

REPHRAIN
Protecting citizens online



Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent

Joseph O'Hagan, University of Glasgow

Pejman Saeghe, University of Glasgow

Jan Gugenheimer, TU-Darmstadt

Daniel Medeiros, University of Glasgow

Karola Marky, Leibniz University Hannover

Mohamed Khamis, University of Glasgow

Mark McGill, University of Glasgow



February 2023



Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent

JOSEPH O'HAGAN, School of Computing Science, University of Glasgow, UK

PEJMAN SAEGHE, School of Computing Science, University of Glasgow, UK

JAN GUGENHEIMER, TU-Darmstadt, Germany

DANIEL MEDEIROS, School of Computing Science, University of Glasgow, UK

KAROLA MARKY, Leibniz University Hannover, Germany

MOHAMED KHAMIS, School of Computing Science, University of Glasgow, UK

MARK MCGILL, School of Computing Science, University of Glasgow, UK

Fundamental to Augmented Reality (AR) headsets is their capacity to visually and aurally sense the world around them, necessary to drive the positional tracking that makes rendering 3D spatial content possible. This requisite sensing also opens the door for more advanced AR-driven activities, such as augmented perception, volumetric capture and biometric identification - activities with the potential to expose bystanders to significant privacy risks. Existing Privacy-Enhancing Technologies (PETs) often safeguard against these risks at a low level e.g., instituting camera access controls. However, we argue that such PETs are incompatible with the need for always-on sensing given AR headsets' intended everyday use. Through an online survey (N=102), we examine bystanders' awareness of, and concerns regarding, potentially privacy infringing AR activities; the extent to which bystanders' consent should be sought; and the level of granularity of information necessary to provide awareness of AR activities to bystanders. Our findings suggest that PETs should take into account the AR activity type, and relationship to bystanders, selectively facilitating awareness and consent. In this way, we can ensure bystanders feel their privacy is respected by everyday AR headsets, and avoid unnecessary rejection of these powerful devices by society.

CCS Concepts: • **Human-centered computing** → **Mixed / augmented reality**; **Virtual reality**; **Ubiquitous and mobile devices**; • **Security and privacy** → **Social aspects of security and privacy**; **Privacy protections**; **Usability in security and privacy**.

Additional Key Words and Phrases: Augmented Reality; Privacy; Bystanders; Altered Reality; Extended Perception; Biometrics;

ACM Reference Format:

Joseph O'Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. 2022. Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 4, Article 177 (December 2022), 35 pages. <https://doi.org/10.1145/3569501>

Authors' addresses: [Joseph O'Hagan](mailto:Joseph.OHagan@glasgow.ac.uk), School of Computing Science, University of Glasgow, UK, Joseph.OHagan@glasgow.ac.uk; [Pejman Saeghe](mailto:Pejman.Saeghe@glasgow.ac.uk), School of Computing Science, University of Glasgow, UK, Pejman.Saeghe@glasgow.ac.uk; [Jan Gugenheimer](mailto:Jan.Gugenheimer@tu-darmstadt.de), TU-Darmstadt, Germany, Jan.Gugenheimer@tu-darmstadt.de; [Daniel Medeiros](mailto:Daniel.PiresdeSaMedeiros@glasgow.ac.uk), School of Computing Science, University of Glasgow, UK, Daniel.PiresdeSaMedeiros@glasgow.ac.uk; [Karola Marky](mailto:karola.marky@itsec.uni-hannover.de), Leibniz University Hannover, Germany, karola.marky@itsec.uni-hannover.de; [Mohamed Khamis](mailto:Mohamed.Khamis@glasgow.ac.uk), School of Computing Science, University of Glasgow, UK, Mohamed.Khamis@glasgow.ac.uk; [Mark McGill](mailto:Mark.McGill@glasgow.ac.uk), School of Computing Science, University of Glasgow, UK, mark.mcgill@glasgow.ac.uk.

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.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2474-9567/2022/12-ART177 \$15.00

<https://doi.org/10.1145/3569501>

1 INTRODUCTION

As Augmented Reality (AR) devices tend towards everyday/all-day, ubiquitous [64] wearable and fashionable form factors (e.g. smart glasses [39]), paired with advances in the underlying technology (e.g. increasing field of view) and software driving compelling AR experiences, consumer uptake will likely increase significantly. Such devices are typically equipped with a variety of sensors that collect data necessary to drive key functionality—henceforth, referred to as *requisite sensing*. For instance, on-board cameras are required to enable 6DoF positional tracking and render exocentric spatial virtual content [94] and microphones are required to enable voice interactions [14, 24]. While necessary, the requisite sensing capabilities of an AR device have the potential to open significant privacy and security risks for users and bystanders [10, 52, 124]. For instance, in addition to being able to uniquely identify their users [98], AR platforms will be able to gain insights into their users' mental/cognitive processes and phenomenological experiences [56, 63], infer their stress/arousal levels and affective states [3], and infer sensitive personal information and characteristics such as gender, age, ethnicity, and more [76]. This is possible due to the wealth of data captured by such devices from on-board cameras and microphones [92], as well as physiological and biometric sensing (e.g. see Kröger *et al.* [76] for an exploration of what eye-tracking alone can contribute here).

The user of an AR headset may, to some degree, 'opt-in'/consent to requisite sensing activities, either actively through giving permission, or passively through agreeing to terms of service. Crucially, however, *bystanders*—i.e., those people physically within sensing range of an AR headset—typically have no capacity to consent to, or be made aware of, the activities of a user's AR headset that pertain to them. As a result, bystanders may find themselves deliberately or inadvertently subject to AR-enabled surveillance [116, 145], with it being possible to sense and process biometric data (e.g. inferring identity from gait [98]), non-contact physiological data [136], volumetrically capture appearance and augment identity [77, 123], determine protected characteristics [18], instrument behaviour and actions [101] and more [92]—predominantly from camera data alone. Research has shown that users may in fact be more concerned with risking bystanders' privacy, rather than their own; since bystanders could change their behaviour towards the AR user, if they perceive that their privacy is at risk [119]. Moreover, bystanders typically do not have a say in the user's decision to buy and wear an AR device, and users may not be aware of the privacy risks that they are exposing bystanders to in using said technology [41]. These risks are amplified when bystanders may not even be aware that there is a nearby AR headset, given the expectation that in time these devices will be indistinguishable from existing glasses.

The implications of AR-driven privacy risks, for users and bystanders alike, represent pertinent issues beyond face-value concerns regarding recording activities [5, 114]. AR devices' capability for "*persistent, ubiquitous recording*" [21] and *veillance* of others [86] has the potential to erode a bystander's "*reasonable expectation of privacy*" [45], supporting "*cyborg stalkers*" [68], and facilitating a "*global panopticon society of constant surveillance*", where "*the possibility of being recorded looms over every walk in the park, every conversation in a bar, and indeed, everything you do near other people*" [124]. As a result, the anticipated adoption of AR headsets may put AR users in opposition to bystanders—with users wanting to leverage the full potential of these devices, at the cost of infringements to bystanders' privacy.

1.1 Limitations of Existing Work

Whilst Privacy-Enhancing Technologies (PETs) have been proposed to resolve this tension between AR users and bystanders, prior research into PETs in relation to everyday AR has four key limitations:

(1) There is a predominant focus on "data-centric categorization" of privacy properties of AR devices (e.g. see De Guzman's foundational [30]), identifying vulnerabilities from a technical perspective without sufficient consideration of the user, and crucially bystander, experience of these vulnerabilities and associated PETs addressing them.

(2) Where user/bystander PETs are explored, there is a tendency to focus on “all-or-nothing” access control as mitigation strategy to AR privacy risks. Examples include using context-aware physical shutters on cameras [73, 147]—such an approach is incompatible with AR requisite sensing, given the continual need for positional tracking.

(3) Because of the focus on camera access control, there is a lack of a granular understanding of the public’s perception of and attitudes towards the risks posed by AR-driven processing activities. Suppose we assume that requisite sensing will always be active. In that case, the onus is instead on understanding what is permissible at an activity level, rather than a hardware level, if we are to create PETs that protect bystanders whilst still allowing users to take advantage of the powerful capabilities of AR headsets.

(4) By focusing on face-value camera data alone, the importance of what can be inferred from camera data—e.g., biometric identification [145] and biometric psychography [57]—is significantly undermined. Moreover, support for such invasive processing activities may not be apparent to the public [106, 151, 152]. Consequently, attitudes towards such activities may differ from those toward the underlying sensing.

1.2 Contribution

We first review a range of pertinent threats posed by AR headsets to privacy. In particular, we focus on AR activities previously demonstrated to be feasible and have the ability to generate privacy infringing data regarding bystanders’ internal processes, external activities, and identity/personal characteristics. We summarise PETs that have been developed to counteract risks to bystander privacy. We report on the results of an online survey ($N = 102$) investigating the public’s attitude towards AR-driven processing activities from two perspective: (1) the user of an AR headset, and (2) as a bystander to a nearby AR headset. We explore how the needs for awareness and consent vary by activity, and whether the relationship between user and bystander influence these needs. Our results provide novel insights into the public’s attitudes towards AR headsets and the activities enabled through AR requisite sensing:

- There is a low awareness and high concern regarding AR headset activities pertaining to bystander data (e.g., biometric identification, estimating internal and physiological state). As a result of being exposed to the AR capabilities presented in the survey, the majority of respondents were now more concerned with public use of AR devices.
- Attitudes towards opt-in/out consent are influenced both by AR activity type and social relationship.
- There is a strong desire to be aware of the AR headset’s activity, in particular when strangers are enacting these activities. The desired awareness of specific AR activities was influenced by activity type and prior consent. Where prior consent was not sought, awareness of specific AR activities in action was more greatly desired.

Our findings suggest new directions for everyday AR PETs. We provide an evidentiary basis regarding the need to support (1) usable mechanisms for providing consent, and (2) greater transparency regarding AR activities that pertain to bystanders. We argue if the risks posed by AR devices to the public’s privacy and security are not identified, addressed, and communicated transparently and clearly, that everyday wearable AR risks another mass rejection [58].

