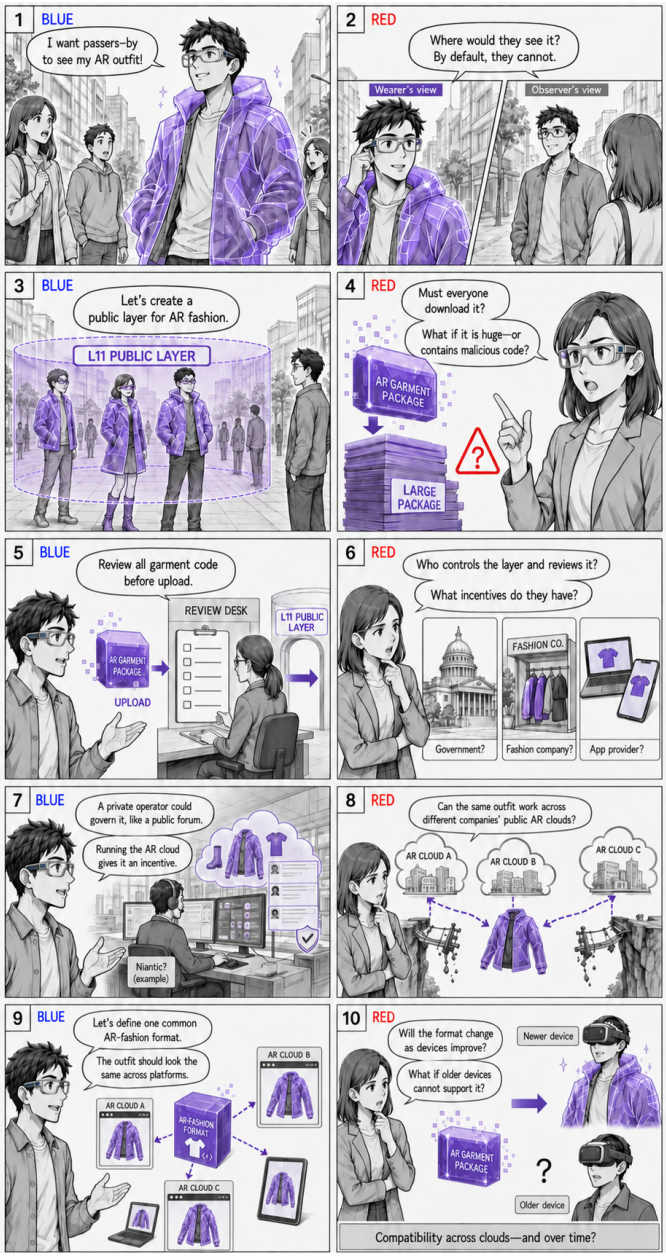


Fig. 2. G1: an author-facilitated AR-fashion discussion with students, reconstructed as a comic. Blue proposals and red challenges develop from visibility to review, governance, and compatibility; purple denotes virtual content.

How the confrontation develops the design space. Answering visibility creates a reception problem; answering reception through review creates an authority problem; selecting an operator exposes interoperability and continuity questions. The wearer's wish to be seen does not establish the receiver's permission, and reviewing code leaves the reviewer's authority open. Each resolution gives the next challenge something concrete to examine. Protocol Futuring thus produces candidate arrangements and connected questions through which to compare them.

4.5 From Protocol Futuring to a general map

Protocol Futuring opened an initial **general map of the design space** by making the conditions of an encounter contestable.

How we analyzed the records. The first author carried out the analysis on the complete set of recorded exchanges, working from the verbatim recorded text. Our unit of examination is a proposal, the challenge it encounters, and the revision or unresolved question that follows. Each such sequence was first articulated as what was at issue in it: an interaction, an affected party, a permitted operation, an information flow, a temporal dependency, or a rule-making authority. Sequences making the same demand were then grouped across teams, which is why one row of Appendix B (Table B1) can carry exchanges from more than one team working the same scenario. Each group was rewritten as a question that literature could be asked to answer, and the analysis stopped when every recorded sequence had been placed in at least one such question. The co-authors reviewed the resulting questions and their mapping back to the records, and disagreements were resolved by returning to the recorded text. **A single analyst produced the initial articulation; we make no claim of independent coding or intercoder agreement, and Section 6.3 treats this as a limitation.** We then developed each question through the literature comparison described below, which supplied the mechanisms, alternatives, and qualifications that the recorded exchanges raise but do not settle. The design space is therefore the product of both steps: the records fix which questions are asked, and the literature comparison fixes how finely each one can currently be answered. For example, G1 distinguishes publishing an outfit, receiving it, and reviewing its code, while G5's response to user admission through code review exposes a related distinction between approving a person and approving an artifact. These became questions for literature comparison: what mechanisms address each decision, which conditions do they assume, and which alternative arrangements remain possible?

Appendix B (Table B1) makes this elicitation visible. G1–G10 identify group discussions; B1/R1, B2/R2, and so forth locate each role's successive contributions. Descriptors distinguish exchanges addressing the same group situation. The final column articulates the comparison question and points to the facets that develop it. These locators refer to the source accounts described in Section 4.2; Appendix A's expanded dialogues illustrate possibilities separately.

From questions to literature comparisons. We used these articulated problems to guide targeted retrieval across the relevant research traditions, drawing on existing source collections, Consensus and Undermind searches, and checks of original papers and technical documentation. For discoverability, searches connected public AR and spectatorship to signage, mediation, and distribution; for access, they connected consent and selective sharing to refusal and revocation; for interoperability, they connected spatial services and formats to admission, custody, and rule-making. Sources were used where they supplied a mechanism, exposed an assumption, or offered a consequential alternative. Search records and revision ledgers preserve the query, source, comparison, and evidence boundary. An implemented component, a reported concern, and a speculative proposal support different kinds of inference.

We compared each selected source with the problem that motivated its use, splitting decisions that a proposal combined and qualifying what a mechanism could establish. For example, G6's sequence led us to distinguish making a cue available from someone noticing it and knowing how to proceed [222]. G1's reception question led to outgoing

sharing and incoming acceptance as separately controlled operations [182]. Existing participation research also retained observation as an outcome, extending the comparison beyond successful entry [130]. The preparatory leakage atlas and earlier intention concepts supplied starting resources; the discussions and literature comparison developed their implications and relationships.

Why five facets. Comparison organized the concerns around five distinguishable questions: **what becomes perceptible** (leakage); **how the relationship develops and ends** (crossing process); **what involved and affected people want** (intentionality); **which operations are authorized** (Reality Access Control); and **which systems and authorities recognize and sustain the arrangement** (interoperability and governance). These are analytical views, not five independently discovered classroom themes. Appendix B shows why concerns span them: time affects entry and grant expiry, privacy concerns cues and later data use, and equipment affects noticing and participation. The five facets organize the questions encountered so far, open to further futuring and comparison.

5 Stage Two: Mapping the Protocol Design Space

Building on the general map in Section 4.4, we develop five complementary facets: leakage channels and discoverability; the crossing process; intentionality; Reality Access Control; and interoperability across plural realities. These views connect how an encounter becomes possible to what it permits and what must sustain it.

We use the concerns in Appendix B (Table B1) to guide targeted literature comparison. For each concern, we distinguish the proposed operation, affected parties, and dependencies, then compare relevant mechanisms, empirical findings, and normative arguments. Literature can support a proposal, expose a limitation, or introduce an alternative. The preparatory channel atlas and existing intentionality categories provide starting vocabularies; the synthesis develops their relationships with the questions raised in the discussions. The five facets organize these comparisons around perceptual access, encounter progression, intentions, authorization, and cross-system coordination. Consensus and Undermind assisted literature retrieval and source examination; generative AI also assisted in organizing candidate comparisons, whose interpretations are presented in the analysis below.

5.1 Facet 1: Leakage channels and discoverability

Spontaneous crossing begins with a problem of discoverability: how can someone encounter a reality whose content is not yet available to them? We use **leakage channels** to describe the routes through which cues about an otherwise private experience become perceptible outside it. A cue may reveal content, signal that a layer exists, or make another person's activity noticeable, with different consequences for interpretation and privacy [103, 171, 219]. This facet maps what crosses that perceptual boundary, through which medium, and to whom. It asks how much an observer can discover before joining, what equipment or prior knowledge discovery requires, and how exposure can remain limited when either party prefers not to engage.

Figure 3 maps thirteen main channels and two subordinate forms. Table 2 attaches precedents to each receiving arrangement; the four families distinguish where information becomes available. Channels can combine, and neither their availability nor their technical operation determines notice, intention, or consent [171, 216, 219].

ID	Channel	Family	Receiving arrangement and supporting precedents
L01	Projection	Environmental Externalization	Shared-surface content and interaction without a personal display [76, 80, 98, 212]

ID	Channel	Family	Receiving arrangement and supporting precedents
L01A	Shadow Projection	Environmental Externalization	Interface footprint rather than full content [105]; compare graduated object awareness [104]
L02	Ambient Effects	Environmental Externalization	Physical actuation as an adjacent precedent [61]; public AR apparatus and perceptible effects [172]
L03	Markers / QR / NFC	Environmental Externalization	Garment/image targets [96, 152], QR entry/alignment [86], and NFC touch-to-fetch [31]
L04	Architectural Framing	Environmental Externalization	Situated signs and apparatus [38, 41]; spatially organized approach and interaction zones [42]
L05	Shared / Spectator Displays	Device-Mediated Signaling	Third-person views [50], prepared remote inspection/annotation [118], and AR-to-phone streams/copies [221]
L06	Device Audio Spill	Device-Mediated Signaling	Audible output and open-ear leakage [148]; designed light/tone concepts [28]
L07	Device Form Factor	Device-Mediated Signaling	Equipment, posture, and sensing indicators [28, 46, 103]
L07A	Face Displays / Reverse Passthrough	Device-Mediated Signaling	Wearer-state/attention display [33], outward touch displays [77], and perspective-correct eyes [133]
L08	Bodily Behavioral Signaling	Social-Behavioral Transmission	Play and audience activity [89, 218]; task-switch cues relevant to interruption [66]
L09	Social Invitation	Social-Behavioral Transmission	Explanation, handover, and co-located social approaches [89, 153, 172]
L10	Internet / Social-Media Circulation	Social-Behavioral Transmission	Distribution and mediation of AR art [41, 190]; online accounts of public encounters [60]
L11	Public Layer	Device-to-Device Leaking	Discoverable spatial resources [45], shared anchors [85], and federated mapping [16, 17]
L12	Proximity Discovery	Device-to-Device Leaking	Service discovery [219], gradual proximity exchange [129], gaze/spatial pairing [40, 193]
L13	Wearable Responses	Device-to-Device Leaking	Nearby social matching [26, 52]; visual, audio, and haptic guardian alerts on phones/watches [220]

Table 2. Leakage channels and supporting mechanisms. L01A and L07A are subordinate forms. Physical actuation provides an adjacent precedent for the illustrative airflow channel.

Environmental externalization varies the disclosed payload: content and controls [98], spatial occupancy [105], physical actuation [61], or entry/context information [42, 86]. A marker can announce an opportunity while its resolution requires another device; NFC supplies a deliberately close interaction [31]. The comparison concerns reach, equipment, interpretability, and control over effects in shared space.

For L01, projecting selected content can make an otherwise private view available on a shared surface [76, 80, 212]. The designer then chooses how much content to externalize and whether the surface also accepts interaction. L01A instead communicates a footprint: the shadow precedent concerns spatial awareness and collision avoidance, leaving

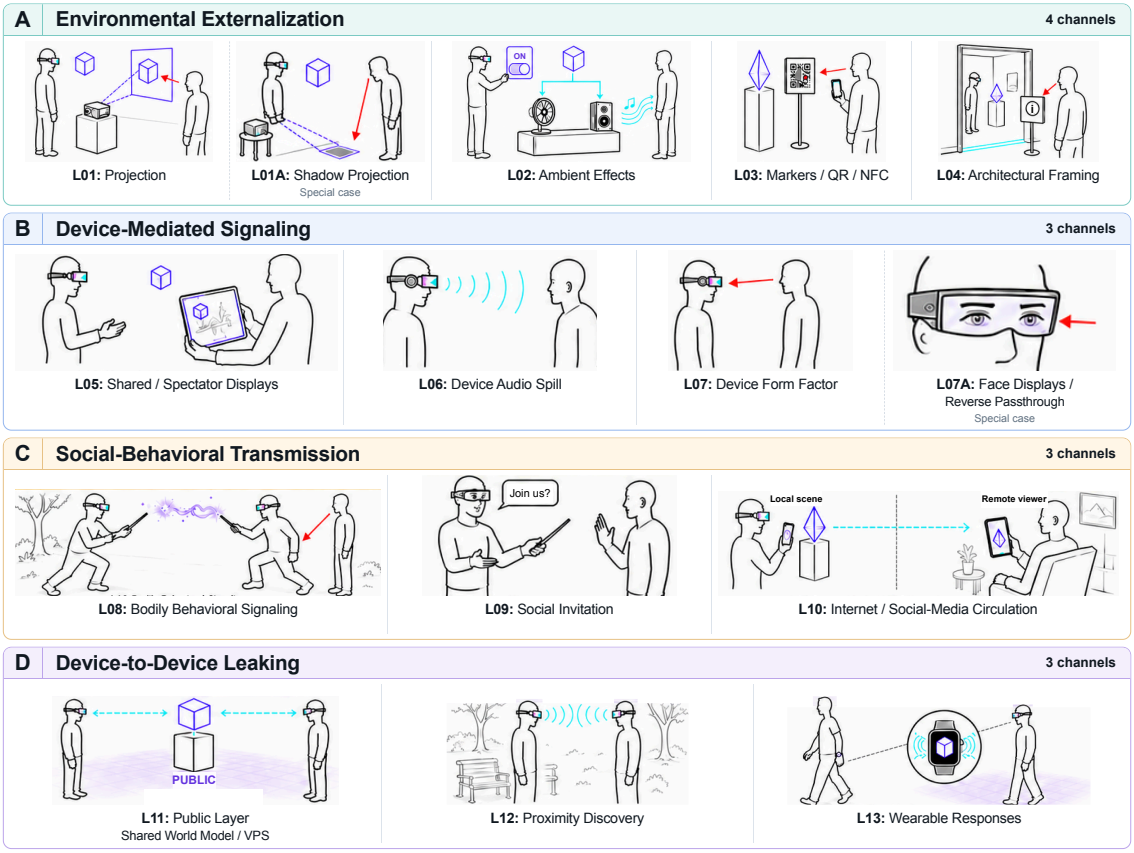


Fig. 3. Leakage channels across four families: thirteen main channels and two subordinate forms, L01A and L07A. Multiple channels can combine within an encounter.

the underlying interface partly undisclosed [105]. L02 raises a related interpretive question: a physical effect can reach someone who cannot identify its digital cause. L03 can preserve detail behind a deliberate scan or touch, whereas L04 places the encounter within signs, apparatus, and approach zones [38, 42]. These alternatives distribute discovery work between the environment and the observer; conspicuousness alone does not establish an understandable invitation.

Device-mediated signaling separates recognizing equipment [46], hearing activity [148], understanding wearer attention [33], and inspecting content [50]. These disclose different information to different audiences. Concept evaluations of multimodal indicators [28] and prototypes of outward displays [77, 133] offer different ways to make wearer state and content perceptible.

An L05 spectator display also involves a choice of perspective: selected content, an external view, and a streamed wearer view support different forms of inspection [50, 118, 221]. L06 may carry recognizable content or merely audible evidence of activity; open-ear design makes outward leakage relevant alongside inward awareness [148]. L07 identifies equipment without necessarily revealing what it is doing, while L07A adds a designed account of wearer state or a surface for interaction [33, 77]. We consequently separate the device being noticeable from the experience being legible.

Choosing a channel also means choosing whether to disclose activity, availability, attention, or content, and to whom that disclosure is directed.

Social-behavioral transmission can be incidental conduct [66], audience attraction [218], or an addressed invitation [172]. Circulation through AR-art networks [41, 190] may precede arrival: already knowing the layer exists changes the starting condition. We therefore retain when awareness arose and whether a signal actually addressed the observer.

The fashion premise makes this distinction concrete: a friend's reaction may reveal that something is happening without revealing the garment or establishing the wearer's willingness to share it. L08 can initiate curiosity, after which L09 supplies an explanation or explicit offer [89]. L10 can similarly supply context before a local encounter, but the circulated account and the experience available on arrival may differ. Our comparison asks what information travels, through whose interpretation, and with what continuation path. Social mediation can bridge missing interface information; it also makes another person's availability and explanation part of the crossing arrangement [153, 172].

Device-to-device leaking separates finding a resource [45], recognizing an endpoint [40], authenticating it [193], and receiving an alert [220]. Discovery itself can disclose identity or service information [200, 219]; privacy-preserving authentication supplies an adjacent countermeasure [81]. Compatibility, enrollment, notification settings, and permission determine who can continue.

L11 makes a resource addressable within some shared arrangement, raising the question of which devices can discover and interpret that arrangement [16, 45]. L12 narrows discovery through proximity or association, but physical nearness and a selected counterpart remain different criteria [40, 129]. L13 makes a machine-detected condition perceptible through a wearable response [26, 52, 220]. The alert can announce a possibility while withholding its content. A single encounter may therefore traverse several channels: bodily curiosity, a spoken explanation, proximity association, and a received preview. Mapping this chain exposes where an observer needs additional equipment or information, and where a new disclosure occurs before the next decision.

5.2 Facet 2: The crossing process, from discovery to exit

Spontaneous crossing unfolds through changes in the relationship between a person and another reality: becoming aware of it, understanding an opportunity, negotiating access, participating, and eventually disengaging. Each transition raises a different protocol question about what the person knows, what has been offered or agreed, and what they can now do. Work on public engagement and cross-reality transitions motivates attending to these changing relationships and their possible reversals [42, 130, 210]. This facet maps the decisions and dependencies along that process, including routes that end with observation or refusal. It also asks how changes during participation reopen negotiation, and what remains after access ends.

Figure 4 organizes **Discover → Make sense → Request or invite → Agree on scope → Scoped crossing → Withdraw or exit**. Engagement can transform or recede [42, 130, 218]; transitions can cross different reality configurations [13, 210]. These are analytical decision points, with observation, refusal, return, and renegotiation retained as branches.