2 THE RISKS POSED TO BYSTANDERS BY ‘AR REQUISITE’ SENSING

Critical to the function of AR headsets is head-mounted, mobile, always-on pervasive sensing of the user and their environment — from visual (e.g., RGB and non-visible multi-spectral imaging [136]) to auditory sensing (e.g., directional microphones [93]); from neural activity (e.g., EEG for emotion detection [46]) to physiological signals (e.g., eye/gaze tracking [147]). The incorporation of such sensing capabilities into mobile, everyday form factors is unique to AR (and Extended Reality (XR)) headsets and goes beyond other forms of body-worn cameras [73] and Internet-of-Things (IoT) sensing [35]. This sensing is *fundamentally necessary* to provide AR functionality; we

refer to this as *AR requisite* sensing. For example, cameras are required to enable a headset's understanding of its position in a 3D space, to track hand movements, etc. Combined with microphones (and other sensing), cameras enable applications to understand the users' context, behaviour, needs, and surrounding environment. However, such sensing capabilities pose obvious privacy risks to both users and *bystanders* within physical range of these sensors. For bystanders, AR “amplifies and combines existing privacy issues” [34] through deliberate *sousveillance* or inadvertent surveillance [86] of AR headset activity and its supporting sensing, potentially causing an “erosion of the concept of reasonable expectation of privacy” [45].

Building on the challenges in XR privacy identified by the IEEE Ethics of XR Initiative [92], and key essays on AR privacy (e.g., [34, 35, 82, 124]), we provide an overview of notable privacy and security risks facilitated by AR requisite sensing. The risks posed to bystanders by everyday XR devices include, but are not limited to: (1) identity, anonymity, and biometric ID, (2) mental privacy (e.g., behaviour, internal state, and biometric psychography), (3) physiological privacy and health, (4) augmented perception and personal surveillance, and (5) capture, appropriation, and alteration of appearance. The activities underpinning these risks have been demonstrated to be feasible in research or practice. While the risks presented here are not, and cannot, be exhaustive, they are illustrative of the variety (and types) of XR-facilitated risks that go beyond simple image capture and recording alone.

2.1 Identity, Anonymity and Biometric ID

Prior work has demonstrated, based on basic captured positional tracking data, users can be uniquely identified with an accuracy of 95% [98], while soft biometric traits (e.g., body shape) can be obtained at a distance without an individual's cooperation [120]. Driven by advances in machine learning and computer vision, it will therefore be trivial for an AR device to segment, classify, and track individuals. Moreover, devices will soon be capable of volumetrically capturing an individual to generate a 3D mesh of their body and likeness [23]. At its most basic level, this data will unlock the ability to pseudo-anonymously identify and track nearby individuals whenever they are within sensing distance of the AR user. Backed by social media platforms, other publicly available biometric data sets and aided by cloud computing platforms, our capacity to break the veil of public anonymity will never be greater [6, 7]; a prospect exacerbated by the public's limited awareness of where/how such data is used [159]. Consequentially, applications will have all the necessary data to strip bystanders of their anonymity. As Meta recently acknowledged, such platforms will likely incorporate facial recognition due to its ‘obvious’ business benefits but suggested this would only be supported “*if it could be done in a way the public and regulators were comfortable with*”, such as enabling bystanders to “mark their faces as unsearchable” [53]. However, beyond compromising anonymity, such sensing could also be used to estimate characteristic and protected traits [1], such as gender, age, sexuality, accessibility needs, race, and other personally identifiable information, potentially without the user's knowledge or consent. The consequences of this may be significant for the individual concerned—for example by being used for the purposes of discrimination and profiling [4] further cementing stereotypes and biases.

2.2 Mental Privacy - Behaviour, Internal State, and Biometric Psychography

The data collected by AR headsets will unlock the ability to develop a sophisticated, longitudinal understanding of users and bystanders—from their behaviours, intentions and actions [101]; to mental and cognitive processes and phenomenological experiences [56]; to stress/arousal and affective states [3]; and personality traits [59]. This provides third parties with unique insights into intimate details of our lives, such as our preferences, and attitudes—a class of processing Heller defined as “*biometric psychography*”, where biometric data is used to identify a person's interests [56]. Consequently, in addition to an erosion in an individual's body integrity [99], physical privacy [100], and outward presentation, their mental privacy will also be eroded [63].

2.3 Physiological Privacy and Health

Beyond what can be seen by the naked eye, AR headsets can also feature multispectral, non-visible optical sensing. For example, prior work has examined the use of near-IR cameras for non-contact physiological measures, such as heart rate variability at-a-distance [136]. When coupled with other soft biometric data [120], such as information regarding gait or skin pallor, or compared to prior historic captured data (in the case of identified individuals), AR headsets will have a significant capacity to generate physical health-related data pertaining to bystanders, which could in-turn form the basis of more sophisticated insights into e.g., mental privacy.

2.4 Augmented Perception and Personal Surveillance

AR enables users not just to capture reality but also to alter, augment, and diminish their perception of reality. Termed *augmented perception* [61, 133, 134], this enables individuals to extend and amplify their sensory range by, for example, providing “superhearing” (e.g., selectively enhancing/suppressing audio to improve speech perception [28]) or “supersight” (e.g., visualising out-of-view objects [48]). Consequentially, these technologies offer significant benefits for overcoming situational, temporary and permanent impairments, however, they also pose a significant security risk [30, 50, 127] potentially bestowing individuals with super-sensory capabilities and memories that could be used for personal surveillance (e.g., supporting sophisticated “shoulder surfing”-type observation attacks [155]).

2.5 Capture, Appropriation, and Alteration of Appearance

The sensing and tracking of bystanders also suggests new methods of losing control over how (and when) we are perceived by others. At its base level, visual sensing can trivially segment and volumetrically capture the appearance of a bystander, going beyond what was previously possible with 2D image/video capture. This can unlock significant potential for abuse, for example seeing strangers capture and appropriate your appearance for a later VR experience [112], or result in “identity hacking”, such as identity theft [141].

Furthermore, the real-time tracking, segmentation and identification of bystanders would unlock headset-based augmentation and alteration of identity, digital self-presentations, and face-to-face interactions [77, 123]. Without control, however, this ability for others to augment how we are perceived could enable new forms of abuse. For example, Lebeck *et al.* highlighted concerns around individuals regarding AR overlays based on sensitive information drawn from elsewhere [79]. But it is evident more extreme abuses will soon be feasible, such as a convergence of AR sensing and deep fake technology [22] enabling users to sexualise [67] or appropriate someone’s identity for socially unacceptable reasons. So great are concerns for this that Lemley *et al.* have already considered the legality of this ability to augment our personal sensecape and the sensecapes of others, asking: “What if people use this... to make your avatar appear ridiculous... without your knowledge or consent? Or what if they want to make you appear naked?” [81].

2.6 Emergent and Future Processing Risks

The risks described thus far represent some of the key concerns, predominantly around what has been termed *input protection* [30] or *input privacy* [126], pertaining to the security and privacy of data gathered by an AR device. However, underlying these risks is a common theme—our increasing capacity to process captured data to reveal new, often unanticipated, insights. Novel processing capabilities build on and extend what is possible (see Figure 1 for examples). These at-times imperfect [18] composite processing activities increasingly employ sophisticated machine learning algorithms and AI-driven approaches which can be trained to predict/infer information about identity, behaviour, activity, internal state—and then are increasingly offered to developers as cloud services that can trivially enhance the capability of an application to process sensed data (e.g., Microsoft Cognitive Services, Apple’s CoreML, Google Cloud). Consequently, the risks detailed here are not exhaustive;

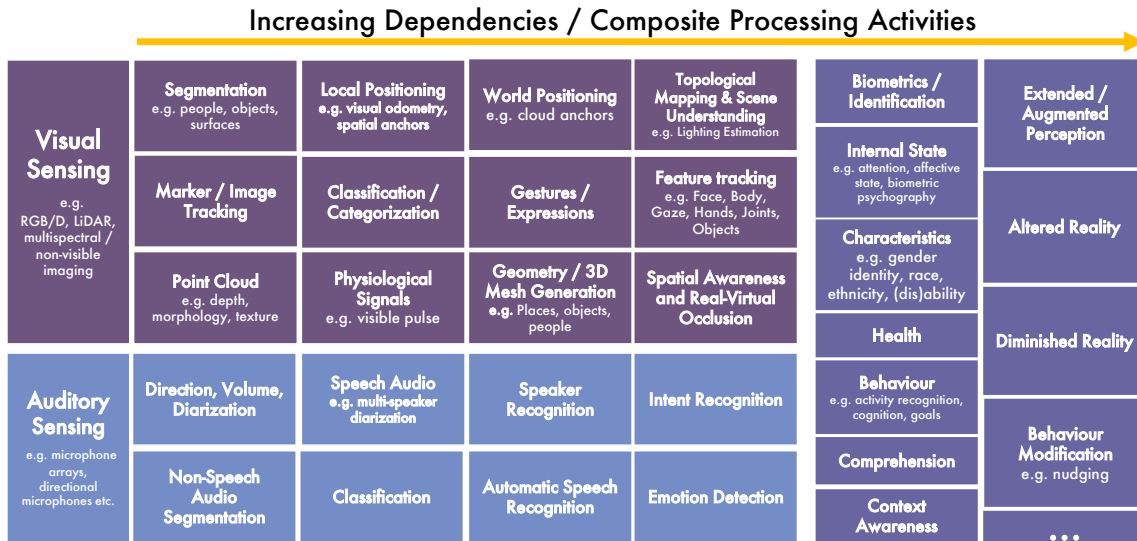


Fig. 1. Examples of processing activities / outputs based on visual, auditory and visual+auditory sensing, presented broadly in order of increasing inter-dependencies and consequent complexity.

further risks will emerge in time. However, they do highlight significant themes within the AR–bystander risks detailed thus far: of compromising anonymity and identity; of estimating cognitive, affective and physiological states; of veillance and recording of our actions and identities; and of the capacity of AR to alter how we appear to others. Where prior research has focused on PETs as applied to the underlying sensor hardware, we argue there is a research gap around the need for more granular PETs as applied to the breadth of emerging AR activities enabled by that hardware.

3 THE PRIVACY RISK-MITIGATION GAP BETWEEN AR USERS AND AR BYSTANDERS

Given the breadth of potential violations of anonymity, privacy, and identity, it is pertinent to ask: *should AR headsets have such capabilities?* If such capabilities are removed from AR headsets (e.g., because they are deemed too risky), the aforementioned privacy risks are eliminated. However, we argue it is highly likely many of the aforementioned capabilities may find their way into consumer mass-market AR headsets in the future. There are a number of drivers for this possibility: benefiting accessibility use cases through sensing and revealing pertinent bystander data to address a given impairment [80]; supporting the latest advances in augmented intelligence [162]; empowering individuals to better curate/control how they are perceived by others, etc. In effect, it is not enough to suggest these activities will not occur. Instead, there is a need to anticipate their potential occurrence and understand to what extent we can integrate protective measures into AR technology, to better safeguard or empower bystanders. This is particularly pertinent as for some of these activities, there exist both benign, and distinctly abusive use cases.

When considering AR risks to privacy and security, AR users by default have some capacity for agency and control over whether such outcomes occur. For example, AR users will have the capacity to opt-in or out of such data capture, be made aware of when this is occurring and who their data is shared with, provided the

headset and/or application(s) supports this. In contrast to the AR user, bystanders may have little-to-no awareness that a headset is active; what sensing is currently active; whether said sensing and subsequent processing activities pertain to them; and whether they consent to being sensed by the AR headset. Addressing this gap are *Privacy-Enhancing Technologies (PETs)*, evolving social norms, formal guidelines and legislative measures, which were historically applied to CCTV, IoT and other “ubiquitous intelligent cameras” [71], more recently applied to lifelogging and body-worn video [25, 60], and are increasingly being considered from the perspective of AR bystanders.

3.1 Privacy-Enhancing Technologies (PETs)

We focus on PETs with the greatest relevance to bystander–AR user interactions, namely those pertaining to *input* and *data* protection approaches that facilitate the properties of access control, confidentiality, awareness, and consent as described by DeGuzman *et al.* [30]. Note: we exclude consideration of *output*, *user interaction*, and *device* protection approaches. The threats associated with output and device protection do not pose risks to bystanders, whilst user interaction protection predominantly pertains to secure, authenticated shared multi-user XR experiences. We focus on activities that are not shared, where the AR device acts upon sensed data pertaining to bystanders for the purposes of the AR user. See Table 1 on the following page for an overview, and see [30] for a broader overview of XR-related PETs.

3.1.1 Access Control and Contextual Privacy. Access control refers to the ability to control access to raw sensor data or the underlying hardware generating this data. If we consider visual and auditory sensing (the predominant means used by AR headsets), access to camera and microphone data streams gives a platform/application uninhibited scope to extract a wealth of information. With respect to intrinsic access control, PrivacEye [147] employed automatic detection of the “privacy sensitivity” of a given context to determine the use of a mechanical shutter to physically disable the head-mounted camera. Participants noted the visibility of the physical occlusion provided by the shutter was preferred over visual feedback, such as LEDs, because the physical coverage increased the trustworthiness of the system to bystanders. However, intrinsic access control through camera shutters brings with it significant challenges when considered for AR headsets. Cameras are used to drive 6 Degrees of Freedom (DoF) visual inertial positional tracking [94]. If a shutter completely occludes the headset camera(s), this immediately necessitates the headset fallback to 3DoF tracking, and removes its capacity to augment the environment with exocentric world-locked AR content. Such a degradation in capability would significantly undermine many key use cases envisioned for everyday AR.

Various contextual information has also been suggested to drive access control, taking into account the location and environment (e.g., “world-driven access control” [128]), existing social connections [156, 158], social signals [96], and accessibility needs [158]. For accessibility in particular, there is evidence individuals “*would support special permissions to be granted to specific groups of persons when there is a justified reason to do so*” [39], if “*higher security assurances can be made by improving their control over how their information is shared*” [10]. This reflects the influence of societal roles on the perception of the underlying technology, and suggests the possibility of “social acceptability calculus” [39] where activities may be more or less justifiable from a bystander perspective, despite the potential for misuse, based on personal relationship or societal benefits. However, this picture is not clear cut. Ahmed *et al.* [10] found despite visual information about bystanders already available to sighted individuals, there was nonetheless a reticence to capture and convey that information (e.g. demographics, activity, emotion) for accessibility needs.

Finally, in terms of extrinsic bystander led access control, Ahmad *et al.* [9] elicited a variety of bystander both social and technology-driven coping and controlling mechanisms for privacy regulation, including masking behaviours (e.g. saturating a microphone through loud music); ignoring “low-priority” inputs such as audio sensing; blocking behaviours such as physically covering the sensors on a device; withdrawing from the social

Table 1. Overview of input protections and input privacy mitigations. We summarize PETs from two perspectives - those that are enacted by or on the AR User intrinsically; and those that are enacted by or otherwise related to the AR bystander extrinsically [30].

Mitigation	Perspective	Examples from Literature and Consumer Products	Risks to Effectiveness for AR HMDs
Access Control Blocking or circumventing sensing	User / Intrinsic	<ul style="list-style-type: none"> • Social pressure and norms e.g. removing headset in someone else's home [79] • Context sensitivity e.g. <i>PrivacEye</i> mechanical shutter [147] and need for physical occlusion [9, 73]; use of social signals to determine muting audio or disabling cameras [96]; manufacturer-enforced recording bans [79] 	<ul style="list-style-type: none"> • Physical occlusion of sensors incompatible with need for 'always-on' access for e.g. visual-inertial tracking, voice assistants.
	Bystander / Extrinsic	<ul style="list-style-type: none"> • Blocking behaviours such as withdrawal from the sensing environment/range [9] or otherwise disabling the sensing. • Social pressure toward users e.g. asking them to remove their headset [75], aggression toward sensing and/or user [9] • Transferring control to bystanders [73] e.g. to turn off device • Embedding or encoding privacy information in the environment that would prevent auditory capture [158] such as "world-driven" access controls [128]. 	<ul style="list-style-type: none"> • Social and blocking mitigations risk socially unacceptable AR-bystander interactions. • Embedded privacy restrictions necessitate the infrastructure to define and pseudo-anonymously share said restrictions.
Data Obfuscation, Sanitization and Masking	User / Intrinsic	<ul style="list-style-type: none"> • Selective obfuscation [121] and image degradation [36] of bystander data e.g. blurring faces and bodies [11, 54] • Incorporating noise in data through differential privacy [146] 	<ul style="list-style-type: none"> • Data can still be reconstructed and inferred e.g. using machine learning • Trust gap regarding obfuscation being applied [122]
	Bystander / Extrinsic	<ul style="list-style-type: none"> • Active obfuscation e.g. masking audio with additional noise [9]; defensively subverting visual sensing [115] e.g. using IR LEDs to overcome facial ID with FacePET [113] and physiological sensing with InPhysible [91] - a means of overcoming the "trust gap" [122]; 	<ul style="list-style-type: none"> • Mitigations require changes in behaviour from bystanders, and wearing active defensive countermeasures. Such measures could be circumvented if they see significant adoption.
Consent Directing access control and obfuscation	User / Intrinsic	<ul style="list-style-type: none"> • Pre-activity through checklist permissions during application installation (e.g. see any Android-based XR headset); • During activity through permission prompts and direct UI control of headset activity e.g. Microphone prompts on Android-based XR headsets. 	<ul style="list-style-type: none"> • Users may not fully engage with considering the need for requested permissions
	Bystander / Extrinsic	<ul style="list-style-type: none"> • Pre-activity such as proactive opt-in and opt-out designs [32]; preference toward pre-determined consent [17, 140], typically enacted through tracked physical artefacts [137] or biometric ID combined with cloud-based preferences [75, 138] • During activity through interaction, from verbal social indications that direct the AR user's activity [69] to gestural technology-mediated interactions conveying opt-in/out to the headset directly [69], examples of "user empowerment" approaches [44, 109]. 	<ul style="list-style-type: none"> • Implicit consent activities have similar requirements to embedded privacy restrictions - there is a need for infrastructure to support defining and conveying those permissions. • Explicit consent activities risk interrupting social interactions and placing onerous requirements on bystanders to actively manage their consent.
Situational Awareness and Activity Transparency	User / Intrinsic	<ul style="list-style-type: none"> • UI and notifications e.g. using notifications and nudging to raise awareness of device camera activity [55, 105, 108] 	<ul style="list-style-type: none"> • Change blindness and inattention blindness may lead to ignoring such prompts.
	Bystander / Extrinsic	<ul style="list-style-type: none"> • Presence and Privacy Notices e.g. "peacocking" visibility of cameras [73] and emphasizing noticeability [70]. • Device Status such as color indicator LEDs conveying device on/off state and basic activity [9, 147] • Mode transparency through privacy labels [88, 89], conveying information regarding intent and activity through iconographic, textual, and color-coded displays [73] • Scope and activity using on-device displays to convey detail regarding capture area and field of view [73] or intent (e.g. EyeCam [150]), toward full presentations of activity such as MirrorCam [72]. 	<ul style="list-style-type: none"> • Bystanders may lack prior knowledge regarding the meaning of LED and other status indicators, or overlook them entirely • As we convey more information regarding the AR user's activity, we may risk exposing personal details to bystanders, degrading their privacy. [88]

interaction by e.g. leaving the room; or exhibiting a degree of aggressive behaviour toward the device, such as actively seeking to tamper or disable it.

3.1.2 Data Obfuscation, Sanitization and Differential Privacy. An alternative to binary access control is to intrinsically or extrinsically obfuscate/capture/sanitize what data is made available to applications. De Guzman *et al.* discussed input sanitization in terms of intrinsic user-defined policies, context-based sanitization that identifies the presence of sensitive objects, and extrinsic sanitization based on privacy preferences [30].

Intrinsic, selective obfuscation is intended as a means to overcome the “trust gap” [121, 122] in potentially privacy-infringing technology. Alharbi *et al.* [11] and Hasan *et al.* [54] used image obfuscation to remove bystanders whilst still enabling non-bystander processing activities. Such an approach exposes a tension between obfuscating sensor data versus employing access control coupled with separate APIs/data streams for specific processed data sources. Other approaches employ inpainting to completely remove bystanders in the user’s view and fill in the missing parts of the image in a manner visually consistent with the background [160]. Alternatively, differential privacy can be used to add noise to the data “so as to minimise chances to infer privacy-sensitive information... whilst, at the same time, still allow[ing] use of the data for desired applications” [146], e.g., through image degradation [36].

However, extrinsic obfuscation controlled by bystanders is possible, by employing personal privacy preferences which inform intrinsic obfuscation, or through active countermeasures. A variety of wearable systems have been developed to enable bystanders to actively/defensively subvert any camera-based facial sensing [115]. For example, FacePET [113] effectively blinded cameras by using infrared spectrum LEDs integrated into glasses, whilst InPhysible [91] actively camouflaged bystander physiological signals using IR LEDs to fake signals imperceptible to the human eye.