Discover. A crossing opportunity first has to become perceptible through something outside the otherwise private experience. The leakage choices and perceptual-boundary detail in Figure 4 connect this first transition to the channel map in Section 5.1. Relevant precedents include visible actions and effects [171], signage [38], mobile-AR ambiguity [103], and public-XR curiosity [155, 198]. We distinguish a cue being available, someone encountering it, and that person noticing that another experience exists. These distinctions locate different design failures: an inaccessible cue, an overlooked cue, or an ambiguous activity. Discovery need not disclose the experience itself. A protocol might expose

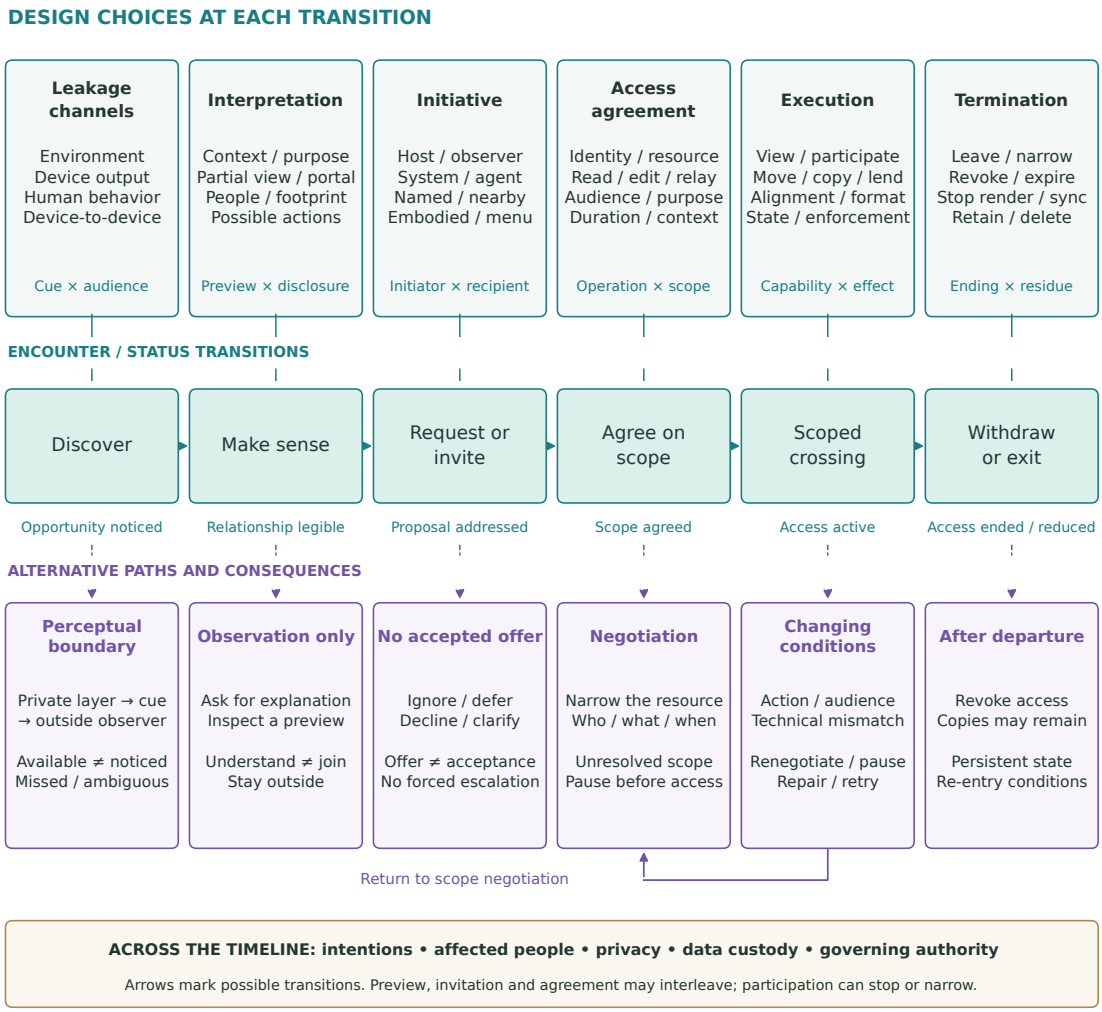


Fig. 4. The crossing process: design choices above, encounter transitions in the middle, and alternative paths below. Observation, refusal, renegotiation, and retained state remain possible; transitions can interleave or stop.

only its existence, an approximate footprint, or a requestable preview [104, 105]. The question is what becomes available to which potential observer, before any agreement to participate.

Make sense. Noticing an unusual reaction or display leaves open what is happening, whether observation is welcome, and how to continue. Interpretation can use progressive disclosure [129], a view across a boundary [15, 114], or an asymmetric spectator view [50, 98]. A preview therefore has at least two possible purposes: explaining the content and explaining the relationship being offered. An observer might understand the outfit yet remain unsure whether touching

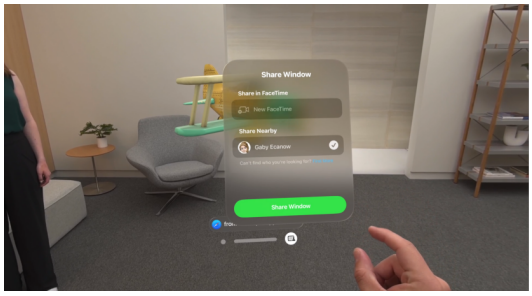
it changes the wearer's appearance. As G6 suggests, a cue also needs an intelligible next action. We retain observation as a possible endpoint: understanding enough to appreciate an experience need not culminate in joining it [130].

A **portal** makes sensemaking spatially concrete. In *Allow Me Into Your Dream*, a handshake opens a preview into the other person's environment; the opening expands before passage, while the previous world remains visible across the boundary [91]. This lets the proposed relationship be inspected through a partial view rather than inferred from a device list. For our map, the design variables are what the preview includes, whether it permits inspection without transfer, and whether releasing contact cancels the approach. A portal can explain an experience while still withholding particular objects, identities, or capabilities; the visible scene alone does not explain every permission attached to entry.

Request or invite. An offer can involve explanation and handover [89, 172], an embodied proposal [91], or cues supporting an interruption [166]. The initiative may come from the host, an interested observer, or a mediating system. We ask how that initiative identifies its recipient and communicates the proposed operation. Endpoint selection [36, 40] and authentication [62, 194] help establish which device or counterpart is involved; neither alone specifies whether the offer is to look, interact, or copy. This stage also permits silence, refusal, and clarification. A technically successful request can leave the social encounter unresolved, especially when the proposed recipient does not recognize an invitation as addressed to them.

Figure 5 compares two ways of making an offer legible and acting on it, covering both sensemaking and invitation. Apple's official visionOS demonstration uses a recipient list and Share Window action; its guide separately requires recipients to select Join [7, 10]. TouchPort instead proposes a bilateral handshake and pull [91]. The comparison concerns how a counterpart is addressed, how an offer becomes legible, and where someone can withdraw. Physical contact is itself a consequential request; a handshake does not authorize every subsequent action.

(a) Select a recipient and issue an invitation



Apple WWDC25, session 318, 01:15

(b) Preview through an expanding portal



TouchPort, Figure 1c (detail)

Fig. 5. Invitation forms, shown side by side: (a) the sender-facing Share Nearby selection and Share Window action in Apple WWDC25 session 318 at 01:15 [7]; (b) the expanded portal preview following the handshake in TouchPort, Figure 1c [91]. The images illustrate distinct ways to propose entry.

Agree on scope. Consent mechanisms depend on context and the operation being proposed [167, 216]. Agreement must distinguish recipient, sender, represented person, and sensed bystander [46, 147, 180]. These positions can overlap, but approval from one does not describe everyone else's position. In fashion, receiving an outfit, changing its local rendering, and changing what other observers see are different candidate agreements [121, 182]. A protocol can negotiate

these together or progressively as further actions become relevant. Our mapping asks which commitments are explained before crossing, which remain adjustable, and what happens when the parties agree to different scopes. An unresolved negotiation need not default to either full access or complete disconnection.

Cross within scope. Access can grant viewing or editing [182], a stream or independent copy [221], or an object transition [39]; spatial alignment remains a separate dependency [95, 134, 138]. We therefore specify what actually crosses: a person's participation, a representation, an object, or a capability over it. Sharing a view does not necessarily make both participants members of the same session, and a transferred object need not retain an ongoing connection to its source. Comparing configurations requires identifying the receiving environment, the permitted effects there, and the parties able to change them. A later request to edit or relay content returns the encounter to a scope decision, even when the initial crossing was seamless.

Withdraw or exit. Departure separates experiential disengagement [110, 158], ending temporary sharing [129], revoking future access [182], and managing stored or persistent material [45, 126]. Withdrawal can occur before entry: in TouchPort, releasing the handshake can abort the approach [91]. After entry, a person can stop looking while a session remains connected, or disconnect while a received copy persists. A protocol must specify which of these ends when someone leaves, what action communicates that ending, and whether the other party receives an intelligible indication. Exit can also be partial: a participant might stop interacting while retaining a spectator view. The process map consequently compares termination and reduced participation alongside continued involvement, without treating deeper engagement as the preferred outcome.

5.3 Facet 3: Intentionality spectrum

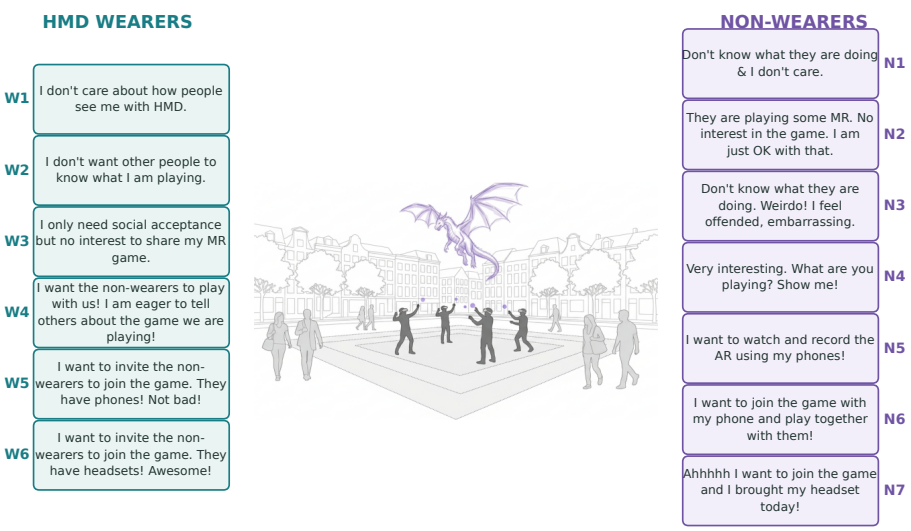
The encounter involves more positions than willing and unwilling. Hu et al. [90] distinguish six wearer and seven non-wearer positions. Figure 6a illustrates these varied intentions; Figure 6b develops them into a protocol-oriented repertoire. Here, **spectrum denotes possible relationships**, with observation, refusal, and participation treated as distinct outcomes rather than ranks.

The host's spectrum includes indifference to being seen, keeping activity private, seeking social acceptance without sharing, explaining the experience, and inviting participation through different devices [90]. These alternatives distinguish social visibility from access to content. A wearer may accept bodily observation but withhold a spectator stream, or explain a game without opening its session. Public gesture acceptance and performer preferences likewise vary with setting and audience [3, 4, 179]. Broadcasting an outfit can consequently express a wish to be appreciated without a wish to converse, be recorded, or allow modification. The protocol should specify which relationship the host is offering.

The observer's spectrum separates unawareness or indifference, tolerance, disturbance, curiosity, mediated spectatorship, and intended participation [90]. Knowing that a game exists does not imply wanting its content; wanting an explanation differs from wanting a phone view, and both differ from becoming a player. An affected passer-by may seek relief from obstruction while having no interest in the experience itself. Public engagement research supports retaining changing and peripheral positions [42, 130, 218]. Interruption can seek the wearer's attention for an unrelated purpose [66, 148, 166]. Treating every approach as recruitment would erase this possibility.

Being affected does not imply wanting to join. A person can be unaware of an augmentation yet be sensed, represented, avoided, or approached because of it. This is a relationship to the encounter's consequences, not a position on a scale of willingness. G2's discussion makes a passer-by an obstacle in someone else's game; the illustrative G9, Invisible Target, explores how invisible game activity can disrupt ordinary movement. Someone may want an

(a) Positions reported by Hu et al., redrawn



Statements are quoted from the positions reported by Hu et al.; the scene is redrawn and does not assign a position to any individual figure.

(b) The same positions developed into a protocol-oriented repertoire

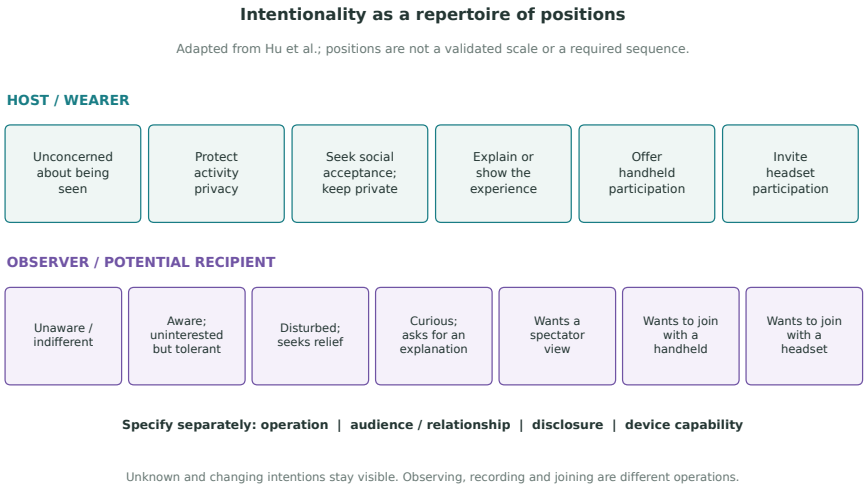


Fig. 6. Wearer and non-wearer intentions. (a) The six wearer (W) and seven non-wearer (N) positions reported by Hu et al. [90], with their statements quoted verbatim. We redrew the scene rather than reproducing the original figure, and it therefore does not assign a position to any individual figure in the drawing. (b) Our analytical development of the same positions: read each row as a repertoire rather than an ordinal ladder. Operation, audience, disclosure preferences and available equipment qualify each position; unknown or changing intention is not silently classified as agreement.

explanation, continued passage, or an effect to stop without wanting a preview or membership [46, 90, 147]. The intentionality facet therefore records both desired participation and desired protection from effects. Unknown intention cannot be replaced by the host's interpretation of a person's presence.

Participation, disclosure, and equipment are separate dimensions. The source distinguishes joining through a handheld device from joining through a headset. We retain this difference as a participation configuration, without interpreting the headset as evidence of stronger willingness. Someone can strongly want to join but lack compatible equipment; someone already wearing a headset may prefer to observe. Likewise, a spectator can want to record, while a player can refuse recording. Depiction [25, 180], sensing [1, 2], and object manipulation [167, 182] each introduce another operation whose intention must be specified. Wearers' balance between immersion and awareness adds a further contextual condition [117, 149].

Bilateral compatibility is negotiated across these positions. An eager sharer can meet an indifferent passer-by; a curious observer can approach someone protecting privacy; two willing participants can disagree about recording. Rather than assign an intention to a channel, we ask what each party seeks from this particular operation, audience, and moment. L08 bodily activity may invite curiosity without intentionally soliciting it; L05 may satisfy spectatorship; L09 may support either an accepted or declined invitation; L12 may advertise a service before its owner expresses interest. G4's mutual-matching proposal makes reciprocal interest relevant but does not settle all later uses. Unknown intention, silence, and inferred availability remain distinct from an explicit agreement [111, 216].

The spectrum therefore changes what counts as a supported outcome. A protocol may enable explanation without access, access without joint activity, participation without recording, or refusal without losing unrelated connectivity. We compare how easily parties can state, revise, or withdraw these positions and how the counterpart learns of that change. This keeps changing or unknown intentions visible in the protocol, including a request to contest an effect without entering. Section 6.2 develops the implications for people whose surroundings remain affected after they refuse or leave.

5.4 Facet 4: Reality Access Control: broadcasting, privacy, permission, and consent

Reality Access Control asks **who can access which part of a reality, perform which operation, under what conditions, and through whose enforcement**. Broadcasting specifies an exposure arrangement; privacy concerns collection and use; permission specifies an authorized capability; consent concerns the relevant parties' agreement [147, 180, 182, 216, 219]. A public announcement, an accepted invitation, and permission to modify an object are consequently different decisions. Figure 7 organizes the dimensions.

Granularity determines what sharing means. ShareAR explicitly distinguishes per-user and per-object viewing and editing permissions, including outgoing sharing and incoming acceptance [182]. These distinctions extend the question from "may this person enter?" to "may this person inspect this object, alter this property, or send it elsewhere?" A window can expose a rendered stream while leaving its underlying application state inaccessible [221]. A spatial object can expose geometry while withholding behavior or ownership metadata. For G1's outfit, publishing to a public layer, transferring garment data, and executing its code involve different permissions; approval for publication does not settle the receiver's consent. User concerns and expert elicitation support examining these boundaries explicitly [121, 167].