3.1.3 Consent and Personal Privacy Preferences. Mechanisms for bystanders to indicate opt-in/out consent preferences, either before or during activity, have been repeatedly explored, with said preferences being used to either enact selective obfuscation/sanitization [137] or access control [10]. Koelle *et al.* [69] examined opt-in/out gestures for privacy mediation in smart camera interactions, finding tensions between identifying clear gestures to convey consent without resulting in false positives or other additional unintended meanings. Shu *et al.* [137] explored how individuals could passively indicate preferences toward recording through using wearable tags, as well as utilizing active ways of communicating preferences such as hand gestures. In both cases, bystander faces were blurred if recording was not permitted. Cardea used biometric ID to retrieve personal privacy preferences regarding visual privacy [138], whilst Krombholz *et al.* [75] explored how preferences regarding privacy could be conveyed through an app tied to cloud-based facial recognition, and worn physical artefacts denoting the wearer did not want to be recorded. However, some participants noted an unease towards this, e.g., “the server behind the app bothers me just as much as [the AR headset] does”.

Notably, Koelle *et al.* [69] found some participants felt the onus should be on users, rather than bystanders, to ensure privacy protection, suggesting instead consent could be managed automatically by default, with additional opt-in/out interactions being available for special cases where more individual or granular control was required. However, they did not determine what might constitute such cases. Singhal *et al.* [140] found participants preferred to give prior permission to recording activities, so they could moderate their actions appropriately. They also noted some anecdotal impact on attitudes based on place, activity, and whether those recording were strangers or known. Moreover, bystanders may be unaware of there being activity to consent to. In the context of live streaming, bystanders noted they were “not fully aware of when their image or speech is being live-streamed in a casual context”, specifically wanting stronger awareness mechanisms and the ability to consent to such activity [43]. Whilst others have repeatedly re-affirmed the need for mechanisms for consent regarding control over identity and interpersonal space (Bye’s ‘consent to augment’ [19]).

3.1.4 Awareness and Activity Transparency. Interleaved with consent is the notion bystanders should be made aware of their exposure to AR sensing [131], and when they are the focus of sensor activity. Marky *et al.* [88] found smart home visitors lacked the ability to “judge [the] consequences of data collection and [a] means to express their privacy preferences”. They suggested this necessitated mechanisms for gaining awareness of how bystanders are sensed, knowledge regarding how sensed data is used, and the sensitivity of what is being captured (e.g. through privacy labels facilitating mode transparency regarding what the sensing was doing), informing any use of consent mechanisms and privacy decisions. Meanwhile, Prange *et al.* [117] proposed a concept to visualise privacy intrusions by sensors and found users are keen to learn about surrounding sensors and more details about the data being collected about them.

Koelle *et al.* [73] examined the design of privacy notices for body-worn cameras to deal with bystander concerns, exploring if designs could satisfy privacy-by-design guidelines regarding *Openness*, *Notice* and *Visibility/Transparency*, specifically conveying *Situation Awareness* and *Justification for Use*. They sought to meet two challenges: such cameras should announce themselves “in a noticeable but not too obtrusive way”, and such cameras should “publicly communicate their purpose of use to bystanders, but not impair the user’s privacy”. Across an expert design study and a UX evaluation, they found preferences towards: visible selective occlusion of lenses; conveying capture area and angle of vision of the camera; conveying device actions e.g., text/iconographic displays on the device; and supporting full visibility of the camera image and subsequent activities. They also raised transferring control to bystanders, e.g. enabling the bystander to request the camera turn off or delete captured imagery. Key resultant design considerations included improving the understandability of what the device was doing, and exploring how devices could “automatically react to privacy sensitive situations in a predictable and reliable way”. Building on this, MirrorCam [72] explored a wearable camera with an in-built mirror display to support bystander situational awareness regarding who was in-view, and whether the camera was on/off. And Eyecam [150] utilized an anthropomorphic webcam to establish a different form of relationship between the sensor and those being observed, in particular enabling the sensor to embody/convey agency and intent.

Ahmad *et al.* [9] proposed the concept of tangible privacy i.e. bolstering the *perception* of and actual effective privacy through tangible affordances to provide “a clear and definite sense of awareness of what data is being collected”. They found perception of privacy was influenced by the device type and its visibility (including the visibility of its sensors), as well as uncertainty around the device’s states and functions (e.g. is it active, could it still work when off) and available prior knowledge about the capability and capacity of a device to infringe upon their privacy. Regarding device feedback, Ahmad *et al.* noted LEDs were necessary but not sufficient for conveying awareness of device activity, and reaffirmed other research around the necessity for shutter mechanisms for cameras.

3.2 Legal and Social Protections for Bystanders

In Europe, the General Data Protection Regulation (GDPR) governs how personal data can be captured, stored and processed, with a particular focus on special categories of personal data, such as biometric and personal characteristics data—balancing “the right to be seen versus the right to be recognised” [29]. GDPR requires a “lawful basis” for processing personal data, typically through seeking consent, or determining “legitimate interest” in processing. As such, many ‘risky’ AR bystander activities discussed in this paper are not necessarily ruled out by GDPR, but rather would have to be legally justified through e.g. garnering user consent. As Koelle *et al.* [69] noted, despite the expectation of privacy-by-default introduced by GDPR, “most body-worn cameras do not provide any privacy mediating procedure. Thus, to date the *de facto* procedure is *Opt-out*, i.e., (verbally) asking the device user to turn the camera off”. In the United States, U.S. law is focused predominantly on personally identifiable information (PII) and biometric data [57], with no equivalent of the EU’s GDPR currently in place. The burden of data protection has fallen largely to “a patchwork of national- and state-level legislation addressing various

concerns” [35]. Around the world, few countries have protections of similar scope as GDPR [139]. Consequently, legal privacy protections for AR [125] will vary significantly from country to country, and even the foremost protections such as GDPR remain largely untested with respect to the multitude of processing activities discussed thus far. In many cases, “*these frameworks do not reflect the development of immersive technology when considering what features are available with hardware, how those features function, what information about users is available, and how that information could be used*” [57].

3.2.1 Consensual and Non-Consensual Erosion of Bystander Rights. Critically, where there is a mechanism for consent for lawful processing of personal data, there exists the scope for consensual erosion of existing rights—legal loopholes whereby access to/usage of an AR platform/application requires the user agree to terms of service or privacy policies which permit extensive capture and processing activities. Whilst these must be justified, nonetheless AR users (who will themselves also be bystanders to other AR users) may find themselves willingly giving permission to capture/processing activities because of the perceived low cost to themselves (e.g. not appreciating the privacy implications), balanced against the high perceived potential benefits (e.g. access to the latest AR headset and its augmented intelligence capabilities). Weighing heavily on this balance will be the capability of well resourced companies to either lobby for changes or omissions in legislation; bend the interpretation of existing legislation in their favour immediately and address the consequences later (e.g. Facebook setting aside €302 million for anticipated fines in the EU [85]); or manipulate the user into compliance (e.g. through ‘dark patterns’ [42, 90] and ‘creepy technology’ [149]).

3.2.2 Social Norms and Contextual Integrity. As AR headsets are an emerging technology, guidelines regarding ethical usage of this technology [82] (e.g. around XR privacy [92], usage [112], human rights [2], neuro-rights [161], and privacy standards [51]) are just beginning to emerge. Moreover, because of the lack of public adoption, social norms are yet to be established although studies have hinted at norms slowly emerging [129], such as removing headsets when in other people’s homes [79]. In effect, we lack a strong model for what constitutes a violation of privacy when discussing AR devices and their activity. One framework which considers this is the theory of privacy as contextual integrity [13, 104], which is concerned largely with information flow (i.e. the social **context** or backdrop upon which information flows); the **actors** involved in this flow; the **attributes** and types of data being shared; the **transmission principles** that condition this flow; and the **contextual ends, purposes and values** underwriting this flow (i.e. its purpose). Privacy norms are then assessed on “how they affect the interests of relevant parties and how they impinge on societal values, such as equality, justice, fairness and political liberties” [16].

However, everyday AR pertaining to bystanders poses challenges to the transfer of any entrenched information norms. The context in this scenario is unbounded, with interactions possible across all social relationships and strata in any place at any time. There are relatively few transmission principles applicable, outside of legally mandated protections, as the underlying data is effectively available to any camera within proximity of the bystander. And the contextual ends and purposes of using this data can vary infinitely, based on how the AR user may want to augment their sensecape, with the intent and outcomes currently entirely opaque to the bystander—the very essence of what is described as a “disruptive flow” necessitating a reconsideration of privacy norms, and the PETs we use to help facilitate them. Moreover, existing entrenched privacy norms are likely to be destabilized by the mass adoption of this technology, and the social and environmental shifts it may bring [16].

Given this, the risk is without adequate protections (e.g. mandated PETs), AR headsets will see public rejection due to their requisite sensing and the privacy implications this brings. Google Glass in particular forewarned the impact of social acceptance. It featured an obvious camera that aroused bystander concerns regarding capture [33, 58]. And, through its usage, it impaired social interaction, interrupting conversational participation [38]. This, in-turn, led to the coining of the conjunction “glassholes” to refer to the perceived obnoxiousness of its users. This setback was significant; the utility of the headset was undermined, and users not only rejected using

the headset, *they rejected the users themselves*. Given the varying legal and social protections, and the strong risk of rejection, there is a need for further research into PETs that embody “reasonable standards for transparency and choice” [34]—this paper directly contributes to this challenge.

4 AR BYSTANDER ATTITUDES SURVEY

We have established there are a breadth of AR-enabled activities that can potentially infringe upon bystander privacy (section 2), and have discussed the technical, social and legal means by which privacy can be safeguarded (section 3). PETs have predominantly considered AR privacy from the perspective of camera access control. However, such an approach is not compatible with the need for AR-requisite sensing, such as cameras for positional tracking.

Consequently, we argue there is a need for a more granular understanding regarding awareness of, and attitudes towards, potentially privacy infringing AR activities. In particular, the aim of this paper is two-fold—firstly, **to capture a snapshot of current awareness of, and concern regarding, AR headset activity**. The suggestion is the general public remain naive to the breadth of capabilities enabled by such devices, and they would be concerned by such capabilities—posing a potential barrier for adoption, and emphasising the need for AR User—Bystander PETs. And secondly, motivated by this concern, **to better understand how we can design PETs that might selectively address bystander awareness of, and consent to, AR activities** that to some degree rely on the capture and usage of their publicly available sensed personal data. Note this data is “publicly available” insofar as the average person lacks the technological means to obfuscate or prevent a suitably equipped AR device from capturing this data, ignoring legal protections that should prevent such capture, or social norms or interventions (e.g. requesting removal of the headset).