Broadcasting requires a specified payload and audience. A protocol may announce only that a layer exists, offer a coarse footprint, publish a preview, or distribute full content [45, 104, 105]. The audience may be named people, a group, nearby compatible devices, or anyone encountering a physical display. A public cue therefore need not expose

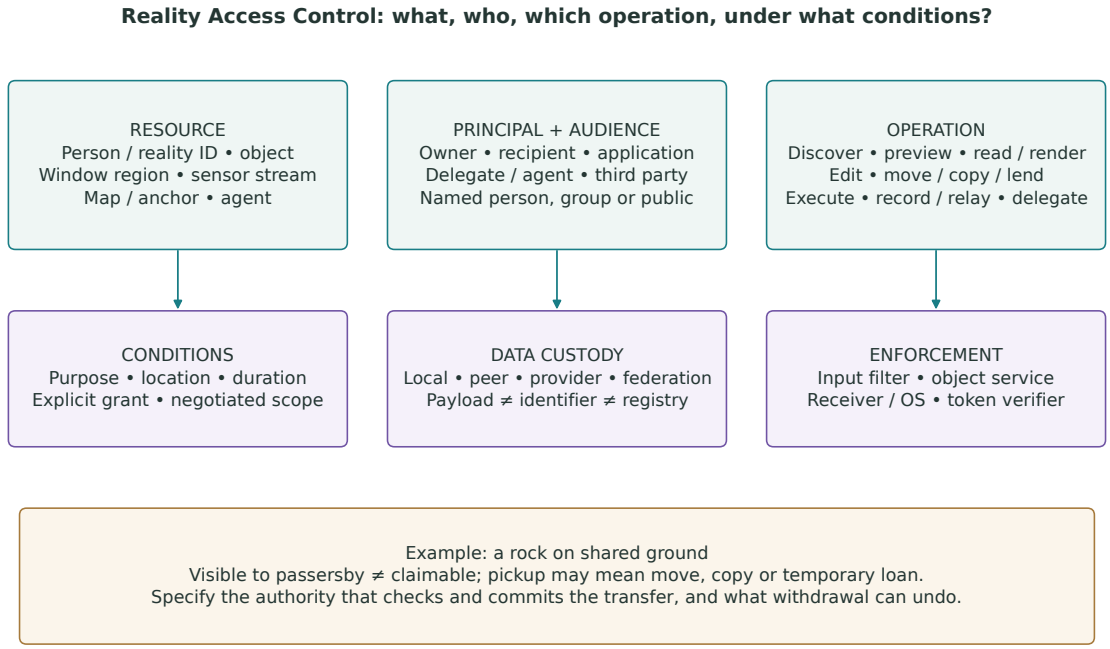


Fig. 7. Reality Access Control as intersecting resource, principal, operation, condition, custody and enforcement choices. The rock example is an author-developed comparison: public visibility alone does not determine whether pickup means moving, copying or borrowing an object, or who can authorize the operation.

a stable personal ID, and identifying a reality need not reveal its contents. Selective discovery can reduce information exposed before admission [81, 219], while projection and audible output can reach people who never opted into the digital session [76, 148]. We ask what crosses the boundary before agreement and whether the resulting exposure can actually be limited at the receiver.

Authority attaches to operations and affected parties. The sender controls an outgoing offer; the recipient controls acceptance and rendering; a represented or sensed person may have interests neither endpoint can settle alone [46, 145, 180, 182]. Privacy Equilibrium explores negotiation among conflicting sensing preferences through simulation and walkthroughs [168], highlighting negotiation among affected parties. For an agent, granting permission to negotiate availability need not grant access to private memory or authority to accept every invitation. Delegation research distinguishes the human principal, acting agent, and permitted scope [197]. Our extension asks which reality-level actions can be delegated and when another decision is needed.

Policy models offer alternative ways to express these choices. Named-user or object lists can specify who may view or edit [182]; attribute-based policies can evaluate the requester, resource, operation and current conditions [92, 107]; scoped tokens can carry delegated authority [91, 197]. A role such as spectator or moderator can summarize permissions, but its meaning still needs to be defined across services. The map asks which model fits a particular

encounter and how grants can be inspected, narrowed and revoked. It does not treat a policy engine as a substitute for understandable offers or agreement.

Custody and enforcement may be distributed. The payload may remain on a device, be streamed peer to peer, reside with a cloud provider, or be retrieved from an independently administered store [16, 45, 91]. Its identity record, permission metadata and spatial anchor may have different custodians. Local storage does not prevent later disclosure, and cloud storage does not imply public access. Input filtering [107], map-query policies [223], object-sharing controls [182], and receiving-side output mediation [120] enforce different boundaries. This separation makes it possible to ask where a refusal is checked and which downstream copy or service remains outside that check.

A dropped object makes the choices concrete. Suppose someone places a virtual rock on shared ground. Another person might see it yet lack permission to pick it up, obtain a copy while the original remains, borrow it until the encounter ends, or acquire the only authoritative instance. TouchPort discusses move, copy and lend semantics and a public-layer transfer scenario [91]; psychological-ownership research shows why visibility and perceived possession should not be conflated [162]. This illustrative extension makes the distinction between visibility and transfer concrete. Making such a protocol concrete requires naming who authorizes pickup, whether two recipients can claim the same instance, where the change is committed, and what happens if the connection fails midway.

Refusal, expiry and revocation have different scopes. A person may decline one invitation, block one source, stop rendering, withdraw a capability, or leave the session. These actions differ in effort, social cost and lost functionality [2, 111, 216]. G2's connectivity-shutdown response makes selective refusal especially salient. A grant may expire while a recorded view remains; a live stream may stop while its source object persists [126, 129, 221]. We therefore record both the operation that ends and the material that survives. Revoking future access cannot by itself undo an independent copy already received.

5.5 Facet 5: Governing interoperability across plural realities

Interoperability asks **which parts of an encounter different systems can recognize, exchange, and govern together**. Spontaneity concerns how people encounter a reality they did not already know about; heterogeneity concerns differences in devices, applications, services, and rules. These conditions vary independently. Two users of the same game can meet unexpectedly while inhabiting separate private layers, and two people can arrange a meeting across incompatible platforms. The design space covers both: discovery and agreement remain necessary within an ecosystem, while crossing its boundaries can introduce further coordination requirements.

Locate the boundaries before specifying access. "The same platform" may mean the same application, operating system, vendor ecosystem, content service, or spatial cloud. These are overlapping arrangements, not equivalent guarantees of compatibility or permission. Figure 8 identifies what may already be shared and which access decisions remain. It separates the recipient's capability to receive something from the authority to provide or use it [27, 182, 214].

Reconcile the information needed for the agreed operation. Figure 9 distinguishes spatial reference, representation, rendering and behavior, permission interpretation, and state continuity. For spatial reference, exchanging coordinates is insufficient when devices use different origins, maps, or tracking assumptions [95, 134, 138]. A shared anchor can align an encounter without granting access to the entire spatial map [223]. For representation, glTF supports asset delivery and OpenUSD supports scene composition, while scene-change exchange supplies another component [47, 159, 202]. A common format does not settle how an object behaves or who controls its authoritative state; an OpenUSD layer is not automatically a socially governed reality layer.

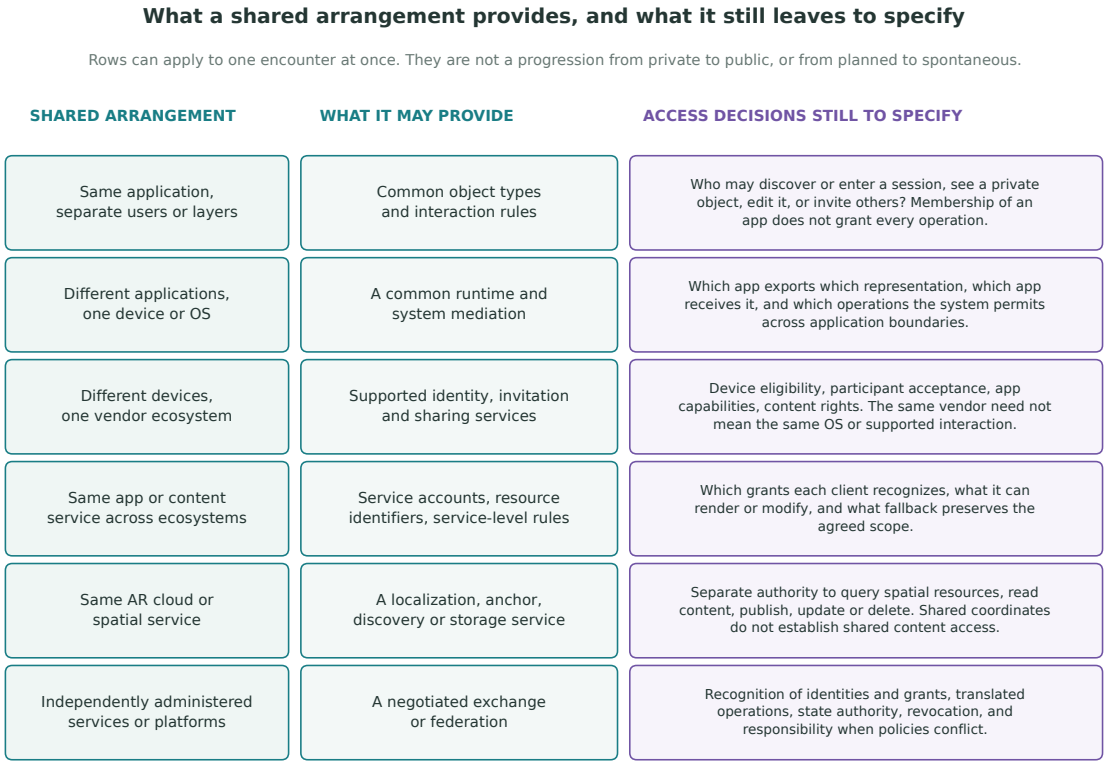


Fig. 8. Infrastructure arrangements and the access questions they leave open. Several rows can apply to one encounter; they are not a progression from private to public or from planned to spontaneous.

Rendering and behavior can therefore require a negotiated fallback. G5's common-format proposal still encounters a device unable to load the scene. A static preview, simplified object, or spectator stream changes what participation permits, and should be compared against the intended interaction [94, 185, 214]. Receiving systems also mediate unsafe output and conflicts between applications [119, 120]. A sender's wish to preserve appearance cannot determine the receiver's accessibility or safety settings.

Recognize a grant across boundaries. A receiver must interpret an authorization's issuer, recipient, resource, operation, conditions, and lifetime [182, 197]. Localization credentials need not authorize downloading or executing a garment. If one provider offers viewing while another permits copying, transfer has not preserved the agreement: a narrower operation, renegotiation, or refusal is needed. The comparison names where each decision is checked; state coordination separately determines whose update counts and how interrupted actions are resolved [47].

Specify custody and continuity. Names, anchors, objects, grants, and sessions can have different lifetimes and custodians. Apple's shared world anchors end with sharing [7], while later invitations can trigger collaborative map building under a previously accepted relationship for up to 45 days [9]. The latter is a condition on subsequent encounters

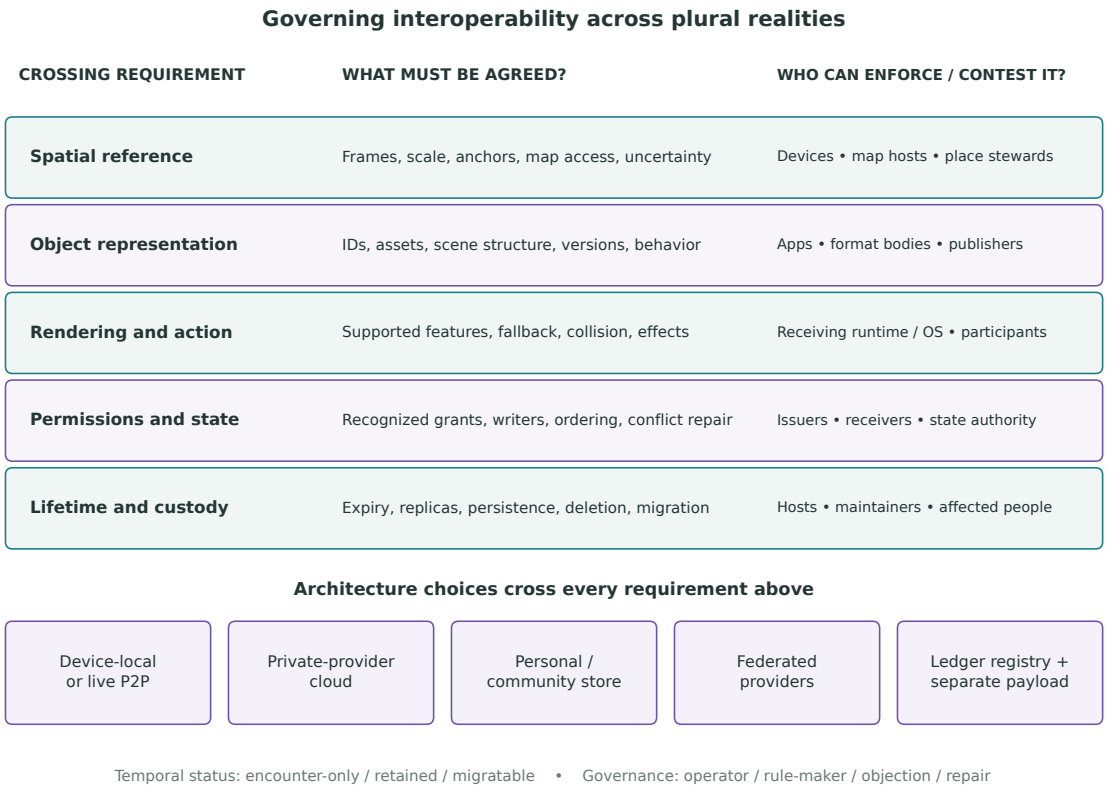


Fig. 9. Interoperability and governance across plural realities. Each row identifies a requirement, the information or rules to reconcile, and potential decision holders. The architecture choices below cut across all rows: a deployment can combine local state, a provider cloud, federation and a registry. Persistence and legitimate authority are not guaranteed by any single architecture.

with the same participant and location, not a map-retention period. Figure 10 separates storage arrangements from maintenance and authority.

Allocate authority beyond service administration. Devices can restrict sensing or output, apps can control membership and objects, and service operators can administer discovery and custody [107, 119, 182]. People represented or affected by an encounter may control none of these systems [142, 145, 163]. For each arrangement we therefore distinguish who sets a rule, who enforces it, who has standing to object, and who can provide a remedy. Technical administration alone does not answer those questions [58, 65, 69, 191].

A limited encounter profile can specify resources, coordinates, fallbacks, grants, state authority, lifetime, and withdrawal while preserving independent layers [34, 87]. Who can contest its effects remains consequential even within one app or cloud [173, 215].

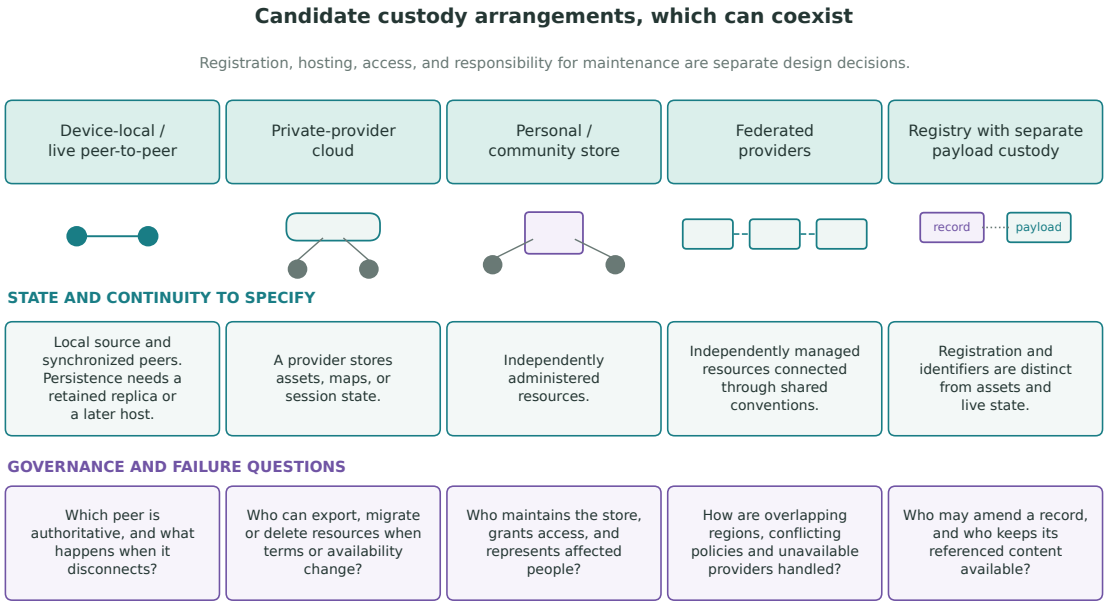


Fig. 10. Candidate custody arrangements, which can coexist. Registration, hosting, access, and responsibility for maintenance are separate design decisions.

5.6 What mapping adds: a workshop proposal revisited

The two stages divide the work differently, and one exchange shows the division concretely. A team working the City & City scenario (G5) set out to let people in two separate ecosystems see the same scene. It proposed a public cloud that either side could reach, and then a common format so that the scene would render identically across them. The red team asked how it would be decided which users may access it. The blue team answered that the code should be uploaded for review, and that access to the public layer would follow once the review passed.

Read through the facets, that single answer merges three decisions. Reviewing an uploaded artifact asks what is checked, against which criteria, and by whom, which is a question about shared representations and the authority that maintains them (Section 5.5). Admitting a person asks whose permission is recognized and how it is enforced against a particular requester (Section 5.4). Choosing what a receiving system renders asks what actually crosses and what fallback applies when a device cannot honor it (Sections 5.2 and 5.5). A passing review can certify that a scene is safe to load. It cannot say who may enter, and it does not settle what any given receiver should display. Appendix B (Table B1) records this concern; separating the three decisions is what the map contributes to it.

The separation changes the arrangement rather than only annotating it. Once admission is its own decision, the proposal needs a party that recognizes a request, a route for refusing one, and an account of what happens to a person who is refused, none of which an artifact review supplies. The team's own next turn indicates that the merged answer does not hold: asked how a public scene and a personal one would be kept apart, the team introduced selectable layers,

so the rendering decision reappeared as soon as the admission answer was pressed. Mapping makes that separation available before a further challenge forces it, and it carries the unresolved part forward as a question the next round can take up.

Futuring produced the question and a first answer; mapping separated the decisions that the answer had collapsed and returned a revised arrangement to the next round. We present this as an exploratory approach. It is a resource for further Protocol Futuring rather than a validated design procedure, and its usefulness in design practice remains to be tested.

5.7 Walking through SharePlay and TouchPort

We apply the five facets to **Nearby SharePlay** and to **TouchPort's handshake-and-pull proposal**. This author-led walkthrough compares documented platform features with a reported proof of concept whose permission-negotiation interface uses Wizard-of-Oz [7, 8, 10, 91]. It examines decisions and open questions; it does not evaluate performance or user experience, and neither system was studied with participants here.

Consider a passer-by who sees someone gesturing toward apparently empty space and asks to encounter the augmentation. The initial clue is bodily activity (L08), followed by a social approach (L09). Neither a recipient-selection menu nor a portal preview explains this prior noticing by itself. Figure 11 follows the encounter across all five facets.

The five facets applied to a deployed platform and a research proposal

	NEARBY SHAREPLAY / WINDOW SHARING documented platform behavior	TOUCHPORT reported two-device proof of concept
Leakage and discoverability	System UI lists eligible nearby recipients. Code-based pairing covers non-contacts; human communication makes that route discoverable.	Presence cues precede the hand gesture, and the handshake opens a portal preview. Dedicated bystander indicators are proposed but not implemented.
Crossing process	Select or pair, invite, accept, share. Shared world anchors have a session-bound lifetime; application data needs its own lifecycle decision.	Approach, handshake and preview, pull, explore, release. Release can abort before entry; later termination is tied to the proposed sharing terms.
Intentionality	A sharer initiates and recipients choose whether to join. Moving a shared window can be available without interacting with its underlying app.	Reciprocal grasp is treated as agreement. What a greeting, a refusal of touch, or a wish only to inspect would mean is left open.
Reality Access Control	Selected-window sharing differs from an app-supported collaborative activity. Shared coordinates do not synchronize application state; the app specifies that.	Scoped rights are distinguished, with move, copy and lend operations. The negotiation interface remains Wizard-of-Oz.
Interoperability and governance	The platform supplies session and spatial coordination; applications remain responsible for shared state. Compatible ecosystem support is a condition.	The implementation uses Meta Quest 3 and Meta spatial anchors. Grant recognition across independent providers remains a further design problem.

Neither column reports how an initially unaware bystander experiences these arrangements.
The columns carry different evidential weight and are set side by side for comparison, not as equivalent evidence.

Fig. 11. An analytical walkthrough across the five facets. TouchPort descriptions refer to its protocol and implementation account [91]; questions about alternative interpretations and extensions are our analysis. Neither column establishes how an initially unaware public bystander experiences these arrangements.

First, make the encounter discoverable without assuming entry. Code pairing supports SharePlay encounters between non-contacts; it does not explain how someone learns that sharing is possible. TouchPort's embodied offer also needs to be understood before it can be accepted or declined. Adding an outward cue could improve noticing while disclosing activity to others. A further futuring round can examine the cue's audience, content, and refusal route together.

Second, change the operation and revisit the agreement. Suppose the passer-by wants to handle an object. Repositioning a SharePlay window differs from changing its application's state; selecting a loan, copy, or move in TouchPort changes state and permission requirements. A gesture or an invitation cannot explain all these consequences. A candidate repair is an operation-specific offer, examined jointly through intention, resource, authority, and exit: who authorizes it, and which receiving systems can honor it?

Third, test departure and effects on others. Ending participation can stop live exchange while leaving edited or copied material, platform relationships, or physical disruption to be addressed separately. A person without compatible equipment may still need to pass through the occupied space. The next protocol question is therefore not only how a participant leaves, but how someone outside the session can request an explanation or relief. An accountless reporting route or a physically perceptible contact point would introduce its own operator and response obligations. These are proposed extensions for further futuring, not capabilities established by this comparison.

The map supports starting with an offer, changing one condition, and tracing which decisions need revision. Further Protocol Futuring can contest platform defaults, application choices, and participant agreements with affected stakeholders, exposing omissions before a repair commits the encounter to a wider arrangement.

6 Discussion

6.1 Spontaneous crossing in the wild

Spontaneous crossing brings questions of awareness, willingness, and the meaning of an invitation into the encounter itself. Shared infrastructure can prepare some conditions, but an unplanned encounter within one app still requires these understandings. Heterogeneity adds questions about whether representations, permissions, and state can travel between systems; it can also arise in planned meetings. Our contribution connects these concerns across collaborative MR, participation journeys, spontaneous consent, and encounter protocols [54, 91, 130, 216]. The design space makes their dependencies available for examination when neither prior agreement nor compatible infrastructure can be assumed.

Across the discussions, a local repair reallocates work and decision rights. G1's code-review proposal introduces a reviewer and an operating incentive; its common format introduces obligations toward changing devices. G2's suggestion to refuse participation by disabling Bluetooth and Wi-Fi instead places the cost on the passer-by, affecting connectivity beyond the encounter. These moves show why an arrangement's convenience must be considered alongside who sustains it and who bears its costs. Comparing facets makes these allocations explicit: a discovery mechanism, access grant, or compatible format leaves responsibilities elsewhere in the encounter that require further decisions.

The map supports a repeatable inquiry: hold an encounter goal explicit, compare arrangements, and trace the obligations introduced by each. If a brief view meets the goal, a bounded offer can avoid making custody of an editable object and modification rights prerequisites; handling, retaining, or redistributing it requires reconsidering those choices and their exit conditions. This connects mapping back to Protocol Futuring: subsequent discussion can address specific

unresolved decisions with those expected to authorize, maintain, or experience their consequences. Whether this analytical use improves design practice remains an empirical question.

6.2 Discoverability, permission, and governing plural realities

Discoverability requires some disclosure before someone can decide whether to engage. In G6, signs and lightweight notifications respond to the unwanted interruption of automatic history, yet still reveal activity and claim attention. G1 likewise distinguishes publishing an outfit from permission to receive it. These cases motivate staged, limited disclosure whose scope can be reconsidered as an encounter develops [147, 216]. Complete invisibility has costs too: public expression needs routes to an audience [190]. The design question is what becomes perceptible to whom before each decision, and what remains available to refuse.

Agreement among participants does not cover everyone affected. G2's play discussion represents a passer-by as an obstacle; the illustrative G9 extends this concern to players obstructing passage while following invisible objects. Someone can remain outside the session while being sensed, represented, or impeded. Independent platforms still share physical surroundings, where their effects can accumulate. Following DiSalvo [49], we understand publics through the consequences around which people may organize, extending attention beyond a platform's members or audience. We therefore argue that being affected should establish standing to seek explanation and review, independently of membership. This is a normative implication of the analysis, rather than an evaluated provision.

Such standing requires an accessible response. A person without an account or glasses could report a location and effect through a place steward or ordinary contact route, receive an explanation from the responsible operator, and request review. Discoverability here includes finding who can respond, without exposing an entire private layer. An objection cannot automatically confer a right to delete: obstruction, unwanted representation, and disagreement with public expression require different remedies. Uninvited Gallery and Public Canvas highlight the danger of giving a venue or platform unilateral control over expression, a tension also raised in site-specific AR activism [142, 191]. Response procedures must therefore make both responsibility and grounds for decisions contestable.

Protocols can give these provisional arrangements a concrete form [88]. G5's question about who sets fair relay rules extends to specifying access, custody, responsibility, and procedures for contesting and revising decisions. Implementing those procedures requires institutions able to respond. Infrastructure dependencies in Internet Protocol revision [141] and resistance to human-rights advocacy within the IETF [32] caution against expecting smooth convergence. Successive protocols may emerge through trials, disputes, and uneven adoption; widespread implementation would not establish public acceptance. The map provides a resource for examining each revision: which encounters it enables, whose claims it recognizes, and where it relocates burdens or leaves effects unaddressed.

6.3 Limitations and future work

One workshop with 40 students in an HCI classroom provided a situated opening of the space. Author facilitation, prompts, and preparatory materials shaped the discussion; selective literature comparison shaped its organization into five facets. Other perspectives and arrangements may remain outside this vocabulary. Four prepared scenarios were not taken up by any team, so the recorded material covers six of the ten tracks. The first author produced the initial articulation of the recorded exchanges, with co-author review but without independent coding, so the grouping in Appendix B reflects one reading of the records. Most participants reported no prior XR use, which limits what the discussions can say about experienced wearers and about people who have already encountered public AR. The cases

and analytical walkthrough support reasoning about dependencies, but do not establish saturation, technical feasibility, practical usefulness of the map, or public acceptance of a protocol.

Successive Protocol Futuring workshops should bring students and HCI researchers into dialogue with industry practitioners, governance specialists, place stewards, and affected nonparticipants. They should examine whether the map reveals overlooked parties or consequences, supports reasoned revisions, and encounters questions its facets cannot express. Iterations can assess conceptual saturation within a stated scope. Public-setting prototypes should then examine whether offers are understood, withdrawal has the promised effects, and people who never join can obtain a response. This would evaluate both the map's contribution to deliberation and the consequences of arrangements developed through it.

7 Conclusion

Spontaneous crossing offers a way for people to encounter one another's mixed realities while retaining different views, preferences, and boundaries. Through protocol futuring [88] across ten scenarios and literature-informed mapping, we developed an initial protocol design space of spontaneous crossings into mixed reality. The project shows how a desirable encounter can generate questions about compatibility, authority, and responsibility as its proposed rules develop.

References

[1] Melvin Abraham, Mohamed Khamis, and Mark McGill. 2024. Don't Record My Private pARts: Understanding The Role of Sensitive Contexts and Privacy Perceptions in Influencing Attitudes Towards Everyday Augmented Reality Sensor Usage. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 749–758. doi:10.1109/ismar62088.2024.00090

[2] Melvin Abraham, Mark Mcgill, and Mohamed Khamis. 2024. What You Experience is What We Collect: User Experience Based Fine-Grained Permissions for Everyday Augmented Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–24. doi:10.1145/3613904.3642668

[3] David Ahlström, Khalad Hasan, and Pourang Irani. 2014. Are you comfortable doing that?: acceptance studies of around-device gestures in and for public settings. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. Association for Computing Machinery, New York, NY, USA, 193–202. doi:10.1145/2628363.2628381

[4] Fouad Alallah, Ali Neshati, Yumiko Sakamoto, Khalad Hasan, Edward Lank, Andrea Bunt, and Pourang Irani. 2018. Performer vs. observer: whose comfort level should we consider when examining the social acceptability of input modalities for head-worn display?. In *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3281505.3281541

[5] Amanda Almond. 2021. *We Are All Related: (Re)Storying With Augmented Reality to Build Indigenous-Settler Relations*. Master's thesis. University of Alberta. doi:10.7939/r3-mszj-nh48

[6] Apple. 2021. AirDrop security. https://support.apple.com/guide/security/airdrop-security-sec2261183f4/web Apple Platform Security; published 18 February 2021; accessed 9 September 2026.

[7] Apple. 2025. Share visionOS experiences with nearby people. WWDC25, session 318. https://developer.apple.com/videos/play/wwdc2025/318/

[8] Apple. 2026. Configure your visionOS app for sharing with people nearby. https://developer.apple.com/documentation/groupactivities/configure-your-app-for-sharing-with-people-nearby Apple Developer Documentation. Accessed 9 September 2026.

[9] Apple. 2026. Nearby sharing security. https://support.apple.com/guide/security/nearby-sharing-security-sec201b7a876/web Apple Platform Security. Published 3 August 2026. Accessed 9 September 2026.

[10] Apple. 2026. Share apps and experiences with people nearby on Apple Vision Pro. Apple Vision Pro User Guide. https://support.apple.com/guide/apple-vision-pro/share-with-people-nearby-tanbccb085c1/visionos Official living documentation; accessed 9 September 2026.

[11] Sally A. Applin and Catherine Flick. 2021. Facebook's Project Aria indicates problems for responsible innovation when broadly deploying AR and other pervasive technology in the Commons. *Journal of Responsible Technology* 5 (2021), 100010. doi:10.1016/j.jrt.2021.100010

[12] Jonas Auda, Sarah Faltaous, Uwe Gruenefeld, Sven Mayer, and Stefan Schneegass. 2023. The Actuality-Time Continuum: Visualizing Interactions and Transitions Taking Place in Cross-Reality Systems. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, New York, NY, USA, 35–40. doi:10.1109/ISMAR-Adjunct60411.2023.00015

[13] Jonas Auda, Uwe Gruenefeld, Sarah Faltaous, Sven Mayer, and Stefan Schneegass. 2024. A Scoping Survey on Cross-reality Systems. *Comput. Surveys* 56, 4 (2024), 1–38. doi:10.1145/3616536

[14] Dirk Balfanz, Diana K. Smetters, Paul Stewart, and H. Chi Wong. 2002. Talking to Strangers: Authentication in Ad-Hoc Wireless Networks. In *Network and Distributed System Security Symposium*. Internet Society, Reston, VA, USA. https://www.ndss-symposium.org/ndss2002/talking-strangers-authentication-ad-hoc-wireless-networks/

[15] Steve Benford, Chris Greenhalgh, Gail Reynard, Chris Brown, and Boriana Koleva. 1998. Understanding and Constructing Shared Spaces with Mixed-Reality Boundaries. *ACM Transactions on Computer-Human Interaction* 5, 3 (1998), 185–223. doi:10.1145/292834.292836

[16] Sagar Bharadwaj, Anthony Rowe, and Srinivasan Seshan. 2025. Uniting the World by Dividing it: Federated Maps to Enable Spatial Applications. In *Proceedings of the Workshop on Hot Topics in Operating Systems*. Association for Computing Machinery, New York, NY, USA, 74–79. doi:10.1145/3713082.3730371

[17] Sagar Bharadwaj, Harrison Williams, Luke Wang, Michael Liang, Tao Jin, Srinivasan Seshan, and Anthony Rowe. 2025. OpenFLAME: Federated Visual Positioning System to Enable Large-Scale Augmented Reality Applications. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 706–716. doi:10.1109/ismar67309.2025.00080

[18] Divyanshu Bhardwaj, Alexander Ponticello, Shreya Tomar, Adrian Dabrowski, and Katharina Krombholz. 2024. In Focus, Out of Privacy: The Wearer's Perspective on the Privacy Dilemma of Camera Glasses. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3613904.3642242

[19] Arpita Bhattacharya, Travis W. Windleharth, Rio Anthony Ishii, Ivy M. Acevedo, Cecilia R. Aragon, Julie A. Kientz, Jason C. Yip, and Jin Ha Lee. 2019. Group Interactions in Location-Based Gaming: A Case Study of Raiding in Pokémon GO. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300817

[20] B. C. Biermann. 2018. Augmented Reality, Art, and Public Space. In *ACM SIGGRAPH 2018 Talks*. Association for Computing Machinery, New York, NY, USA, 1–2. doi:10.1145/3214745.3214747

[21] Mark Billinghurst and Hirokazu Kato. 2002. Collaborative augmented reality. *Commun. ACM* 45, 7 (2002), 64–70. doi:10.1145/514236.514265

[22] Mark Billinghurst, Hirokazu Kato, and Ivan Poupyrev. 2001. The MagicBook: a transitional AR interface. *Computers & Graphics* 25, 5 (2001), 745–753. doi:10.1016/s0097-8493(01)00117-0

[23] M. Billinghurst, S. Weghorst, and T. Furness. 1998. Shared space: An augmented reality approach for computer supported collaborative work. *Virtual Reality* 3, 1 (1998), 25–36. doi:10.1007/BF01409795

[24] Thomas Birtchnell, Pauline McGuirk, Christopher Moore, and Loren Vettoretto. 2020. Pay to play? Subverting the digital economy of Pokémon Go in the smart city. *Digital Geography and Society* 1 (2020), 100004. doi:10.1016/j.diggeo.2020.100004

[25] Jolie Bonner, Florian Mathis, Joseph O'Hagan, and Mark Mcgill. 2023. When Filters Escape the Smartphone: Exploring Acceptance and Concerns Regarding Augmented Expression of Social Identity for Everyday AR. In *29th ACM Symposium on Virtual Reality Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3611659.3615707

[26] R. Borovoy, M. McDonald, F. Martin, and M. Resnick. 1996. Things that blink: Computationally augmented name tags. *IBM Systems Journal* 35, 3.4 (1996), 488–495. doi:10.1147/sj.353.0488

[27] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt. 2019. Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–28. doi:10.1145/3290605.3300792

[28] Syed Ibrahim Mustafa Shah Bukhari, Maha Sajid, Bo Ji, and Brendan David-John. 2025. Rethinking Privacy Indicators in Extended Reality: Multimodal Design for Situationally Impaired Bystanders. In *2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, New York, NY, USA, 265–272. doi:10.1109/ismar-adjunct68609.2025.00059

[29] Andreas Butz, Clifford Beshers, and Steven Feiner. 1998. Of Vampire Mirrors and Privacy Lamps: Privacy Management in Multi-User Augmented Environments. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 171–172. doi:10.1145/288392.288598

[30] A. Butz, T. Hollerer, S. Feiner, B. MacIntyre, and C. Beshers. 1999. Enveloping users and computers in a collaborative 3D augmented reality. In *Proceedings 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. IEEE Comput. Soc, New York, NY, USA, 35–44. doi:10.1109/IWAR.1999.803804

[31] Qijun Cao, Junqi Zhang, Shengtao Fan, Jiaqi Rong, Menghao Qi, Zhuowen Duan, Peikun Zhao, Ling Liu, Zihao Zhou, and Wenjie Chen. 2024. NearFetch: Enhancing Touch-Based Mobile Interaction on Public Displays with an Embedded Programmable NFC Array. In *International Conference on Multimodel Interaction*. Association for Computing Machinery, New York, NY, USA, 234–243. doi:10.1145/3678957.3685749

[32] Corinne Cath. 2021. The technology we choose to create: Human rights advocacy in the Internet Engineering Task Force. *Telecommunications Policy* 45, 6 (2021), 102144. doi:10.1016/j.telpol.2021.102144

[33] Liwei Chan and Kouta Minamizawa. 2017. FrontFace: facilitating communication between HMD users and outsiders using front-facing-screen HMDs. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*. Association for Computing Machinery, New York, NY, USA, 1–5. doi:10.1145/3098279.3098548

[34] Sze Wing Chan. 2026. Fragmented Spatial Realities: Urban AR and the Fracture of Public Space. In *Companion Publication of the 2026 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 343–347. doi:10.1145/3802974.3809443

[35] EunJeong Cheon, Stephen Tsung-Han Sher, Šelma Sabanović, and Norman Makoto Su. 2019. I Beg to Differ: Soft Conflicts in Collaborative Design Using Design Fictions. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 201–214. doi:10.1145/3322276.3322350

[36] Ming Ki Chong, Rene Mayrhofer, and Hans Gellersen. 2014. A Survey of User Interaction for Spontaneous Device Association. *Comput. Surveys* 47, 1 (2014), 1–40. doi:10.1145/2597768

[37] Ji Won Chung, Xiyu Jenny Fu, Zachary Deocadiz-Smith, Malte F Jung, and Jeff Huang. 2023. Negotiating Dyadic Interactions through the Lens of Augmented Reality Glasses. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 493–508. doi:10.1145/3563657.3595967

[38] Ashley Colley, Leena Ventä-Olkkonen, Florian Alt, and Jonna Häkkilä. 2015. Insights from Deploying See-Through Augmented Reality Signage in the Wild. In *Proceedings of the 4th International Symposium on Pervasive Displays*. Association for Computing Machinery, New York, NY, USA, 179–185. doi:10.1145/2757710.2757730