For the bystander need for consent, we focus on the influence of two parameters: the AR activity type (i.e. the outcome of the use of bystander data), and the social relationship between the AR user and bystander (equivalent to the role state in the contextual integrity framework, limited to intimacy groups). In doing so, we seek to understand whether it is possible to design minimally invasive mechanisms for seeking consent only when necessary i.e. for activities deemed problematic, or where there is not a sufficient social connection between AR user and bystander that indicates a prior basis for consent.

For the bystander need for awareness, we again consider the influence of the AR activity (i.e. do there exist such activities that, whilst they are privacy invasive in theory, in practice are perceived as not meriting informing bystanders of their occurrence); and also the influence of a prior basis for consent (i.e. if the bystander consents, do they need any on-going knowledge of the activity). But we also address what awareness users desire—that the device is active [73]; that they know the activity type and data being captured [73]; or that they have full awareness of what the device is doing with that sensed/captured data [72]. Again, we seek to understand whether it is possible to design minimally invasive mechanisms for providing activity awareness as-and-when necessary.

If we understand how attitudes vary by activity type, social relationship, and prior consent, we can propose PETs that better take into account the context of the AR activity, and are more flexible in balancing the need for bystander privacy against the want to utilize the full capabilities of AR headsets in the future. Accordingly, we established three research questions:

- **RQ1:** Is the public aware of, and concerned about, a wearable AR device’s capabilities?
- **RQ2:** How do bystanders wish to consent to AR activities pertaining to them, and is the desire to convey consent universal or contextual e.g. based on activity type (**RQ2.1**) or social relationship to the AR user (**RQ2.2**)?
- **RQ3:** How do bystanders wish to be made aware of AR activities pertaining to them, and is the desire for this awareness universal or contextual e.g. based on activity type (**RQ3.1**) or the prior basis for consent (**RQ3.2**)?

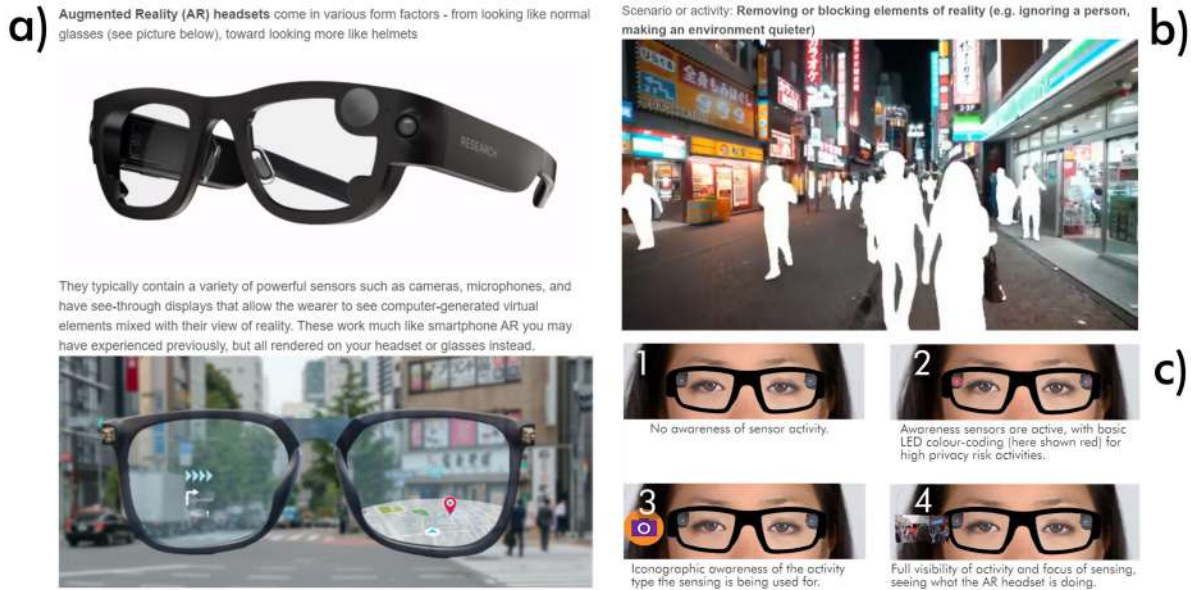


Fig. 2. **a)** Introduction to AR headsets provided to respondents, showing a front image of Facebook’s project aria to illustrate the existence of cameras, and a user perspective image of generic navigation augmentations. **b)** Still image example of one of the activities shown, diminished reality, in the survey illustrated by a video clip from @amako0609 [84]. **c)** Still image of the awareness archetypes video shown to respondents. Depicts a user wearing AR-type glasses with two in-built cameras, and shows **c1)** no awareness; **c2)** basic activity awareness through LED color; **c3)** activity awareness through a floating icon; **c4)** full awareness through a floating video feed of what the AR headset is doing. These archetypes were not intended to be practical, but rather illustrative of the information being conveyed in a way that even a respondent unfamiliar with AR technology could comprehend.

4.1 Procedure

Participants first read an information sheet informing them about survey’s data handling and procedure. Respondents were provided with a description of what AR headsets are (see Figure 2), to ensure all respondents had a baseline understanding of what we meant by AR headsets. Next, we focused on the hypothesised future of AR, presenting wearable everyday AR glasses allowing the respondents to see a view of reality intermixed with computer-generated virtual content, equipped with a variety of sensors, such as cameras and microphones. We then told respondents our focus was to investigate their attitudes towards a series of activities/scenarios that might be enabled in the near future if AR glasses became commonplace. We focus on nine bystander-pertinent AR activities based on our review. These activities were selected as they each can be tied to anticipated use cases of AR technology demonstrated in research. They span activities respondents could reasonably expect based on the current capabilities of smartphones (e.g., volumetric capture, biometric identification); activities closer to bleeding edge research that respondents would be less likely to be aware of being feasible (e.g. inferring internal state, non-contact measures of physiological state); towards activities utilizing both the AR sensing and rendering capability to fundamentally alter how/when the bystander is perceived (diminished reality, augmented appearance, augmented perception). We also include 2 baseline activities most respondents would be familiar with—general usage of a camera, and microphone, giving a total of 11 activities:

- **Camera Usage:** An application generally using the camera(s) on the device
- **Microphone Usage:** An application generally using the microphone(s) on the device

- **Volumetric Capture:** Capturing 3D imagery that could later be viewed or repurposed (e.g. a 3D model of your body or home)
- **Activity Tracking:** the physical movements, behaviour and activity of nearby people
- **Personal Characteristics:** Identifying and inferring personal characteristics of other people such as gender identity, age, race, sexuality etc.
- **Biometric ID:** Identifying who other people are (e.g. through facial ID or other biometric data)
- **Internal State:** Understanding the internal state of other people e.g. their emotions, likes, dislikes, and mental processes
- **Physiological State:** Understanding the physiological state of other people e.g. sensing health-related data such as pulse/heart rate, dilation of pupils etc.
- **Diminished Reality:** Removing or blocking elements of reality (e.g. ignoring a person, making an environment quieter)
- **Augmented Appearance:** Augmenting or altering others appearance (e.g. applying snapchat or instagram-like filters to your view of others) - pertaining to altered reality as applied to people.
- **Augmented Perception:** e.g. cancelling noise, selectively enhancing speech in a noisy room; and Super sight or other vision enhancements e.g. zooming, magnification, night/thermal vision

Each activity was represented with both a textual description (as above), and a video/pictorial description to assist in comprehension (see Figure 2 and the associated video figure for more details) and was presented in turn to the respondents. The order of the first two AR activities presented was consistent across all respondents (*Camera usage* then *Microphone usage*), so as to capture general attitudes towards sensing whilst not biasing responses based on subsequent exposure to potentially unknown camera/microphone-based AR activities. The remaining nine were presented in a randomised order. For each of these activities, we asked questions regarding:

- **Prior Awareness of Capability:** “Did you know that AR headsets have this capability?”
- **Concern Regarding Activity:** Concern as felt if a stranger were to use an AR headset to perform this activity on, or near, you; and concern relative to this activity being performed on a smartphone.
- **Attitude Towards Consent:** Opt-in by default, no consent required; opt-in by default, with ability to withdraw your consent; opt-out by default, with ability to request your consent; and opt-out. Consent was broken down by social relationship (close friend, friend, familiar stranger, stranger, stranger with accessibility needs).
- **Awareness of Activity as AR Bystander:** Broken down into four awareness archetypes (see Figure 2): no awareness; basic awareness device sensing is active; awareness of device and sensing activity type; and full awareness of what the device is doing, and explored based both on whether consent had been or had not been sought. Illustrated using video (made with Snapchat’s Lens Studio [142]) of an AR user wearing glasses that could convey this information.
- **Awareness of Activity as AR User:** Using the same awareness archetypes, describing how your headset would preferably inform others of your activity.

After respondents completed all AR activities we closed by asking two 5-point Likert scale questions to capture whether general attitudes towards AR headsets had changed as a result of the survey, and generally how comfortable they were with the possibility of AR headsets being used in public. We also provided open text fields to allow respondents to provide comments regarding each activity, as well as any final comments they had regarding the survey. The survey took under 15 minutes to complete, see <https://doi.org/10.5281/zenodo.7244156> for an archived version of the survey questions. Prior to its distribution the survey was reviewed and approved by our institution’s ethics committee.

4.2 Demographics

We distributed the survey through mailing lists and social media. Respondents who completed the full survey were given the chance to enter a prize draw to win an online voucher (a £30 (or local equivalent) Amazon voucher). To strengthen ecological validity, we advertised the questionnaire on a variety of different platforms. These were relevant XR-related subreddits, XR discord groups, XR mailing lists, Facebook groups and Twitter (with relevant hashtags). 102 respondents (self-identified as 40% female, 58% male, 2% non-binary / third gender) completed the

survey. The age ranges of respondents were: 42% 18-24, 38% 25-34, 13% 35-44, 6% 45-54, and 2% 55 and above. The majority of respondents were from the UK (56%) and United States of America (22%).

4.3 Limitations

4.3.1 Experience with AR. Prior work has indicated that the public has an incomplete understanding of AR's capabilities [151] and is still relatively unfamiliar with what AR technologies are conceptually and what they are capable of [152]. Therefore, we opted to advertise our survey primarily on XR-related channels to ensure respondents had some baseline awareness of what AR is conceptually. While this meant our respondents likely had more awareness of AR activities/capabilities than the general public, it ensured they would understand the described technologies/scenarios in our survey and could more confidentially discuss how they envisioned future iterations of AR technologies would work.

To ensure respondents had a clear understanding of what constitutes an AR headset and the described activity scenarios we provided a detailed textual description and a video/pictorial description both for AR headsets (at the beginning of our survey) and for every AR activity scenario. Whilst a survey does not provide hands-on experience of the AR activities, this approach was chosen so we could capture representative data from a large sample, with prior research suggesting that members of the public can reasonably reflect on ethical challenges given adequate presentation of the concerns [159].

4.3.2 Scope of AR Activities. We limit our survey to 11 AR activities of interest, including two effective baselines regarding general attitudes to cameras and microphones. The selected activities are well motivated by our review—all have been demonstrated as feasible by prior research and are within the scope of privacy challenges discussed by McGill *et al.* [92], meaning bystanders are exposed to significant risks. This list is not exhaustive, nor can it be given the continual emergence of AR-enabled capabilities. However, it is illustrative of a wide range of activities that go beyond general use of the camera or microphones for recording activity, and is thus sufficient to answer our intended RQs.

4.3.3 AR Activity Interpretation. Although we included both a textual and video/pictorial description for each described AR activity, respondents' interpretation of a described activity may have differed from our intended presentation. For example, for our "Personal Characteristics" activity a respondent may have focused on a particular instance of this (e.g. identifying/infering gender identity) rather than activity as a whole. Steps were taken, however, within the survey's design to reduce the risk of this occurring, e.g. activity descriptions were written to be illustrative in nature and an open text field included for each activity to allow respondents to comment on individual problematic instances for a given activity.

4.3.4 Privacy Attitudes. We specifically asked participants about different privacy-related aspects of AR. Since the study relied on self-report, participants might have reported a behaviour that might not match their actions in reality. This is explained by the so-called privacy paradox [47]. To validate our findings, future work should investigate real usage. However, since AR headsets are rather scarce, our investigation should be seen as an important first step.

4.3.5 Demographics. The demographics of our survey responses skewed towards western respondents under the age of 34. It is reasonable, however, to expect both cultural and age-based effects to have some had influence on our results. Consequentially, it is therefore necessary for future research to investigate demographics (both cultural and generational) where expectations and attitudes may differ significantly. To assist further research in this endeavour, we make available the full question set and anonymized data set for further analysis.

5 RESULTS

For statistical significance testing an Aligned-Rank Transform (ART) [40] was used to transform non-parametric data prior to conducting a repeated measures ANOVA, using the *ARTool* R package [65]. ART enables the use of parametric tests on non-parametric data (e.g. Likert-type responses [132], preference tallies). When using ART, as noted by Wobbrock *et al.* [157] “the response variable may be continuous or ordinal, and is not required to be normally distributed”, making it well suited to our dataset. Where two factors of concern existed, a two-way ANOVA was conducted, again using ARTool. Where feasible to report, pairwise contrasts for main effects were also conducted [66]. However, given the number of pairwise comparisons available (particularly where two factors exist), we do not report contrasts for all results. In lieu of this, all plots show 95% confidence intervals (visualized with red bars) based on conversion of dependent interval/ordinal variables to numeric ranks, allowing a by-eye estimation of significant pairwise differences (where the confidence intervals do not overlap). This approach is favoured by those in HCI that believe reporting should be done with confidence intervals and visualisations [37]. Respondent qualitative answers were coded using initial coding [26] where respondents’ statements were assigned emergent codes over repeated cycles with the codes grouped using a thematic approach. A single coder performed the coding and reviewed the coding with one other researcher to resolve unclear codes and discuss the depth and specificity of codes. See <https://doi.org/10.5281/zenodo.7244156> for the survey dataset [107].

5.1 Participant Prior Experience

Regarding their understanding, and prior experience, of AR technology, as it can be seen in Figure 3, respondents were largely familiar with AR as a technology, but their experiences were predominantly around smartphone-based AR.




Fig. 3. Respondents were asked (left) if they had used AR before (Yes, but only smartphone AR (e.g. Instagram filters, Snapchat lenses, IKEA furniture app), Yes including AR headsets, No, I don’t know); and (right) whether they knew what an AR headset was.

5.2 Knowledge of AR Headset Capabilities

There was a significant effect on *AR Activity* ($F=48.6$, $p<0.01$). Outside of *Camera/Microphone Usage*, *Biometric ID*, and *Augmented Appearance*, respondents were predominantly unaware of key AR Activities (Figure 4). In particular, over 50% of respondents had no awareness of AR headset’s potential capacity to identify *Internal State*, *Physiological State*, nor diminish (*Diminished Reality*) or enhance (*Augmented Perception*) the user’s perception of reality. This is despite the fact that many of our respondents were familiar with AR (see Figure 3), which further affirms that there is low awareness of sophisticated activities of AR headsets.

5.3 Managing Consent Towards AR Activity

For each AR Activity, we questioned respondents regarding their attitudes towards consent to this activity as a bystander (Figure 5). A two-way ANOVA found significant effects both on *AR Activity* ($F=83.6$, $p<0.01$) and *Relationship* ($F=170.2$, $p<0.01$), with no interaction ($F=1.2$, $p=0.16$). Regarding *Relationship*, we found significant

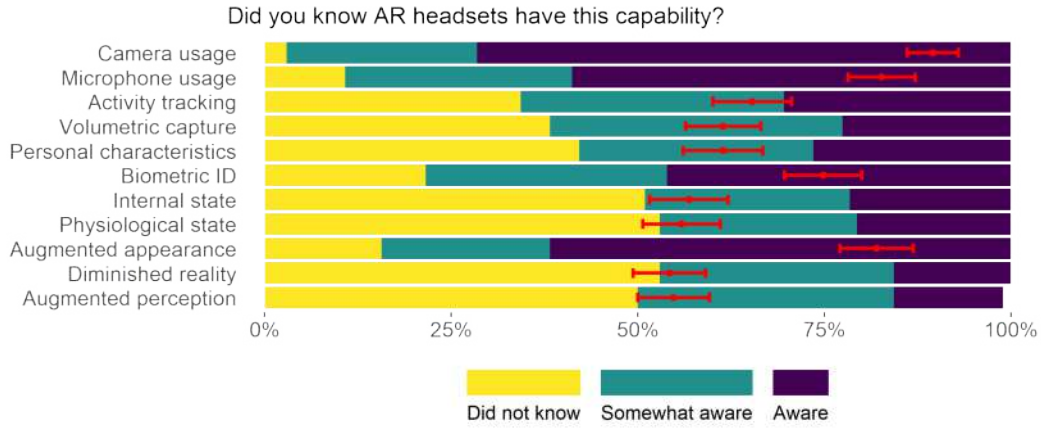


Fig. 4. Responses to “Did you know that AR headsets have this capability?” – outside of *Camera/Microphone Usage*, *Augmented Appearance*, and *Biometric ID*, respondents were largely unaware of key AR capabilities.




Fig. 5. Responses to our question on consent, broken down by activity type and social relationship.

effects ($p < 0.01$) between all pairwise contrasts aside from *Close friend—Friend*, with preferences towards opt-out increasing both based on the strength of the relationship to the AR user and based on the perceived accessibility needs of the AR user. Attitudes were largely split regarding the ability to request or withdraw consent within the opt-in and opt-out options. However, when treated in combination, significant portions of respondents (>60%) were varying in favour of requesting or being able to withdraw consent—suggesting a strong willingness to manage consent beyond opt-in/out alone. With regards to AR activity, *Volumetric Capture*, *Internal State* and *Physiological State* featured 80%+ weightings towards requiring opt-out when strangers were performing a said activity. *Diminished Reality* consistently had >50% reporting for opt-in regardless of relationship to the AR user, and *Augmented Appearance* featured opt-in rates of approaching 70% when dealing with friends.

We also asked participants to consider whether, as a bystander, they would want an XR headset to automatically opt them in/out of AR activities based on their pre-provided AR privacy preferences (see Q26 in [107]). In this regard, we found a significant effect on AR Activity ($F = 2.52$, $p < 0.01$), with contrasts suggesting differences regarding *Augmented Appearance* – {*Augmented Perception*, *Physiological State*, *Volumetric Capture*}. Respondents were largely favourable (typically >70%) toward pre-providing AR preferences (see Figure 6).

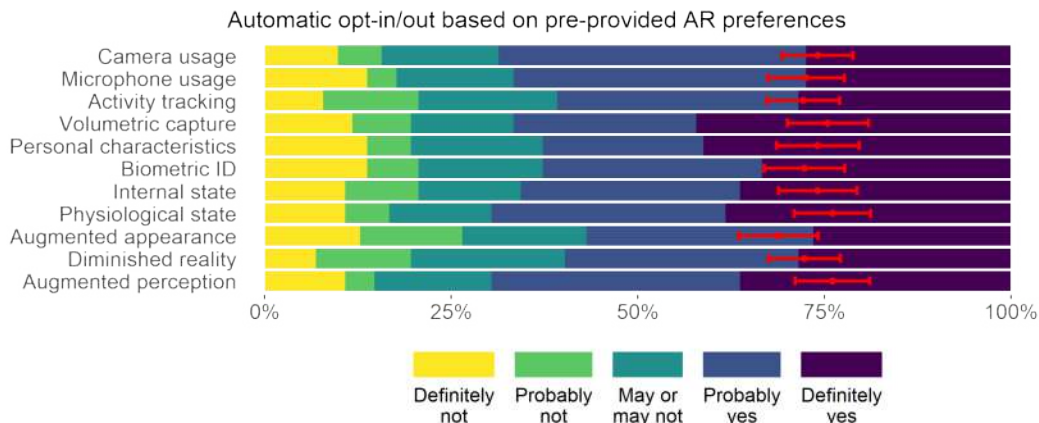


Fig. 6. Desire for automatic opt-in/out consent based on pre-provided preferences, broken down by activity.