[39] Robbe Cools, Augusto Esteves, and Adalberto L. Simeone. 2022. Blending Spaces: Cross-Reality Interaction Techniques for Object Transitions Between Distinct Virtual and Augmented Realities. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 528–537. doi:10.1109/ISMAR55827.2022.00069

[40] Matthew Corbett, Jiacheng Shang, and Bo Ji. 2024. GazePair: Efficient Pairing of Augmented Reality Devices Using Gaze Tracking. *IEEE Transactions on Mobile Computing* 23, 3 (2024), 2407–2421. doi:10.1109/tmc.2023.3255841

[41] Kristof Crolla and Garvin Goepel. 2022. Entering hyper-reality: Resonance-In-Sight, a mixed-reality art installation. *Frontiers in Virtual Reality* 3 (2022), 1044021. doi:10.3389/frvir.2022.1044021

[42] Peter Dalsgaard, Christian Dindler, and Kim Halskov. 2011. Understanding the Dynamics of Engaging Interaction in Public Spaces. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Cham, Switzerland, 212–229. doi:10.1007/978-3-642-23771-3_17

[43] Ryan Anthony J. de Belen, Huyen Nguyen, Daniel Filonik, Dennis Del Favero, and Tomasz Bednarz. 2019. A systematic review of the current state of collaborative mixed reality technologies: 2013–2018. *AIMS Electronics and Electrical Engineering* 3, 2 (2019), 181–223. doi:10.3934/electreng.2019.2.181

[44] Jaybie A. De Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2020. Security and Privacy Approaches in Mixed Reality: A Literature Survey. *Comput. Surveys* 52, 6 (2020), 1–37. doi:10.1145/3359626

[45] Maxim Van de Wynckel and Beat Signer. 2024. Discoverable and Interoperable Augmented Reality Environments Through Solid Pods. SoSy. https://ceur-ws.org/Vol-3947/short2.pdf

[46] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2377–2386. doi:10.1145/2556288.2557352

[47] Jaspreet Singh Dhanjan and Anthony Steed. 2021. Revisiting the Scene-Graph-as-Bus Concept: Inter-networking Heterogeneous Applications Using glTF Fragments. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, New York, NY, USA, 342–346. doi:10.1109/vrw52623.2021.00068

[48] John David N. Dionisio, William G. Burns III, and Richard Gilbert. 2013. 3D Virtual worlds and the metaverse: Current status and future possibilities. *Comput. Surveys* 45, 3 (2013), 1–38. doi:10.1145/2480741.2480751

[49] Carl DiSalvo. 2009. Design and the Construction of Publics. *Design Issues* 25, 1 (2009), 48–63. doi:10.1162/desi.2009.25.1.48

[50] Nam Hoai Do, Khanh-Duy Le, Duy-Nam Ly, Morten Fjeld, and Minh-Triet Tran. 2024. XRPublicSpectator: Towards Public Mixed Reality Viewing in Collocated Asymmetric Groups. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3613905.3650796

[51] Alix Ducros, Aurélien Tabard, Peter Dalsgaard, Raphaelle Bats, and Eva Eriksson. 2024. A design space for event-centric displays in public libraries. *Frontiers in Computer Science* 6, Article 1250861 (2024). doi:10.3389/fcomp.2024.1250861

[52] N. Eagle and A. Pentland. 2005. Social Serendipity: Mobilizing Social Software. *IEEE Pervasive Computing* 4, 2 (2005), 28–34. doi:10.1109/MPRV.2005.37

[53] Chloe Eghtebas, Gudrun Klinker, Susanne Boll, and Marion Koelle. 2023. Co-Speculating on Dark Scenarios and Unintended Consequences of a Ubiquitous(ly) Augmented Reality. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 2392–2407. doi:10.1145/3563657.3596073

[54] Barrett Ens, Joel Lanir, Anthony Tang, Scott Bateman, Gun Lee, Thammathip Piumsomboon, and Mark Billinghurst. 2019. Revisiting collaboration through mixed reality: The evolution of groupware. *International Journal of Human-Computer Studies* 131 (2019), 81–98. doi:10.1016/j.ijhcs.2019.05.011

[55] Leighton Evans and Michael Saker. 2019. The Playeur and Pokémon Go: Examining the Effects of Locative Play on Spatiality and Sociability. *Mobile Media & Communication* 7, 2 (2019), 232–247. doi:10.1177/2050157918798866

[56] Arthur Fages, Cédric Fleury, and Theophanis Tsandilas. 2022. Understanding Multi-View Collaboration between Augmented Reality and Remote Desktop Users. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–27. doi:10.1145/3555607

[57] Steven K. Feiner. 1999. The Importance of Being Mobile: Some Social Consequences of Wearable Augmented Reality Systems. In *Proceedings of the 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. IEEE Comput. Soc, New York, NY, USA, 145–148. doi:10.1109/IWAR.1999.803815

[58] Benjamin Feldman. 2018. Agency and governance: Pokémon-Go and contested fun in public space. *Geoforum* 96 (2018), 289–297. doi:10.1016/j.geoforum.2018.08.025

[59] Sophie Fiennes and Slavoj Žižek. 2006. The Pervert's Guide to Cinema. Documentary film. https://www.dfi.dk/en/viden-om-film/filmdatabasen/film/perverts-guide-cinema Directed by Sophie Fiennes; written and presented by Slavoj Žižek.

[60] Joshua A. Fisher. 2025. Virtual Movements, Real Consequences: Confrontations of Augmented Reality Social Media in Public Spaces. Oxford Academic. doi:10.1093/9780198945253.003.0015

[61] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. 2013. inFORM: Dynamic Physical Affordances and Constraints through Shape and Object Actuation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 417–426. doi:10.1145/2501988.2502032

[62] Mikhail Fomichev, Flor Alvarez, Daniel Steinmetzer, Paul Gardner-Stephen, and Matthias Hollick. 2018. Survey and Systematization of Secure Device Pairing. *IEEE Communications Surveys & Tutorials* 20, 1 (2018), 517–550. doi:10.1109/comst.2017.2748278

[63] B. Fröhler, C. Anthes, F. Pointecker, J. Friedl, D. Schwajda, A. Riegler, S. Tripathi, C. Holzmann, M. Brunner, H. Jodlbauer, H.-C. Jetter, and C. Heinzl. 2022. A Survey on Cross-Virtuality Analytics. *Computer Graphics Forum* 41, 1 (2022), 465–494. doi:10.1111/cgf.14447

[64] A. Fuhrmann, H. Loffelmann, D. Schmalstieg, and M. Gervautz. 1998. Collaborative visualization in augmented reality. *IEEE Computer Graphics and Applications* 18, 4 (1998), 54–59. doi:10.1109/38.689665

[65] Ana María Cárdenas Gasca, Payton Croskey, Tobias Höllerer, Kai M Thaler, Emilia Yang, and Jennifer Jacobs. 2025. Lessons from AR Memorialization: Artists' and Activists' Approaches to Responsible AR Development. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 40–58. doi:10.1145/3715336.3735815

[66] Ceenu George, Philipp Janssen, David Heuss, and Florian Alt. 2019. Should I Interrupt or Not?: Understanding Interruptions in Head-Mounted Display Settings. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 497–510. doi:10.1145/3322276.3322363

[67] Ceenu George, An Ngo Tien, and Heinrich Hussmann. 2020. Seamless, Bi-directional Transitions along the Reality-Virtuality Continuum: A Conceptualization and Prototype Exploration. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 412–424. doi:10.1109/ismar50242.2020.00067

[68] VanJessica Gladney, Breanna Moore, and Kathleen Brown. 2024. Penn & Slavery Project's Augmented Reality Tour: Augmenting a Campus to Reveal a Hidden History. Position paper for the CHI 2024 workshop With or Without Permission: Site-Specific Augmented Reality for Social Justice. doi:10.48550/arXiv.2404.14379

[69] Mark Graham, Matthew Zook, and Andrew Boulton. 2013. Augmented reality in urban places: contested content and the duplicity of code. *Transactions of the Institute of British Geographers* 38, 3 (2013), 464–479. doi:10.1111/j.1475-5661.2012.00539.x

[70] Jeronimo Gustavo Grandi, Henrique Galvan Debarba, and Anderson Maciel. 2019. Characterizing Asymmetric Collaborative Interactions in Virtual and Augmented Realities. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, New York, NY, USA, 127–135. doi:10.1109/VR.2019.8798080

[71] Saul Greenberg, Sebastian Boring, Jo Vermeulen, and Jakub Dostal. 2014. Dark Patterns in Proxemic Interactions: A Critical Perspective. In *Proceedings of the 2014 Conference on Designing Interactive Systems*. Association for Computing Machinery, New York, NY, USA, 523–532. doi:10.1145/2598510.2598541

[72] Jens Grubert, Tobias Langlotz, Stefanie Zollmann, and Holger Regenbrecht. 2017. Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 23, 6 (2017), 1706–1724. doi:10.1109/tvcg.2016.2543720

[73] Uwe Gruenefeld, Jonas Auda, Florian Mathis, Stefan Schneegass, Mohamed Khamis, Jan Gugenheimer, and Sven Mayer. 2022. VRception: Rapid Prototyping of Cross-Reality Systems in Virtual Reality. In *CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3491102.3501821

[74] Uwe Gruenefeld, Torge Wolff, Niklas Diekmann, Marion Koelle, and Wilko Heuten. 2019. ChalkboARd: exploring augmented reality for public displays. In *Proceedings of the 8th ACM International Symposium on Pervasive Displays*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3321335.3324929

[75] Jan Gugenheimer, Christian Mai, Mark McGill, Julie Williamson, Frank Steinicke, and Ken Perlin. 2019. Challenges Using Head-Mounted Displays in Shared and Social Spaces. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3290607.3299028

[76] Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. 2017. ShareVR: Enabling Co-Located Experiences for Virtual Reality between HMD and Non-HMD Users. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 4021–4033. doi:10.1145/3025453.3025683

[77] Jan Gugenheimer, Evgeny Stemasov, Harpreet Sareen, and Enrico Rukzio. 2018. FaceDisplay: Towards Asymmetric Multi-User Interaction for Nomadic Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173628

[78] Kim Halskov, Graham Dove, and Aron Fischel. 2021. Constructing a Design Space from a Collection of Design Examples. *She Ji: The Journal of Design, Economics, and Innovation* 7, 3 (2021), 462–484. doi:10.1016/j.sheji.2021.07.001

[79] Rodney Harrison, Alexander Rose, and Colin Sterling. 2025. Ghosts of Solid Air: Contested Heritage and Augmented Reality in Public Space. *International Journal of Heritage Studies* 31, 4 (2025), 425–443. doi:10.1080/13527258.2024.2437367

[80] Jeremy Hartmann, Yen-Ting Yeh, and Daniel Vogel. 2020. AAR: Augmenting a Wearable Augmented Reality Display with an Actuated Head-Mounted Projector. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 445–458. doi:10.1145/3379337.3415849

[81] Alexander Heinrich, Matthias Hollick, Thomas Schneider, Milan Stute, and Christian Weinert. 2021. PrivateDrop: Practical Privacy-Preserving Authentication for Apple AirDrop. In *30th USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 3577–3594. https://www.usenix.org/conference/usenixsecurity21/presentation/heinrich

[82] Jaylin Herskovitz, Yi Fei Cheng, Anhong Guo, Alanson P. Sample, and Michael Nebeling. 2022. XSpace: An Augmented Reality Toolkit for Enabling Spatially-Aware Distributed Collaboration. *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (2022), 277–302. doi:10.1145/3567721

[83] Ken Hinckley. 2003. Synchronous gestures for multiple persons and computers. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*. Association for Computing Machinery, New York, NY, USA, 149–158. doi:10.1145/964696.964713

[84] Lars Erik Holmquist, Friedemann Mattern, Bernt Schiele, Petteri Alahuhta, Michael Beigl, and Hans-W. Gellersen. 2001. Smart-Its Friends: A Technique for Users to Easily Establish Connections between Smart Artefacts. In *Ubicomp 2001: Ubiquitous Computing*. Springer Berlin Heidelberg, Cham, Switzerland, 116–122. doi:10.1007/3-540-45427-6_10

[85] Jubayer Hossain, M. Rasel Mahmud, Honghui Xu, and Chenyu Wang. 2026. Blockchain-Based Spatial Registry and Management in Multi-User Mixed Reality. In *2026 35th International Conference on Computer Communications and Networks (ICCCN)*. IEEE, New York, NY, USA, 1–9. doi:10.1109/icccn69946.2026.11662632

[86] Botao Hu, Yuchen Zhang, Sizheng Hao, and Yilan Tao. 2023. InstantCopresence: A Spatial Anchor Sharing Methodology for Co-located Multiplayer Handheld and Headworn AR. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, New York, NY, USA, 762–763. doi:10.1109/ismar-adjunct60411.2023.00165

[87] Botao Amber Hu. 2025. Autonomous Realities: A Journey into Protocolizing Digital Object Permanence in a Future of Many Mixed Realities. In *Proceedings of the sixth decennial Aarhus conference: Computing X Crisis*. Association for Computing Machinery, New York, NY, USA, 290–302. doi:10.1145/3744169.3744197

[88] Botao Amber Hu, Samuel Chua, and Helena Rong. 2026. Protocol Futuring: Speculating Second-Order Dynamics of Protocols in Sociotechnical Infrastructural Futures. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3772318.3790527

[89] Botao Amber Hu, Rem RunGu Lin, Yilan Elan Tao, Samuli Laato, and Yue Li. 2025. Towards Immersive Mixed Reality Street Play: Understanding Co-located Bodily Play with See-through Head-Mounted Displays in Public Spaces. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–46. doi:10.1145/3757679

[90] Botao Amber Hu, Yilan Elan Tao, Rem RunGu Lin, and Yue Li. 2024. On Intent Inclusivity in Spontaneous Cross Realities. In *2024 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, New York, NY, USA, 181–185. doi:10.1109/ismar-adjunct64951.2024.00046

[91] Botao Amber Hu, Yilan Elan Tao, Bernhard Riecke, and Yue Li. 2026. Allow Me Into Your Dream: A Handshake-and-Pull Protocol for Sharing Mixed Realities in Spontaneous Encounters. arXiv. doi:10.48550/arxiv.2604.19423

[92] Vincent C. Hu, David Ferraiolo, Rick Kuhn, Adam Schnitzer, Kenneth Sandlin, Robert Miller, and Karen Scarfone. 2014. *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*. Technical Report. National Institute of Standards and Technology. doi:10.6028/nist.sp.800-162 Includes updates as of 2 August 2019.

[93] Andrew Hudson-Smith. 2026. Spatial Computing and the Phygital City: Towards Digital Habitation. *Urban Planning* 11, Article 12412 (2026). doi:10.17645/up.12412

[94] Seungyeon Huh, Shapna Muralidharan, Heedong Ko, and Byounghyun Yoo. 2019. XR Collaboration Architecture based on Decentralized Web. In *The 24th International Conference on 3D Web Technology*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3329714.3338137

[95] Ke Huo, Tianyi Wang, Luis Paredes, Ana M. Villanueva, Yuanzhi Cao, and Karthik Ramani. 2018. SynchronizAR: Instant Synchronization for Spontaneous and Spatial Collaborations in Augmented Reality. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 19–30. doi:10.1145/3242587.3242595

[96] Jonna Häkkilä, Ashley Colley, Paula Roinesalo, and Jani Väyrynen. 2017. Clothing integrated augmented reality markers. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia*. Association for Computing Machinery, New York, NY, USA, 113–121. doi:10.1145/3152832.3152850

[97] Shahram Izadi, Harry Brignull, Tom Rodden, Yvonne Rogers, and Mia Underwood. 2003. Dynamo: A Public Interactive Surface Supporting the Cooperative Sharing and Exchange of Media. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 159–168. doi:10.1145/964696.964714

[98] Pascal Jansen, Fabian Fischbach, Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. 2020. ShARe: Enabling Co-Located Asymmetric Multi-User Interaction for Augmented Reality Head-Mounted Displays. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 459–471. doi:10.1145/3379337.3415843

[99] Lucy Jiang. 2024. ARtivism: AR-Enabled Accessible Public Art and Advocacy. Position paper, arXiv:2404.13285. doi:10.48550/arXiv.2404.13285

[100] Ross Johnstone, Neil McDonnell, and Julie R. Williamson. 2022. When Virtuality Surpasses Reality: Possible Futures of Ubiquitous XR. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3491101.3516396

[101] Alan Joy and Joseph O'Hagan. 2025. Acceptance of an Augmented Society: Initial Explorations into the Acceptability of Augmenting Real World Locations. arXiv. doi:10.48550/arxiv.2502.06378

[102] Pavel Karpashevich, Eva Hornecker, Nana Kesewaa Dankwa, Mohamed Hanafy, and Julian Fietkau. 2016. Blurring boundaries between everyday life and pervasive gaming: an interview study of ingress. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia*. Association for Computing Machinery, New York, NY, USA, 217–228. doi:10.1145/3012709.3012716

[103] Ryan M. Kelly, Hasan Shahid Ferdous, Niels Wouters, and Frank Vetere. 2019. Can Mobile Augmented Reality Stimulate a Honeypot Effect?: Observations from Santa's Lil Helper. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300515

[104] Talha Khan and David Lindlbauer. 2024. Don't Block My Stuff: Fostering Personal Object Awareness in Multi-user Mixed Reality Environments. *Proceedings of the ACM on Human-Computer Interaction* 8, ISS (2024), 20–43. doi:10.1145/3698126

[105] Talha Khan, Abigail Zimmerman, Edward Andrews, David Lindlbauer, and Jacob Biehl. 2025. Shadows of Reality: Enhancing Bystander Awareness of Mixed Reality Interfaces. In *Proceedings of the 2025 ACM Symposium on Spatial User Interaction*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3694907.3765920

[106] Raphael Kim, Conor Linehan, and Larissa Pschetz. 2022. Navigating Imaginaries of DNA-Based Digital Data Storage. In *CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 616, 15 pages. doi:10.1145/3491102.3501911

[107] Yoonsang Kim, Sanket Goutam, Amir Rahmati, and Arie E. Kaufman. 2023. Erebus: Access Control for Augmented Reality Systems. USENIX Security. https://www.usenix.org/conference/usenixsecurity23/presentation/kim-yoonsang

[108] Afshan Kirmani and Rowanne Fleck. 2014. "People don't bump": sharing around mobile phones in close proximity. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services*. Association for Computing Machinery, New York, NY, USA, 549–554. doi:10.1145/2628363.2634231

[109] Kiyoshi Kiyokawa, Haruo Takemura, and Naokazu Yokoya. 1999. Seamless Switching of Shared Virtual and Augmented Environments for Collaboration Support. *Transactions of the Virtual Reality Society of Japan* 4, 4 (1999), 657–664. doi:10.18974/tvrsj.4.4_657

[110] Jarrod Knibbe, Jonas Schjerlund, Mathias Petraeus, and Kasper Hornbæk. 2018. The Dream is Collapsing: The Experience of Exiting VR. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3174057

[111] Marion Koelle, Swamy Ananthanarayan, Simon Czupalla, Wilko Heuten, and Susanne Boll. 2018. Your smart glasses' camera bothers me!: exploring opt-in and opt-out gestures for privacy mediation. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction*. Association for Computing Machinery, New York, NY, USA, 473–481. doi:10.1145/3240167.3240174

[112] Marion Koelle, Torben Wallbaum, Wilko Heuten, and Susanne Boll. 2019. Evaluating a Wearable Camera's Social Acceptability In-the-Wild. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3312837

[113] N. Kohtake, J. Rekimoto, and Y. Anzai. 1999. InfoStick: An Interaction Device for Inter-Appliance Computing. In *Handheld and Ubiquitous Computing*. Springer Berlin Heidelberg, Cham, Switzerland, 246–258. doi:10.1007/3-540-48157-5_23

[114] Boriana Koleva, Holger Schnädelbach, Steve Benford, and Chris Greenhalgh. 2000. Traversable interfaces between real and virtual worlds. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 233–240. doi:10.1145/332040.332437

[115] Gerd Kortuem, Zary Segall, and Thaddeus G. Cowan Thompson. 1999. Close Encounters: Supporting Mobile Collaboration through Interchange of User Profiles. In *Handheld and Ubiquitous Computing*. Springer Berlin Heidelberg, Cham, Switzerland, 171–185. doi:10.1007/3-540-48157-5_17

[116] Olya Kudina and Peter-Paul Verbeek. 2019. Ethics from Within: Google Glass, the Collingridge Dilemma, and the Mediated Value of Privacy. *Science, Technology, & Human Values* 44, 2 (2019), 291–314. doi:10.1177/0162243918793711

[117] Yoshiki Kudo, Anthony Tang, Kazuyuki Fujita, Isamu Endo, Kazuki Takashima, and Yoshifumi Kitamura. 2021. Towards Balancing VR Immersion and Bystander Awareness. *Proceedings of the ACM on Human-Computer Interaction* 5, ISS (2021), 1–22. doi:10.1145/3486950

[118] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Bjoern Hartmann. 2020. TransceiVR: Bridging Asymmetrical Communication Between VR Users and External Collaborators. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 182–195. doi:10.1145/3379337.3415827

[119] 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*. Association for Computing Machinery, New York, NY, USA, 81–86. doi:10.1145/3301293.3302362

[120] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2017. Securing Augmented Reality Output. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, New York, NY, USA, 320–337. doi:10.1109/sp.2017.13

[121] 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)*. IEEE, New York, NY, USA, 392–408. doi:10.1109/sp.2018.00051

[122] DoYoung Lee, Youryang Lee, Yonghwan Shin, and Ian Oakley. 2018. Designing Socially Acceptable Hand-to-Face Input. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 711–723. doi:10.1145/3242587.3242642

[123] Lik-Hang Lee and Zijun Lin. 2024. Danger, Nuisance, Disregard: Analyzing User-Generated Videos for Augmented Reality Gameplay on Hand-held Devices. *Proceedings of the ACM on Human-Computer Interaction* 8, CHI PLAY (2024), 1–33. doi:10.1145/3677063

[124] Tony Liao. 2019. Future directions for mobile augmented reality research: Understanding relationships between augmented reality users, nonusers, content, devices, and industry. *Mobile Media & Communication* 7, 1 (2019), 131–149. doi:10.1177/2050157918792438

[125] Nicolette Little and Tom Divon. 2024. The Unsolicited Algorithm: Unveiling Gendered Harms and (Non)Consent in Apple iOS Features. *Journal of Gender Studies* (2024). doi:10.1080/09589236.2024.2433683

[126] Andrés Lucero, Jussi Holopainen, and Tero Jokela. 2011. Pass-Them-Around: Collaborative Use of Mobile Phones for Photo Sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1787–1796. doi:10.1145/1978942.1979201

[127] Angella Mackey, Stephan Wensveen, Ron Wakkary, Annika Hupfeld, and Oscar Tomico. 2019. Wearing Digital Shimmers: A Fashion-Centric Approach to Wearable Technology. In *Proceedings of the 4th Biennial Research Through Design Conference (RTD 2019)*. https://www.semanticscholar.org/paper/76be61910b2288df9980294 3b0988bfa4d1a0aee

[128] Steve Mann. 2004. Continuous lifelong capture of personal experience with EyeTap. In *Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/1026653.1026654

[129] Nicolai Marquardt, Till Ballendat, Sebastian Boring, Saul Greenberg, and Ken Hinckley. 2012. Gradual Engagement: Facilitating Information Exchange between Digital Devices as a Function of Proximity. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces*. Association for Computing Machinery, New York, NY, USA, 31–40. doi:10.1145/2396636.2396642

[130] Danica Mast, Sanne I. de Vries, Joost Broekens, and Fons J. Verbeek. 2021. The Participant Journey Map: Understanding the Design of Interactive Augmented Play Spaces. *Frontiers in Computer Science* 3, Article 674132 (2021). doi:10.3389/fcomp.2021.674132

[131] Keiichi Matsuda. 2016. Hyper-Reality. Speculative film. https://vimeo.com/166807261

[132] Keiichi Matsuda. 2023. Agents. Short film, Liquid City for Niantic. https://www.linkedin.com/posts/keiichi-matsuda-914b5214_how-will-ai-agents-change-the-way-we-live-activity-7288609214720077824-nqFx Creator-posted film and transcript.

[133] Nathan Matsuda, Brian Wheelwright, Joel Hegland, and Douglas Lanman. 2021. Reverse Pass-Through VR. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Emerging Technologies*. Association for Computing Machinery, New York, NY, USA, 1–4. doi:10.1145/3450550.3465338

[134] Mark McGill, Jan Gugenheimer, and Euan Freeman. 2020. A Quest for Co-Located Mixed Reality: Aligning and Assessing SLAM Tracking for Same-Space Multi-User Experiences. In *26th ACM Symposium on Virtual Reality Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3385956.3418968

[135] Daniel Medeiros, Romane Dubus, Julie Williamson, Graham Wilson, Katharina Pöhlmann, and Mark McGill. 2023. Surveying the Social Comfort of Body, Device, and Environment-Based Augmented Reality Interactions in Confined Passenger Spaces Using Mixed Reality Composite Videos. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–25. doi:10.1145/3610923

[136] Metaverse Standards Forum. 2026. The Metaverse Standards Forum. https://metaverse-standards.org/ Living documentation; accessed 9 September 2026.

[137] China Miéville. 2025. *The City & The City*. Picador, London, UK. https://www.panmacmillan.com/authors/china-mieville/the-city-the-city/9781035060245 Paperback edition published 6 February 2025.

[138] John Miller, Elahe Soltanaghai, Raewyn Duvall, Jeff Chen, Vikram Bhat, Nuno Pereira, and Anthony Rowe. 2022. Cappella: Establishing Multi-User Augmented Reality Sessions Using Inertial Estimates and Peer-to-Peer Ranging. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, New York, NY, USA, 428–440. doi:10.1109/ipsn54338.2022.00041

[139] Mark Roman Miller, Hanseul Jun, Fernanda Herrera, Jacob Yu Villa, Greg Welch, and Jeremy N. Bailenson. 2019. Social interaction in augmented reality. *PLOS ONE* 14, 5 (2019), e0216290. doi:10.1371/journal.pone.0216290

[140] MoMAR. 2018. Hello, we're from the internet. Augmented-reality exhibition. https://momar.gallery/exhibitions/werefromtheinternet.html Exhibition project page; 2 March–10 May 2018.

[141] Eric Monteiro. 1998. Scaling Information Infrastructure: The Case of Next-Generation IP in the Internet. *The Information Society* 14, 3 (1998), 229–245. doi:10.1080/019722498128845

[142] Erica L. Neely. 2019. Augmented reality, augmented ethics: who has the right to augment a particular physical space? *Ethics and Information Technology* 21, 1 (2019), 11–18. doi:10.1007/s10676-018-9484-2

[143] Fabrizio Nevola, Tim Coles, and Cristina Mosconi. 2022. Hidden Florence revealed? Critical insights from the operation of an augmented reality app in a World Heritage City. *Journal of Heritage Tourism* 17, 4 (2022), 371–390. doi:10.1080/1743873x.2022.2036165

[144] Nels Numan and Anthony Steed. 2022. Exploring User Behaviour in Asymmetric Collaborative Mixed Reality. In *Proceedings of the 28th ACM Symposium on Virtual Reality Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–11. doi:10.1145/3562939.3565630

[145] Joseph O'Hagan, Jan Gugenheimer, Jolie Bonner, Florian Mathis, and Mark McGill. 2023. Augmenting People, Places & Media: The Societal Harms Posed by Everyday Augmented Reality, and the Case for Perceptual Human Rights. In *Proceedings of the 22nd International Conference on Mobile and Ubiquitous Multimedia*. Association for Computing Machinery, New York, NY, USA, 225–235. doi:10.1145/3626705.3627782

[146] Joseph O'Hagan, Jan Gugenheimer, Florian Mathis, Jolie Bonner, Richard Jones, and Mark McGill. 2024. A Viewpoint on the Societal Impact of Everyday Augmented Reality and the Need for Perceptual Human Rights. *IEEE Security & Privacy* 22, 1 (2024), 64–68. doi:10.1109/MSEC.2023.3333988

[147] 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. *Proceedings of the ACM on*

Interactive, Mobile, Wearable and Ubiquitous Technologies 6, 4 (2022), 1–35. doi:10.1145/3569501

[148] Joseph O'Hagan, Julie R. Williamson, and Mohamed Khamis. 2020. Bystander interruption of VR users. In *Proceedings of the 9TH ACM International Symposium on Pervasive Displays*. Association for Computing Machinery, New York, NY, USA, 19–27. doi:10.1145/3393712.3395339

[149] Joseph O'Hagan, Julie R. Williamson, Florian Mathis, Mohamed Khamis, and Mark McGill. 2023. Re-Evaluating VR User Awareness Needs During Bystander Interactions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3544548.3581018

[150] Open AR Cloud Association. 2026. Open Spatial Computing Platform. https://openarcloud.org/ Living documentation; accessed 9 September 2026.

[151] Open Geospatial Consortium. 2023. OGC GeoPose 1.0 Data Exchange Standard. https://docs.ogc.org/is/21-056r11/21-056r11.html OGC 21-056r11, version 1.0.0; published 8 September 2023.

[152] Nikita Menon K. P., Eldy S. Lazaro Vasquez, Hannah Curran, Sasha De Koninck, and Laura Devendorf. 2023. Weaving Augmented Reality Markers. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*. Association for Computing Machinery, New York, NY, USA, 310–314. doi:10.1145/3594739.3610786

[153] Susanna Paasovaara, Pradthana Jarusriboonchai, and Thomas Olsson. 2017. Understanding collocated social interaction between Pokémon GO players. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia*. Association for Computing Machinery, New York, NY, USA, 151–163. doi:10.1145/3152832.3152854

[154] Janne Paavilainen, Hannu Korhonen, Kati Alha, Jaakko Stenros, Elina Koskinen, and Frans Mäyrä. 2017. The Pokémon GO Experience: A Location-Based Augmented Reality Mobile Game Goes Mainstream. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2493–2498. doi:10.1145/3025453.3025871

[155] Leonardo Pavanatto, Verena Biener, Jennifer Chandran, Snehanjali Kalamkar, Feiyu Lu, John J. Dudley, Jinghui Hu, G. Nikki Ramirez, Per Ola Kristensson, Alexander Giovannelli, Luke Schlueter, Jörg Müller, Jens Grubert, and Doug A. Bowman. 2025. Working in Extended Reality in the Wild: Worker and Bystander Experiences of XR Virtual Displays in Public Real-World Settings. *IEEE Transactions on Visualization and Computer Graphics* 31, 10 (2025), 9104–9123. doi:10.1109/tvcg.2025.3589283

[156] Kushani Tharushika Perera, Tobias Langlotz, Nadia Pantidi, and Holger Regenbrecht. 2024. What you see is (not necessarily) what I see – Pervasive AR for Public Displays. In *Proceedings of the 36th Australasian Conference on Human-Computer Interaction*. Association for Computing Machinery, New York, NY, USA, 59–73. doi:10.1145/3726986.3727005

[157] Thammathip Piumsomboon, Arindam Dey, Barrett Ens, Gun Lee, and Mark Billinghurst. 2019. The Effects of Sharing Awareness Cues in Collaborative Mixed Reality. *Frontiers in Robotics and AI* 6, Article 5 (2019). doi:10.3389/frobt.2019.00005

[158] Thammathip Piumsomboon, Gavin Ong, Cameron Urban, Barrett Ens, Jack Topliss, Xiaoliang Bai, and Simon Hoermann. 2022. Ex-Cit XR: Expert-elicitation and validation of Extended Reality visualisation and interaction techniques for disengaging and transitioning users from immersive virtual environments. *Frontiers in Virtual Reality* 3, Article 943696 (2022). doi:10.3389/frvir.2022.943696

[159] Pixar Animation Studios. 2026. Introduction to USD. Universal Scene Description documentation. https://openusd.org/release/intro.html Official living documentation; accessed 9 September 2026.

[160] Fabian Pointecker, Judith Friedl, Daniel Schwajda, Hans-Christian Jetter, and Christoph Anthes. 2022. Bridging the Gap Across Realities: Visual Transitions Between Virtual and Augmented Reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 827–836. doi:10.1109/ismar55827.2022.00101

[161] Pokémon GO. n.d.. Pokémon GO. Official game website. https://pokemongo.com/en

[162] Lev Poretski, Ofer Arazy, Joel Lanir, and Oded Nov. 2021. Who Owns What? Psychological Ownership in Shared Augmented Reality. *International Journal of Human-Computer Studies* 150 (2021), 102611. doi:10.1016/j.ijhcs.2021.102611

[163] Lev Poretski, Joel Lanir, and Ofer Arazy. 2018. Normative Tensions in Shared Augmented Reality. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–22. doi:10.1145/3274411

[164] Katharina M. T. Pöhlmann, Graham Wilson, Laura Bajorunaite, Mark McGill, and Joseph O'Hagan. 2025. "Everyone Knows You're Watching P*rn": Reflecting on First-Hand User Accounts on the Use and Public Perception of in-Transit Passenger XR. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, New York, NY, USA, 560–570. doi:10.1109/ismar67309.2025.00066

[165] Luis Quintero, António Miguel Beleza Maciel Pinheiro Braga, Noak Petersson, and Uno Gh Fors. 2025. Transitional Portals for Participatory Co-Located Cross-Reality Experiences. In *Proceedings of the 2025 ACM International Conference on Interactive Media Experiences*. Association for Computing Machinery, New York, NY, USA, 368–371. doi:10.1145/3706370.3731708

[166] Adil Rahman, Wen Ying, Md Aashikur Rahman Azim, Michelle Annett, and Seongkook Heo. 2026. "It Feels Like I am Invited to Communicate": Mediating Ad-Hoc Bystander-VR User Interruptions Through Proactive Proxies. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3772318.3791936

[167] Shwetha Rajaram, Chen Chen, Franziska Roesner, and Michael Nebeling. 2023. Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3544548.3581089

[168] Shwetha Rajaram, Jiasi Chen, and Michael Nebeling. 2025. Privacy Equilibrium: Balancing Privacy Needs in Dynamic Multi-User Augmented Reality Scenarios. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–24. doi:10.1145/3746059.3747783

[169] Xukan Ran, Carter Slocum, Maria Gorlatova, and Jiasi Chen. 2019. ShareAR: Communication-Efficient Multi-User Mobile Augmented Reality. In *Proceedings of the 18th ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 109–116. doi:10.1145/3365609.3365867

[170] Johan Redström, Per Dahlberg, Peter Ljungstrand, and Lars Erik Holmquist. 2000. Designing for Local Interaction. In *Managing Interactions in Smart Environments*. Springer London, Cham, Switzerland, 227–238. doi:10.1007/978-1-4471-0743-9_19

[171] Stuart Reeves, Steve Benford, Claire O'Malley, and Mike Fraser. 2005. Designing the Spectator Experience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 741–750. doi:10.1145/1054972.1055074

[172] Stuart Reeves, Mike Fraser, Holger Schnädelbach, Claire O'Malley, and Steve Benford. 2005. Engaging Augmented Reality in Public Places. CHI. https://people.cs.nott.ac.uk/pszsr/files/reeves-2005-engaging-augmented-reality.pdf

[173] Holger Regenbrecht, Alistair Knott, Jennifer Ferreira, and Nadia Pantidi. 2024. To See and be Seen—Perceived Ethics and Acceptability of Pervasive Augmented Reality. *IEEE Access* 12 (2024), 32618–32636. doi:10.1109/access.2024.3366228

[174] Jun Rekimoto. 1996. TransVision: A Hand-held Augmented Reality System for Collaborative Design. In *Virtual Systems and Multi-Media (VSMM)*. https://lab.rekimoto.org/projects/transvision/

[175] Jun Rekimoto. 1997. Pick-and-Drop: A Direct Manipulation Technique for Multiple Computer Environments. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 31–39. doi:10.1145/263407.263505

[176] J. Rekimoto, Y. Ayatsuka, and K. Hayashi. 1998. Augment-able reality: situated communication through physical and digital spaces. In *Digest of Papers. Second International Symposium on Wearable Computers (Cat. No.98EX215)*. IEEE Comput. Soc, New York, NY, USA, 68–75. doi:10.1109/ISWC.1998.729531

[177] Jun Rekimoto, Yuji Ayatsuka, and Michimune Kohno. 2003. SyncTap: An Interaction Technique for Mobile Networking. In *Human-Computer Interaction with Mobile Devices and Services*. Springer, Berlin, Heidelberg, 104–115. doi:10.1007/978-3-540-45233-1_9

[178] Jun Rekimoto and Katashi Nagao. 1995. The world through the computer: computer augmented interaction with real world environments. In *Proceedings of the 8th annual ACM symposium on User interface and software technology*. Association for Computing Machinery, New York, NY, USA, 29–36. doi:10.1145/215585.215639