17 comments discussing consent were made across 8 AR Activities (*Augmented Perception* 5, *Physiological State* 3, *Activity Tracking* 3, *Personal Characteristics* 2, *Camera Usage* 1, *Internal State* 1, *Biometric ID* 1, *Diminished Reality* 1) by our respondents. The majority of these (11 comments) focused on who respondents considered appropriate to access their information. Eight comments said access should be restricted to only trusted individuals: e.g., P11 (Male, 18-24, United States of America), “only trusted people should have access”. Whereas, three respondents said they considered it appropriate for anyone with accessibility needs to have access to their information, P101 (Female, 25-34, Indonesia): “I am okay to share my movement data to someone with visual disability”. The sentiment of these comments is similar to prior work by Lee et al., who found individuals were willing to provide access to their personal information to assist visually impaired individuals but wary of giving it to those without a need [80]. Interestingly, for *Personal Characteristics* and *Physiological State*, five respondents indicated that trusted persons were specifically police officers and medical professionals, P58 (Female, 25-34, Germany): “good to be used by policemen, ambulance [staff]”. Respondents said this was because individuals in these profession would require access to this information as part of their job, whereas access by a lay-user was felt to be unnecessary,

P66 (Female, 25-34, United Kingdom): “I don’t see a benefit for me personally if someone else [a lay user] knows my physiological signals”.

Finally, one comment outlined the difficulty in obtaining consent, P3 (Female, 25-34, United Kingdom): “seems hard to get consent from people when the people affected could be far away (i.e. could use super hearing to zoom in on people far away and see what they’re doing)”. This was similar to two additional comments made regarding consent at the end of our survey (where we asked if respondents had any final comments they wished to make). Both focused on respondents uncertainty over how their consent could be collected, P22 (Male, 25-34, United Kingdom): “Not sure how consent for many of these activities could be gathered. But I still feel that it should be gathered in some way”.

5.4 Need for Awareness of AR Activity

We asked participants to assume the role of a bystander and report their preference in terms of how a stranger’s AR headset should provide awareness of its activity, in two situations: (1) when they had provided prior consent, and (2) when they had not provided prior consent (see Q30 in [107]). We also asked participants to assume the role of the AR user and report their preference in terms of how their headset should provide awareness of its activity to others (see Q32 in [107]), resulting in three perspectives seen in Figure 7.

We found significant effects regarding AR Activity ($F=32.1, p<0.01$) and perspective ($F=76.9, p<0.01$), with contrasts showing differences between *Bystander without consent* and *Bystander with consent/As the AR user*, as well as an interaction effect ($F=1.9, p<0.01$). Respondents perceived a more substantial need to be made aware of AR activity as a bystander whose consent had not been sought prior. The effect was particularly strong for *Volumetric Capture, Personal Characteristics, Biometric ID, Internal State, Physiological State, and Augmented Perception*, where nearly half the participants indicated that full, real-time awareness of the AR user’s activity was required.

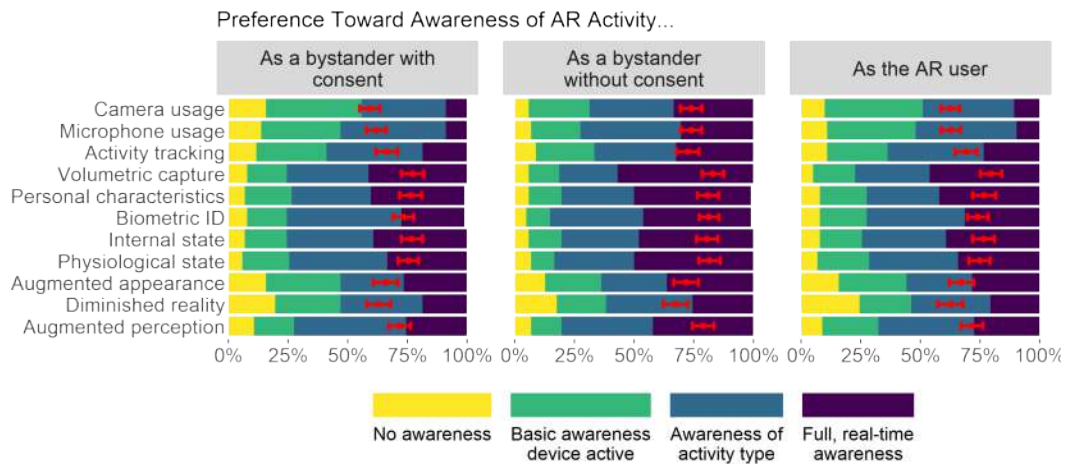


Fig. 7. There is a stronger perceived need for awareness when consent had not been obtained, and preference towards facilitating awareness of the activity type. Without consent, respondents demanded a high degree of oversight, with approaching 50% suggesting that full, real-time awareness of the AR user’s activity was necessary.

5.5 Concerns Regarding AR Activities

To add context to findings around awareness and consent, we wished to establish a baseline for how concerned prospective bystanders were about the proposed AR activities. We asked this from the perspective of what is arguably the worst case—that the capability is being used by a stranger. We also explored the extent to which everyday AR headsets were more or less concerning than the same activities conducted by a stranger using a smartphone i.e. in short bursts and in a more visible manner. The results are presented in Figure 8.

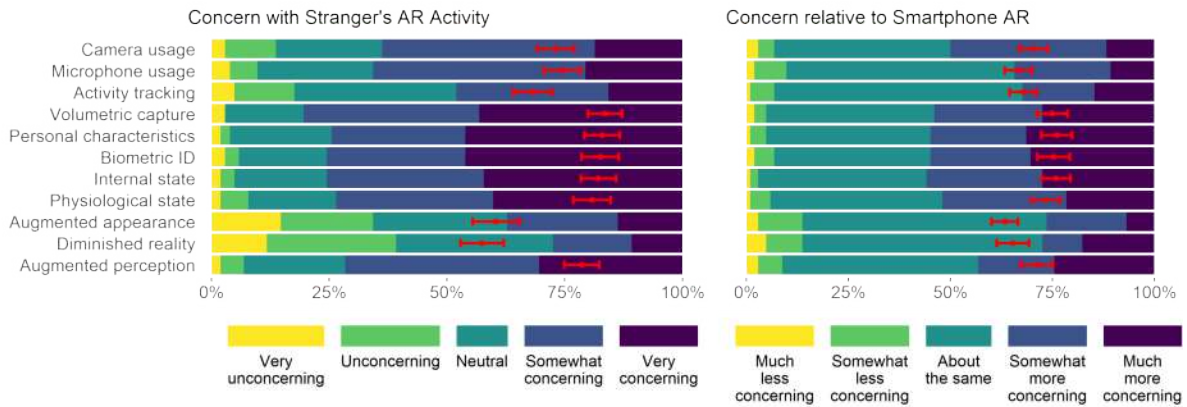


Fig. 8. Perceived concern with a stranger's activity, and relative to smartphone AR generally.

5.5.1 Concern with a Stranger's Activity (left of Figure 8). There was a significant effect on AR Activity ($F=29.7$, $p<0.01$), with obvious high levels of concern featured particularly for *Volumetric Capture*, *Personal Characteristics*, *Biometric ID*, *Internal State*, *Physiological State*, and *Augmented Perception*, with approximately 75% of respondents suggesting these activities were at least somewhat concerning. Interestingly, *Augmented Appearance* and *Diminished Reality* both featured largely neutral or unconcerned responses, despite their relatively significant impact in how the bystander is perceived by the user. These findings may reflect the familiarity with augmented appearance gained through the use of apps including Snapchat and Instagram, as well as an underlying attitude toward accepting augmentations that do not pertain to personal data.

74 comments across ten AR Activities (*Internal State* 11, *Personal Characteristics* 10, *Camera Usage* 8, *Diminished Reality* 7, *Biometric ID* 7, *Physiological State* 7, *Volumetric Capture* 7, *Augmented Appearance* 6, *Augmented Perception* 6, *Microphone Usage* 5) involved some mention of concern or discomfort with the AR Activity described. Notably, *Activity Tracking* was the only activity to receive no comments from any respondent—perhaps due to the already widespread usage of similar technology in devices such as CCTV cameras, or the lack of an obvious means to exploit this data to the detriment of the bystander.

Malicious applications: Eight respondents said they were uncomfortable with their personal data being captured and used by the AR user, P69 (Female, 35-44, United Kingdom): “I am horrified by the intrusion”. Whereas another eight were respondents stating their discomfort was due to uncertainty surrounding how their personal data was being processed, P18 (Male, 25-34, United States of America): “my level of concern would change if the footage captured were to just be stored locally” and were concerned that their data might be manipulated or used out of context, P29 (Male, 25-34, United Kingdom): “it is far easier for observers to decontextualise voice recordings [for malicious purposes]”. However, nineteen comments concerned respondents acknowledging they saw some beneficial applications for the particular activity described, but always with the caveat that they envisioned

malicious applications as well, e.g. for *Augmented Appearance*, P99 (Female, 25-34, United Kingdom): “This one could be fun, but some augmentations could be really bad e.g. blackface”.

Safety: Five respondents were concerned about the safety of the user or bystander, P69 (Female, 35-44, United Kingdom): “what if they [the AR user] are blocking awareness of cars and step in front of mine while I am driving!”, and noted that additional safety measures would be required, P102 (Male, 18-24, United Kingdom): “there should be regulations in place to avoid hazards from unawareness or disregard for what’s transpiring in an individual’s environment”. Five others were respondents who believed the technology would facilitate predatory behaviour, P93 (Male, 18-24, United Kingdom): “this technology would allow predatory behaviour to be carried out far too easily”. Six respondents, on the other hand, had general negative sentiment towards these technologies, for instance, P14 (Female, 45-54, United States of America): “This whole scenario is a dystopian privacy nightmare”.

Personal data and trust: *Personal Characteristics*, *Internal State* and *Biometric ID* were the three AR Activities which respondents were most concerned of, with thirteen comments singling them out to express how dangerous respondents considered them, P69 (Female, 35-44, United Kingdom): “This is really scary from a civil liberties point of view”, with many stating they felt these particular activities should be prohibited, P72 (Male, 25-34, Germany): “it should be illegal”. Respondents made ten further comments discussing their distrust and discomfort with these AR Activities. Three comments contained respondents stating they were uncomfortable with the use of these activities (by either a system or an individual), P69 (Female, 35-44, United Kingdom): “I would be pretty horrified to find out that a friend was using this”. Seven featured respondents being critical of the activities and questioning their validity, P65 (Male, 25-34, United Kingdom): “There is very limited evidence that these models are correct”, and respondents referenced known examples of relevant technical failures, often highlighted in the media [102, 103, 154], to motivate their distrust and concerns, P96 (Female, 35-44, United Kingdom): “Given algorithmic difficulties with even spotting some people’s faces (such as Google’s humiliating gorilla incident) it is hard to believe that this would be a sufficiently inclusive feature”. For these respondents, using systems which were, as our respondents put it, P18 (Male, 25-34, United States of America): “high tech phrenology”, and are, P72 (Male, 25-34, Germany): “prone to bias in the training data and other drawbacks”, is a dangerous proposition and one that could lead to unforeseen consequences, or even discrimination and segregation within society due to malicious actors.