[179] Julie Rico and Stephen Brewster. 2010. Usable gestures for mobile interfaces: evaluating social acceptability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 887–896. doi:10.1145/1753326.1753458

[180] Jan Ole Rixen, Mark Colley, Ali Askari, Jan Gugenheimer, and Enrico Rukzio. 2022. Consent in the Age of AR: Investigating The Comfort With Displaying Personal Information in Augmented Reality. In *CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3491102.3502140

[181] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (2014), 88–96. doi:10.1145/2580723.2580730

[182] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2019. Secure Multi-User Content Sharing for Augmented Reality Applications. USENIX Security. https://www.usenix.org/conference/usenixsecurity19/presentation/ruth

[183] Janíce Tisha Samuels and Kelvin Ramirez. 2021. Building a Virtuous Cycle of Activism Using Art and Augmented Reality: A Community of Practice-Based Project. In *Augmented and Mixed Reality for Communities*. CRC Press, Boca Raton, FL, USA, 213–242. doi:10.1201/9781003052838-14

[184] Dieter Schmalstieg, Anton Fuhrmann, Gerd Hesina, Zsolt Szalavári, L. Miguel Encarnação, Michael Gervautz, and Werner Purgathofer. 2002. The Studierstube Augmented Reality Project. *Presence: Teleoperators and Virtual Environments* 11, 1 (2002), 33–54. doi:10.1162/105474602317343640

[185] D. Schmalstieg and G. Hesina. 2002. Distributed applications for collaborative augmented reality. In *Proceedings IEEE Virtual Reality 2002*. IEEE Comput. Soc, New York, NY, USA, 59–66. doi:10.1109/VR.2002.996505

[186] Leonard Schmidt and Enes Yigitbas. 2024. Development and Usability Evaluation of Transitional Cross-Reality Interfaces. *Proceedings of the ACM on Human-Computer Interaction* 8, EICS (2024), 1–32. doi:10.1145/3664637

[187] Jan-Henrik Schröder, Daniel Schacht, Niklas Peper, Anita Marie Hamurculu, and Hans-Christian Jetter. 2023. Collaborating Across Realities: Analytical Lenses for Understanding Dyadic Collaboration in Transitional Interfaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3544548.3580879

[188] Mickael Sereno, Xiyao Wang, Lonni Besancon, Michael J Mcguffin, and Tobias Isenberg. 2020. Collaborative Work in Augmented Reality: A Survey. *IEEE Transactions on Visualization and Computer Graphics* (2020), 1–1. doi:10.1109/tvcg.2020.3032761

[189] Robert W. Shirey. 2007. *Internet Security Glossary, Version 2*. RFC 4949. RFC Editor. doi:10.17487/RFC4949

[190] Rafael M. L. Silva, Erica Principe Cruz, Daniela K. Rosner, Dayton Kelly, Andrés Monroy-Hernández, and Fannie Liu. 2022. Understanding AR Activism: An Interview Study with Creators of Augmented Reality Experiences for Social Change. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3491102.3517605

[191] Rafael M. L. Silva, Ana María Cárdenas Gasca, Joshua A. Fisher, Erica Principe Cruz, Cinthya Jauregui, Amy Lueck, Fannie Liu, Andrés Monroy-Hernández, and Kai Lukoff. 2024. With or Without Permission: Site-Specific Augmented Reality for Social Justice. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3613905.3636283

[192] Carter Slocum, Yicheng Zhang, Erfan Shayegani, Pedram Zaree, Nael Abu-Ghazaleh, and Jiasi Chen. 2024. That Doesn't Go There: Attacks on Shared State in Multi-User Augmented Reality Applications. USENIX Security. doi:10.48550/arxiv.2308.09146

[193] Ivo Sluganovic, Mihael Liskij, Ante Derek, and Ivan Martinovic. 2020. Tap-Pair: Using Spatial Secrets for Single-Tap Device Pairing of Augmented Reality Headsets. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy*. Association for Computing Machinery,

New York, NY, USA, 61–72. doi:10.1145/3374664.3375740

[194] Ivo Sluganovic, Matej Serbec, Ante Derek, and Ivan Martinovic. 2017. HoloPair: Securing Shared Augmented Reality Using Microsoft HoloLens. In *Proceedings of the 33rd Annual Computer Security Applications Conference*. Association for Computing Machinery, New York, NY, USA, 250–261. doi:10.1145/3134600.3134625

[195] Hannah Snyder. 2019. Literature review as a research methodology: An overview and guidelines. *Journal of Business Research* 104 (2019), 333–339. doi:10.1016/j.jbusres.2019.07.039

[196] Gábor Sörös, James Jackson, Mikel Salazar, Christine Perey, Ali C. Hantal, and Alina Kadlubsky. 2025. Building Interoperable Location-based Augmented Reality with the Open AR Cloud. In *Companion of the 2025 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. Association for Computing Machinery, New York, NY, USA, 1689–1692. doi:10.1145/3714394.3750558

[197] Tobin South, Samuele Marro, Thomas Hardjono, Robert Mahari, Cedric Deslandes Whitney, Dazza Greenwood, Alan Chan, and Alex Pentland. 2025. Authenticated Delegation and Authorized AI Agents. arXiv preprint arXiv:2501.09674. doi:10.48550/arxiv.2501.09674

[198] Helen Stefanidi, Jan-Hendrik Sünderkamp, Markus Tatzgern, Alina Itzlinger, and Alexander Meschtscherjakov. 2025. You're making things AR-kward: Exploring Augmented Reality In-the-Wild MHCI041. *Proceedings of the ACM on Human-Computer Interaction* 9, 5 (2025), 1–24. doi:10.1145/3743740

[199] Jannis Strecker, Simon Mayer, and Kenan Bektaş. 2025. Towards Societally Beneficial Personalized Realities: A Conceptual Foundation for Responsible Ubiquitous Personalization Systems. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 1792–1814. doi:10.1145/3715336.3735709

[200] Milan Stute, Sashank Narain, Alex Mariotto, Alexander Heinrich, David Kreitschmann, Guevara Noubir, and Matthias Hollick. 2019. A Billion Open Interfaces for Eve and Mallory: MitM, DoS, and Tracking Attacks on iOS and macOS Through Apple Wireless Direct Link. In *28th USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 37–54. https://www.usenix.org/conference/usenixsecurity19/presentation/stute

[201] Z. Szalavári, D. Schmalstieg, A. Fuhrmann, and M. Gervautz. 1998. Studierstube: An Environment for Collaboration in Augmented Reality. *Virtual Reality* 3, 1 (1998), 37–48. doi:10.1007/BF01409796

[202] The Khronos Group. 2026. glTF: Runtime 3D Asset Delivery. https://www.khronos.org/gltf/ Living documentation; accessed 9 September 2026.

[203] The Khronos Group. 2026. OpenXR. https://www.khronos.org/openxr/ Living documentation; accessed 9 September 2026.

[204] Tram Thi Minh Tran, Shane Brown, Oliver Weidlich, Soojeong Yoo, and Callum Parker. 2025. Wearable AR in Everyday Contexts: Insights from a Digital Ethnography of YouTube Videos. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713572

[205] Andrea C. Tricco, Erin Lillie, Wasifa Zarin, Kelly K. O'Brien, Heather Colquhoun, Danielle Levac, David Moher, Micah D. J. Peters, Tanya Horsley, Laura Weeks, Susanne Hempel, Elie A. Akl, Christine Chang, Jessie McGowan, Lesley Stewart, Lisa Hartling, Adrian Aldcroft, Michael G. Wilson, Chantelle Garritty, Simon Lewin, Christina M. Godfrey, Marilyn T. Macdonald, Etienne V. Langlois, Karla Soares-Weiser, Jo Moriarty, Tammy Clifford, Özge Tunçalp, and Sharon E. Straus. 2018. PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. *Annals of Internal Medicine* 169, 7 (2018), 467–473. doi:10.7326/M18-0850

[206] Cody Turner. 2022. Augmented Reality, Augmented Epistemology, and the Real-World Web. *Philosophy & Technology* 35, 1, Article 19 (2022). doi:10.1007/s13347-022-00496-5

[207] Diego Vaquero-Melchor, Ana M. Bernardos, and Luca Bergesio. 2020. SARA: A Microservice-Based Architecture for Cross-Platform Collaborative Augmented Reality. *Applied Sciences* 10, 6 (2020), 2074. doi:10.3390/app10062074

[208] Sander Veenhof. 2024. Augmented Reality Gaming in Public Spaces: Dealing with Non-player Characters or Having a Shared Experience? HCI International. doi:10.1007/978-3-031-54475-0_17

[209] Vernor Vinge. 2006. *Rainbows End*. Tor Books, New York. https://us.macmillan.com/books/9780765308511/rainbowsend

[210] Julius von Willich, Frank Nelles, Wen-Jie Tseng, Jan Gugenheimer, Sebastian Günther, and Max Mühlhäuser. 2025. A Qualitative Investigation of User Transitions and Frictions in Cross-Reality Applications. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713921

[211] Chiu-Hsuan Wang, Chia-En Tsai, Seraphina Yong, and Liwei Chan. 2020. Slice of Light: Transparent and Integrative Transition Among Realities in a Multi-HMD-User Environment. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 805–817. doi:10.1145/3379337.3415868

[212] Chiu-Hsuan Wang, Seraphina Yong, Hsin-Yu Chen, Yuan-Syun Ye, and Liwei Chan. 2020. HMD Light: Sharing In-VR Experience via Head-Mounted Projector for Asymmetric Interaction. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 472–486. doi:10.1145/3379337.3415847

[213] Yuntao Wang, Zhou Su, Ning Zhang, Rui Xing, Dongxiao Liu, Tom H. Luan, and Xuemin Shen. 2023. A Survey on Metaverse: Fundamentals, Security, and Privacy. *IEEE Communications Surveys & Tutorials* 25, 1 (2023), 319–352. doi:10.1109/COMST.2022.3202047

[214] Sandro Weber, Linda Rudolph, Sven Liedtke, Christian Eichhorn, Daniel Dyrda, David A. Plecher, and Gudrun Klinker. 2022. Frameworks Enabling Ubiquitous Mixed Reality Applications Across Dynamically Adaptable Device Configurations. *Frontiers in Virtual Reality* 3, Article 765959 (2022). doi:10.3389/frvir.2022.765959

[215] Leonardo Werner, Philip Brey, and Adam Henschke. 2025. Augmented reality and ethics: key issues. *Virtual Reality* 29, 3, Article 122 (2025). doi:10.1007/s10055-025-01172-7

[216] Maximiliane Windl, Petra Zsofia Laboda, and Sven Mayer. 2025. Designing Effective Consent Mechanisms for Spontaneous Interactions in Augmented Reality. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713519

[217] World Wide Web Consortium. 2026. WebXR Device API. https://www.w3.org/TR/2026/CRD-webxr-20260609/ W3C Candidate Recommendation Draft, 9 June 2026; accessed 9 September 2026.

[218] Niels Wouters, John Downs, Mitchell Harrop, Travis Cox, Eduardo Oliveira, Sarah Webber, Frank Vetere, and Andrew Vande Moere. 2016. Uncovering the Honeypot Effect: How Audiences Engage with Public Interactive Systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. Association for Computing Machinery, New York, NY, USA, 5–16. doi:10.1145/2901790.2901796

[219] David J. Wu, Ankur Taly, Asim Shankar, and Dan Boneh. 2016. Privacy, Discovery, and Authentication for the Internet of Things. In *Computer Security – ESORICS 2016*. Springer, Cham, Switzerland, 301–319. doi:10.1007/978-3-319-45741-3_16

[220] Sixuan Wu, Jiannan Li, Maurício Sousa, and Tovi Grossman. 2023. Investigating Guardian Awareness Techniques to Promote Safety in Virtual Reality. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, New York, NY, USA, 631–640. doi:10.1109/VR55154.2023.00078

[221] Shihui Xu, Bo Yang, Boyang Liu, Kelvin Cheng, Soh Masuko, and Jiro Tanaka. 2019. Sharing Augmented Reality Experience Between HMD and Non-HMD User. HCII. doi:10.1007/978-3-030-22649-7_16

[222] Tamlyn Young and Mark T. Marshall. 2023. An Investigation of the Use of Augmented Reality in Public Art. *Multimodal Technologies and Interaction* 7, 9, Article 89 (2023), 89 pages. doi:10.3390/mti7090089

[223] Shaowei Zhu, Hyo Jin Kim, Maurizio Monge, G. Edward Suh, Armin Alaghi, Brandon Reagen, and Vincent C. S. Lee. 2022. Verifiable Access Control for Augmented Reality Localization and Mapping. PoPETs/arXiv. doi:10.48550/arxiv.2203.13308

A Protocol Futuring Portfolio

Table A1 presents the **ten prepared scenarios, G1–G10**, with their inspirations, team prompts, and five blue/red pairs.

This appendix is illustrative, not evidential. G1's five pairs reconstruct the author-facilitated AR-fashion account described in Section 4.2. The dialogues shown for G2–G10 are generative-AI expansions of the prepared prompts: they illustrate the kind of argument each scenario invites, and they do not report what any team said. **The analysis in Section 4 rests on the recorded team exchanges, not on these dialogues.** Four of these scenarios (G7–G10) were prepared but not taken up by any team, so no team recorded an exchange for them.

Table A1. Protocol Futuring portfolio: scenarios G1–G10 with illustrated starting situations and cited inspirations (left), with team prompts and ten-step worked explorations (right).

Scenario	Ten steps
G1. AI Fashion  **Scenario.** A person wears a digital outfit that appears through AR glasses. A friend can see it and stops to admire it, but a passerby sees only the friend's reaction. The wearer would like interested strangers to enjoy the outfit too. Making it visible raises a practical question: how can something published in one person's reality enter another person's view? Inspired by Hyper-Reality [131].	**Blue goal.** Help a willing passerby discover and see the outfit. **Red question.** How do they notice it? Who decides how it appears? **B1.** I want passersby to see my augmented outfit. The clothes express my style, and I would like people around me to appreciate them. **R1.** But others cannot see the outfit by default. Where would they find it, and through which layer could their devices display it? **B2.** We could create a public layer for AR fashion. Wearers could publish their outfits there for other people to encounter. **R2.** Why must passersby download your garment? Who consented to receiving it? What if its code is malicious or its resource package is very large? **B3.** We could set rules and review all garment code before upload, checking what enters the public layer before it reaches viewers. **R3.** Who controls the layer and reviews the code: government, a fashion company, or the app provider? What incentive supports that work? **B4.** A private company could operate the public AR cloud, like an online discussion space. Running the service would provide an incentive to govern its content. **R4.** If several companies operate separate public AR clouds, can I wear the same outfit across them? How would those services become compatible? **B5.** We could adopt a common garment format so that a given outfit retains its appearance when crossing between platforms and clouds. **R5.** Will the format change as devices improve? What happens when older devices cannot support the new version?
G2. Street Magic  **Scenario.** In a public square, people wearing AR glasses dodge and cast spells at one another. An approaching passerby sees their movements but none of the magic, and would like to understand or try the game. Other people simply need to cross the square. Opening this activity to newcomers must also account for different devices and people who never choose to play. Inspired by dueling scenarios in Harry Potter.	**Blue goal.** Let a passerby understand, watch or try the game. **Red question.** Does watching count as joining? What about other devices or refusal? **B1.** A screen beside the game could show spells and explain the rules. Passersby could understand the activity before deciding whether to try it. **R1.** If I approach that screen, am I now a player or target? Wanting to watch does not mean agreeing to be attacked. **B2.** Watching would remain untracked and separate from playing. Only someone who explicitly accepts a temporary player role could become a target. **R2.** My borrowed glasses show spells without team colours or attack ranges. How can I understand the player role I have just accepted? **B3.** Before joining, we could preview the cues each device supports and explain leaving. If essential cues are missing, we could offer a noncombat trial. **R3.** That lets me try something, but people with older glasses remain spectators while others enjoy the real game. Is that meaningful participation? **B4.** The group could choose an unranked round using only spells everyone's devices display. Players would agree together to this simpler version. **R4.** The players understand one another, but their dodging blocks someone pushing a pram. That person has not agreed to negotiate with the game. **B5.** We could mark a play area and keep a passage clear. A named organizer would pause play whenever it spills into that passage. **R5.** Who takes responsibility when friends start an improvised game? If the organizer leaves, does play stop, or does the burden return to passersby?

Table A1 continued: Protocol Futuring portfolio

Scenario	Ten steps
G3. Shop with Agents  **Scenario.** A shopper enters a shop with a personal AI assistant. The shop offers its own assistants for product advice and checkout, which could exchange information with the shopper's assistant. The shopper wants useful help without handling every detail. However, entering the shop or connecting to its network does not clarify which assistants are genuine or what they may agree to. Inspired by the shop-agent encounter in Agents [132].	**Blue goal.** Let your agent read useful information and buy within your authority. **Red question.** Does Wi-Fi imply entry? How far can your agent agree for you? **B1.** My assistant could ask shop assistants about suitable products and compare their answers with my preferences, saving me a search through every shelf. **R1.** Any nearby assistant might claim to represent the shop. How would yours distinguish a genuine sales assistant from someone impersonating one? **B2.** A fixed shop terminal could list recognized assistants and their roles. My assistant would check that list before exchanging information with them. **R2.** A recognized checkout assistant asks for a deposit. Knowing it represents the shop does not mean I authorized my assistant to spend money. **B3.** We could separate permission to read, inquire, reserve, and pay. My assistant would request confirmation before taking on a new cost or commitment. **R3.** If every ordinary purchase interrupts me for confirmation, what work has the assistant saved? I wanted help with routine decisions too. **B4.** I could approve product categories and spending limits in advance, excluding subscriptions and renewals. The assistant could complete purchases within those boundaries. **R4.** An affordable order ignores my dietary requirements. Who fixes the mistake when both assistants insist they followed their instructions? **B5.** The order could retain my conditions and the shop's offer. Shops accepting automatic purchases would agree to returns for orders violating those conditions. **R5.** A small shop may lack resources for that process. Would your assistant stop considering it, making convenience depend on excluding less resourced sellers?
G4. AR Qupid  **Scenario.** Two people sit in the same café without noticing each other. Their personal assistants could identify shared interests and suggest an introduction. Such help might create a welcome encounter, but discovering a possible match already involves information about both people. Someone may want company without publicly advertising that wish, or may prefer to remain undiscovered altogether. Inspired by the Q-pid introduction scene in Agents [132].	**Blue goal.** Help agents introduce people who both want to connect. **Red question.** On whose behalf can it act? What if someone wants to stay undiscovered? **B1.** Our assistants could compare nearby people's interests and ask whether their owners want an introduction. Both people would decide before meeting face to face. **R1.** Before my assistant asks, information about me has already been used to find a match. When did I agree to that comparison? **B2.** People would first enable introductions for this visit and select interests to share. Assistants would compare only information offered for that purpose. **R2.** Announcing that I am available for introductions reveals something personal. Can I seek company without advertising to everyone that I am looking? **B3.** We could exchange temporary invitations without names or seat locations. Identifying details would appear only after both people accepted an invitation. **R3.** In a nearly empty café, the invitation's timing may identify me anyway. Someone could infer who rejected them without seeing a name. **B4.** We could limit repeated attempts and withhold immediate introductions when too few people are present, reducing opportunities to identify someone through invitations. **R4.** You have cancelled the spontaneous encounter I wanted. Why should the assistant decide that a quiet café makes meeting someone unacceptable? **B5.** It could explain that risk and offer a one-time choice about my participation. My choice would not enable introductions for anyone else. **R5.** Who decides how much privacy people must trade for a social opportunity? An explanation informs that choice without making the trade equally acceptable.

Table A1 continued: Protocol Futuring portfolio

Scenario	Ten steps
G5. City & City  **Scenario.** People occupy the same streets but use different iOS and Android realities that hide one another in their device views. They still share physical space, so failing to see someone affects encounters and movement. Residents want to meet across these worlds without abandoning their own. The image shows the earlier premise, in which only virtual objects were hidden. Inspired by The City & The City [137], adapted into separate iOS and Android realities.	**Blue goal.** Help people join across ecosystems while keeping their own worlds. **Red question.** Why can willing people still fail to enter? What changes over time? **B1.** An optional meeting layer could reveal people and shared objects across platforms. Residents could enter it briefly while keeping their usual worlds available. **R1.** How would I invite someone I cannot see? A meeting space helps after we find one another, but how does that first encounter happen? **B2.** We could place a marker that both platforms display. People interested in contact could approach it and request an introduction across worlds. **R2.** People must reach that marker through streets containing invisible bodies. How does a voluntary meeting point address collisions with people who never approach? **B3.** Both platforms could show minimal cues of nearby bodies. Names and fuller appearances would become visible only after a mutual invitation. **R3.** That helps avoid collisions but presents neighbours as obstacles. How would we recognize one another as people who share a street? **B4.** Residents could jointly revise the cues and arrange voluntary events with full visibility. They would help decide how encounters happen locally. **R4.** One platform could delay updates or filter appearances. Residents on different devices would then experience different versions of what they had agreed. **B5.** Shared example scenes and independent comparisons could check both implementations. Residents would have a route to challenge departures from their agreement. **R5.** What if the dominant platform refuses to change? Switching devices makes residents pay for its refusal; what could actually constrain the company?
G6. Live Heritage  **Scenario.** At a historic arch, one visitor sees a reconstructed market while another sees only the present-day street. The reconstruction could help people discover the history of the place, but commuters may not want an interruption. Once visitors enter, the display also makes choices about whose memories count and which accounts receive the most attention. Extends Hyper-Reality [131] into public cultural heritage.	**Blue goal.** Help the public discover and explore the history of this place. **Red question.** Should this public story appear automatically or await discovery? **B1.** We could briefly show the historical market when someone enters the square, making a hidden story discoverable without requiring prior knowledge of it. **R1.** I cross this square to get to work and need today's street. Why should entering a public place trigger a scene I never requested? **B2.** A physical marker could explain that a reconstruction is available. People would deliberately open it; others could continue through the square without interruption. **R2.** Choosing history does not mean accepting the official account. What if the reconstruction celebrates redevelopment but omits the families who were displaced? **B3.** We could present several attributed accounts with supporting material. Visitors could compare the official reconstruction with memories from families connected to the place. **R3.** The first account will still receive more attention. Who chooses that entrance display, and how can less prominent contributors contest the choice? **B4.** A contributors' group could rotate the entrance display and publish reasons for changes. No account would permanently occupy the most visible position. **R4.** Displaced families may live far away or lack time for meetings. A seat in a group does not guarantee a practical voice. **B5.** We could pay for remote editing and accept requests to correct or restrict personal details. Participation would not depend entirely on local meetings. **R5.** One family requests removal of a painful account; another calls that erasure. Who decides what remains publicly visible while they disagree?

Table A1 continued: Protocol Futuring portfolio

Scenario	Ten steps
G7. Uninvited Gallery  **Scenario.** An independent group creates an AR exhibition positioned beside works inside a museum. Visitors can choose to view it through their own devices, although the museum has not approved it. The imagined dispute begins when the venue restricts access. The group wants an audience, while visitors need to understand who made the exhibition and move through the building safely. Inspired by MoMAR's Hello, we're from the internet [140].	**Blue goal.** Let ordinary visitors discover and choose to view the exhibition. **Red question.** If the venue blocks entry, who decides what visitors can see? **B1.** An independent directory could let visitors discover and choose the exhibition. Artists would not need the museum to list or promote their work. **R1.** Art appearing beside a museum object may seem officially endorsed. Finding it through another directory does not necessarily prevent that impression. **B2.** The entry screen and each work could identify its artists and independent status. Visitors could check the source while viewing it. **R2.** Clear authorship does not solve a crowded doorway. How can the museum keep people moving without deciding whether the art deserves to appear? **B3.** We could negotiate viewing positions or presentation changes for that passage, keeping the work accessible while addressing the obstruction it creates. **R3.** What stops the museum declaring the entire gallery sensitive to suppress criticism? Negotiation is unequal when one side controls the building. **B4.** Restrictions could name a location, reason, and duration. Artists could propose alternatives and ask an independent reviewer whether each restriction is justified. **R4.** If review finishes after the exhibition closes, artists have lost their audience. A finding in their favour cannot recreate that moment. **B5.** During review, restrictions could remain narrow and time limited. Visitors could still find the exhibition through the directory and view an outside version. **R5.** This work gains meaning from appearing beside those museum objects. Why should artists bear the loss of moving it while restrictions are reviewed?
G8. Reality Alert  **Scenario.** An emergency develops in a city where people use different digital overlays. Some see entertainment or advertising; others rely on captions or navigation assistance. An urgent warning needs to reach affected people across those views. The challenge is to authorize a necessary interruption without giving an issuer unlimited control over what people see or how long the interruption lasts. Extends Hyper-Reality [131] into cross-layer emergency alerts.	**Blue goal.** Reach affected people with urgent information across their layers. **Red question.** Who may interrupt a reality? Who handles errors, abuse and ending it? **B1.** A common emergency channel could reach every reality platform. People would receive urgent warnings even when their usual digital layers differ. **R1.** Who may use that channel? If a message can override my view, impersonators and advertisers have strong reasons to claim emergency status. **B2.** Platforms could verify approved issuers and limit their permitted locations, message types, and durations. An authenticated sender would still have boundaries. **R2.** An approved issuer can make mistakes, and an emergency may not affect everyone nearby. Why should an authentic message interrupt my entire view? **B3.** Interruption could match urgency and identify the warning's place, time, and uncertainty. A clear view of the physical surroundings would remain available. **R3.** Some overlays help me understand those surroundings. Removing my captions or navigation assistance may make responding to the emergency harder. **B4.** The warning could reserve space for essential assistance and follow local accessibility settings. Receiving an alert would not require losing those functions. **R4.** Who checks that an application needs protected space? What happens on older glasses that cannot arrange the warning and assistance together? **B5.** We could set minimum display requirements and test combinations across devices. Physical signs and non-AR notices would remain available when glasses cannot support them. **R5.** If the issuer disconnects before an all-clear, who restores my ordinary view? A warning needs an ending rule as well as an entrance.

Table A1 continued: Protocol Futuring portfolio

Scenario	Ten steps
G9. Invisible Target  **Scenario.** Players converge on a virtual creature in a public place. Through their devices, the gathering has an obvious purpose; other pedestrians see only a crowd blocking their route. Some outsiders are curious, while others want to pass or object. Explaining the invisible event is only a starting point because its physical effects already extend beyond those who joined. Combines Hyper-Reality [131] and Pokemon GO [161].	**Blue goal.** Let outsiders understand or join, while keeping a route through. **Red question.** How can people understand an event that affects them before they join? **B1.** A physical notice could explain the creature hunt alongside its in-game description. Outsiders could understand why people are gathering without installing anything. **R1.** Now I understand the crowd, but I still cannot get through. How does explaining the game restore a route for people not playing? **B2.** The game could stop adding rewards along an obstructed path and offer targets elsewhere, giving new arrivals a reason to spread out. **R2.** Players may keep waiting or rush to the replacement. Moving a target can move the crowd without reducing the problem. **B3.** Players could earn the same reward at several places and times. Nobody would need to squeeze into this crowd to avoid missing out. **R3.** Players came for a rare discovery they could share together. Are you solving the conflict by removing what made the event appealing? **B4.** Some collective hunts could happen at agreed times with a clear passage. Residents could request changes or complain without creating a game account. **R4.** Three games could each follow their agreement but together overwhelm one street. Who sees the combined effect on people living there? **B5.** A shared local schedule could register events across games and collect residents' requests. Organizers could identify clashes before directing several crowds to one place. **R5.** Who gets desirable times and locations? If established games dominate the schedule, coordination could protect residents while excluding smaller or informal events.
G10. Public Canvas  **Scenario.** A public digital layer mixes advertisements, travel information, and messages from local groups. A pedestrian wants useful updates but would rather not see advertising. Turning off the whole layer would remove both. Creating more selective control raises further questions about who defines essential information, who can publish it, and whose messages disappear when categories or rules change. Inspired by The Pervert's Guide to Cinema [59]	**Blue goal.** Keep public content useful while giving people control over viewing. **Red question.** Who can publish or remove content? Can people opt out and keep services? **B1.** People could disable advertising while keeping transit updates and community information. Avoiding commercial messages would not require abandoning useful public services. **R1.** What if advertising pays for those services and the provider refuses separation? A viewing preference does little when basic information depends on accepting ads. **B2.** Basic information could have separate funding and appear through several viewing applications. People could change applications without losing access to public services. **R2.** Who defines basic? An accessibility feature might become an optional extra even though someone depends on it to travel independently. **B3.** People with different access needs could define the basic service and challenge category changes. Removing a function would require a publicly explained decision. **R3.** Categories overlap. A shop's notice advertises a product and announces a community event; which part must I accept to learn about the event? **B4.** Items could show their publisher, funding, and whether they interrupt. People could filter particular sources instead of relying only on categories. **R4.** Institutions can easily prove their identity, but informal residents' groups may not. Would source checking make public visibility depend on institutional credentials? **B5.** Low-cost temporary publishing through local sponsors could help informal groups contribute. A route to appeal refusals would let them challenge exclusion. **R5.** What if a message criticizes its proposed sponsor? Who helps an unaffiliated resident publish when every available gatekeeper has a reason to refuse?

B Eliciting a General Map through Protocol Futuring

Table B1 records the trace from selected group arguments to the literature-comparison questions developed in Section 5. G1–G10 identify group discussions; B1/R1, B2/R2, and so forth locate successive contributions. Rows are overlapping analytical concerns, not participant frequencies or mutually exclusive codes.

Design-space concern	Proposal, challenge, and revision in the source accounts	Question for literature comparison and resulting facets
Interaction and equipment	**G1 B5/R5:** a common garment format is challenged by device upgrades and older devices that may not support it.	Compare device-dependent entry, supported representations, and fallback. A common format can still exclude people whose equipment cannot load it (**§§5.1–5.2**).
Leakage and discoverability	**G6 B1/R1:** automatic history offers immediate discovery but unwanted interruption. **B3/R3:** visitors may not know to aim at the arch, prompting signs. **B4/R4:** missed signs prompt a lightweight device notification.	Separate disclosing a layer's existence from displaying its content. Compare cue visibility, receiving channel, notice, and an intelligible next action (**§§5.1–5.2**).
The crossing process	**G2, participation exchange, B1/R1:** does watching count as joining? **B2** proposes QR-based voting and viewing. **R3/B4:** refusal is answered by switching off Bluetooth and Wi-Fi.	Distinguish observing, requesting, agreeing, participating, and leaving. Preserve observation as an endpoint and compare refusing one encounter with losing unrelated connectivity (**§§5.2, 5.4**).
Intentions and audiences	**G1 B1/R2:** the wearer wants the outfit seen, but red asks who consented to downloading it. **G4 B1/R1/B2:** interest labels meet an objection to being discovered; blue requires both people to enable matching.	Compare sender and recipient intentions independently, including remaining undiscovered. Willingness to reveal, notice, or meet need not extend to every operation or audience (**§§5.3–5.4**).
Permission management	**G1 B2/R2/B3:** publishing to a public layer raises consent, malicious-code, and loading concerns; blue proposes pre-upload review.	Separate publication, transfer, reception, and code admission. Reviewing content does not establish the receiver's permission; compare who authorizes each operation (**§5.4**).
Delegated agency	**G3 B1/R1:** Wi-Fi-based recommendations and payment require obtaining the customer's needs. **B3/R3:** blue questions hidden agent-to-agent exchanges; red proposes advance notice for sensitive information.	Distinguish connecting agents, disclosing preferences, recommending, and paying. Compare delegated authority, disclosure, and human intervention; notification alone leaves permission unresolved (**§§5.3–5.5**).

Design-space concern	Proposal, challenge, and revision in the source accounts	Question for literature comparison and resulting facets
Privacy of participants and bystanders	**G3 B4/R4:** a preferred service-agent likeness raises another person's portrait interests. **G5, motion exchange, B4/R4/B5:** avatar motion capture may disclose personal details; blue proposes coarse motion with inferred detail.	Identify whose data or likeness is exposed, including people outside the exchange. Compare sensing, inference, representation, and minimization separately from permission to enter (**§§5.4–5.5**).
Bystanders and accessibility	**G2, play exchange, R4/B5:** a passer-by entering the play area is rendered as an obstacle. **G6, access exchange, R4/B5:** a question about visual impairment prompts audio navigation around buildings and obstacles.	Ask how a protocol represents and affects nonparticipants, and who can perceive its cues. Compare continued passage, accessible receiving routes, and opportunities to object without joining (**§§5.1–5.2, 5.5**).
Time, continuity, and exit	**G4 B5/R5:** interest-history monitoring delays eligibility for matching and raises operating costs. **G1 R5:** evolving formats may leave older devices unsupported. **G5, motion exchange, R5:** disconnection remains unanswered.	Compare prerequisites, waiting, interruption, version changes, and continued operation. Extend the map to what exit ends and what persists; the accounts raise continuity questions without specifying a complete exit protocol (**§§5.2, 5.4–5.5**).
Cross-platform interoperability	**G5, format exchange, B2/R2/B3:** a shared cloud does not settle consistent rendering, prompting a common format. **R5** asks what happens when a device cannot load the scene.	Separate service access, shared representations, rendering compatibility, and device support. Compare what crosses, which components must agree, and fallback when capabilities differ (**§§5.2, 5.5**).
Upload review and admission	**G5, format exchange, R3/B4:** a question about which users may enter is answered with code upload and review. **R4/B5:** confusion between public and personal scenes prompts selectable layers.	Separate reviewing an uploaded artifact, admitting a user, and choosing what to render. Specify what is checked, by whom, whose approval is recognized, and how each decision is enforced (**§§5.4–5.5**).
Protocol governance	**G5, relay exchange, B4/R4:** occasional encounters still require a manager. **B5/R5:** reducing the public layer to signal relay leaves protocol authorship and fairness unresolved.	Compare device, application, platform, and cross-provider authority. Connect rule-setting and operating responsibilities to ways of contesting decisions; minimal exchange still requires governance (**§5.5**).

Table B1. Eliciting a general map through Protocol Futuring. Each row connects a group argument to a question for literature comparison and its development in the five facets (§§5.1–5.5). Rows are overlapping concerns, not participant frequencies or mutually exclusive codes.

C Statement of AI Use

Generative AI was used for two bounded purposes. It rendered the comic panels in Figure 2, and it expanded the illustrative portfolio dialogues in Appendix A from the prepared scenario prompts. Those dialogues are identified as generated in Appendix A and carry no analytic weight. Generative AI also assisted in organizing candidate comparisons in the literature synthesis (Section 5).

Generative AI did not produce, edit, translate, or summarize the recorded team exchanges, and it did not perform the analysis reported in Section 4.4. Translations of team text are the authors' own. Consensus and Undermind assisted literature retrieval and source examination. The authors are responsible for all interpretations and supporting citations.

D Participant Characteristics

P01–P40 list the 40 available questionnaire records in source order.

Table D1. Participant overview (N = 40).

Participant	Age	Gender	Field	XR use	XR creation
P01	25	Male	Engineering	No	No
P02	28	Male	Other	No	No
P03	23	Male	Other	No	No
P04	24	Male	Other	No	No
P05	23	Female	Computer science	No	No
P06	22	Male	Computer science	1–2 years	No
P07	24	Female	Computer science	No	No
P08	34	Female	Other	No	No
P09	21	Male	Computer science	No	No
P10	25	Female	Engineering	1–2 years	No
P11	22	Female	Computer science	No	No
P12	25	Male	Engineering	No	<1 year
P13	23	Female	Computer science	No	No
P14	23	Male	Computer science	No	No
P15	23	Male	Other	No	No
P16	22	Female	Computer science	No	No
P17	23	Male	Computer science	No	No
P18	26	Male	Other	<1 year	No
P19	22	Female	Computer science	No	No
P20	22	Female	Computer science	No	No
P21	23	Male	Computer science	No	No
P22	23	Female	Computer science	No	No
P23	23	Male	Computer science	No	No
P24	24	Male	Computer science	No	No
P25	23	Female	Computer science	No	No
P26	22	Male	Computer science	No	No
P27	23	Female	Design	No	No
P28	23	Female	Design	<1 year	<1 year
P29	24	Male	Computer science	No	No
P30	22	Female	Computer science	No	No
P31	25	Male	Computer science	1–2 years	No
P32	23	Male	Computer science	No	No
P33	23	Female	Design	No	No
P34	23	Female	Design	<1 year	<1 year
P35	22	Male	Computer science	No	No
P36	23	Female	Computer science	No	No
P37	31	Male	Computer science	1–2 years	No
P38	22	Female	Computer science	No	No
P39	31	Male	Computer science	1–2 years	No
P40	32	Male	Computer science	1–2 years	No