5.5.2 *Concern Relative to Smartphone AR (right of Figure 8)*. There was a significant effect ($F=11.2$, $p<0.01$) with approximately 50% of respondents noting that headset-based AR was more concerning than smartphone AR for *Volumetric Capture*, *Personal Characteristics*, *Biometric ID*, *Internal State*, *Physiological State*, and *Augmented Perception*. This further suggests that AR headset users will see continuing challenges with respect to how they are perceived by others. Seven comments justified why respondents felt AR headsets were more concerning than smartphones. All indicated that a smartphone AR user was thought to be easier to notice than a headset user, P99 (Female, 25-34, United Kingdom): “With a phone I can see someone holding it up, with a headset this can blend”, believing the smartphone user’s actions would be more explicit when interacting with AR. Headsets meanwhile, particularly if designed with an every-day, wearable, form factor, were thought to more easily “blend” and go unnoticed, possibly even being mistaken for non-AR eyewear.

5.6 Retrospective Comfort with AR Headsets in Public

Approximately 55% of respondents were now more uncomfortable with the prospect of everyday headset-driven AR than they were prior to the survey, and overall nearly 70% of respondents were at least somewhat uncomfortable with the prospect of this technology seeing everyday use in public. At the end of the survey we asked respondents to reflect on their general attitude towards AR—had these changed since the start of the survey

and how comfortable were they now with the prospect of everyday AR having seen the kinds of activities it could potentially support (Figure 9).




Fig. 9. Closing questions of survey, addressing whether respondents were more or less comfortable with AR having been exposed to what it could potentially enable (left), and their overall comfort with the prospect of everyday wearable AR technology (right).

5.7 Summary of Research Question Outcomes

5.7.1 RQ1 - Public Awareness of, and Concern Toward, AR Headsets. Our results expand our understanding of public awareness of, and concern toward, headset-based AR activity. Respondents, including users of smartphone AR and AR headsets, were frequently unaware of key ‘risky’ activities pertaining to bystander privacy, in particular with over half being unaware of the potential for unlocking insights into our affective, cognitive and physiological states. That AR has a unique potential capacity to compromise our individual privacy in this way is not yet fully understood by the public. Indeed, it may be such risks never become part of the public consciousness regarding AR - to be aware of non-contact physiological sensing for example would effectively require every member of the public be a security researcher.

Existing technologies (e.g. smartphones) already offer insights into how such risks may be overlooked by the public. While most individuals are likely aware of some risks associated with the microphone/video recording capabilities of smartphones (e.g. non-consensual audio/video recordings, photo manipulation) many are likely unaware of the risks associated with analysis/processing of captured data. For example, prior works have shown the majority of individuals are unaware of what personal information can be inferred from voice recording [74, 78] or sensory data captured by a smartphone [27, 95]. As in our survey results, while individuals are aware cameras/microphones pose a risk they are often unaware of specific instances of abuse, limiting the extent awareness of these risks enters the public consciousness. Furthermore, existing technologies highlight the difficulty in addressing/combating abusive scenarios once they emerge and are adopted by individuals. For example, as camera-equipped smartphones became more widespread the rise in access to portable, discreet, cameras led to the emergence of “upskirting/creepshots” [148, 153] (non-consensual photos of an individual often focused on sexualized areas). However, responsive actions to prevent this behaviour are slow with laws taking years to introduce [15], technological solutions being limited [143] and despite steps taken to address the issue the problem remains ever present [97].

Nevertheless, a lack of awareness is what underlines the risk we, as consumers, will sleepwalk into the adoption of AR headsets, and their requisite sensing, without full/due consideration of how we should limit the capabilities of these headsets. However, once exposed, our respondents showed clear, elevated concern regarding the use of AR headsets by strangers in public, in particular with respect to activities that revealed insights into who we are (biometric ID, personal characteristics), what we feel (affective/physiological states), or those that could introduce security risks (volumetric capture, augmented perception). These findings are not surprising, building upon past insights regarding wearable cameras [72] amongst others, but crucially they are confirmatory. *The public does not have sufficient awareness of the risks posed by AR headsets beyond the basic presence of cameras and microphones.* And these risks, even when impartially presented as benign activities, result in significant, and in many cases

justified, concern to future AR bystanders. This emphasises the fundamental need for bystander-oriented PETs, if we are to avoid societal rejection of AR headsets.

5.7.2 RQ2 - The Desire for Consent. Respondents opt-in/out preference varied significantly, influenced by activity type and relationship to AR user. Regardless of the fact these activities were predominantly camera-based, the intent behind the use of the sensing, and extent to which the sensing seemingly pierced the veil of bystander privacy, significantly altered attitudes towards consent (**RQ2.1**). In particular, activities enacted by strangers (**RQ2.2**), and those with the capacity to capture personal information, predominantly featured bias toward opt-out by default, whilst some activities involving close friends (e.g. *Augmented Appearance*), or those pertaining to control of an individual's sensecape (e.g. *Diminished Reality*), exhibited a bias towards opt-in by default. Moreover, there was a strong perceived need for AR headsets to be able to request consent or for bystanders to be able to revoke consent. Moreover, few respondents selected opt-in by default with no additional control - despite often being the default mode of operation for many AR devices.

5.7.3 Implications for Consent PETs. Our findings affirm the importance of considering the contextual ends and purposes (as per contextual integrity conceptualisations) around the usage of bystander data in AR. For many of the activities presented, the fundamental "information type" was that of RGBD camera data - however, the processing enacted on this data, and activities derived from it, significantly impacted bystander attitudes towards opt-in/out consent. Privacy around AR devices in particular is infringed upon not just by the data sensed, but in how that data is further processed and used. AR supports privacy infringements which can then lead to forms of misuse or abuse that may not be apparent to the bystander (e.g. inferring traits/characteristics the bystander was unaware they were exposing, or augmenting, altering or appropriating appearance based on body tracking and biometric ID). That privacy norms do not yet exist around these activities exacerbates concerns. This outcome-oriented perspective is often overlooked in Computer Science-grounded privacy contributions [16], and emphasizes AR privacy cannot be reduced to discussions around access control of sensing alone. Indeed, there are activities that are broadly permissible by bystanders based on the context. For example, augmented appearance featured a majority of respondents choosing to opt-in when the AR user was their friend, provided they could to withdraw consent. This offers the possibility we can design context-dependent minimally invasive consent PETs that effectively require no action from bystanders unless they feel it necessary to withdraw consent. However, such an approach is strongly contingent on their awareness of the given AR activity.

5.7.4 RQ3 - The Desire for Awareness. Respondents showed strong preferences towards awareness of the AR headset's activity. In particular, over half of respondents typically desired awareness that went beyond archetypal 'device active' features such as LEDs, desiring awareness of either the type of activity, or in particularly problematic cases (e.g. *Volumetric Capture*, *Personal Characteristics*) full real-time awareness of what the AR headset was doing (**RQ3.1**). This desire for awareness also varied based on consent (**RQ3.2**) - where there is no prior consent, awareness was more greatly desired. Finally, respondents did not exhibit signs of prioritising their privacy over others - with attitudes towards awareness being mirrored between the bystander and AR user perspectives.

5.7.5 Implications for Awareness PETs. Contextual awareness and information overload—The need for awareness (of activity type and the activity itself) expands upon Koelle's research in this area [73], and underlines generic sensor activity awareness (e.g. LED indicators) is not necessarily sufficient. This suggests for problematic activities (i.e. those perceived as privacy infringing) bystanders need further understanding - of how their data is used, and whether it is retained. This need for awareness is also seen as contextually driven, varying based on activity, relationship/social sphere, and prior consent. This is important because if we are to envisage an everyday AR scenario where every device around us is constantly trying to relay full awareness of its activity, we would likely incur an information overload, undermining this awareness feedback, leading to rejection of such mechanisms. Our work suggests we can design context-appropriate minimally invasive awareness cues and

often awareness of device activity or activity type alone may be sufficient. The challenge is in understanding what these minimally invasive awareness cues might be e.g. can AR activities be categorised in a simplified way (e.g. split into bystander augmentations, personal data, movement data, and 3D capture) such that bystanders can clearly/quickly understand what is occurring, and the risk toward them?

Awareness/notice underpins consent— Awareness and consent are strongly interlinked. For a user to give or revoke consent as the activity occurs or is about to occur, they must be notified of, and understand, what it is they are consenting to as a pre-requisite. This dependency complicates the design of awareness and consent PETs, as it asks the question: what degree of activity awareness is necessary to facilitate genuinely informed consent.

Consent influences the need for awareness— However, the prior basis for consent also influences the need for awareness - as our results show, with consent obtained, awareness is less desired. What our results did not capture was the impact of revoking consent, but we would suggest if consent is not obtained, no awareness is necessary as the device should be respecting the bystander's choice, assuming the bystander trusts the AR device to do so. In this way, we move closer to minimally invasive awareness mechanisms, that are perhaps only active for as long as is necessary to enable the bystander to manage their consent to an activity.

6 DISCUSSION

The capabilities of a given sensor (be it a camera, microphone, etc.) are not static. Multi-disciplinary research is continually uncovering new insights into existing sensing, driven by machine learning and cloud computing. That these processing activities are often imperfect, subject to algorithmic bias and approaches which at-times border on phrenology as one respondent noted, does not dispel the notion they will nonetheless be one day deployed to consumer AR devices. That AR places this sensing on the vantage point of the user's head, continually surveying every aspect of their life and the lives of those they encounter, merely amplifies the risks posed in this processing arms race. We have outlined why we believe cameras and microphones will be considered *requisite* sensing for consumer AR. Because of this, there is a strong possibility prior research regarding active shutters and other such occlusive PETs may be fundamentally incompatible with consumer AR. Access control alone (e.g. preventing access to camera data through API permissions, or preventing camera usage through physical shutters) is insufficient as access to a given sensor can be used for many different activities, of varying concern to bystanders. That these risks will be further amplified in time (e.g. through the potential for distributed mass surveillance) if left unchecked suggests the need for privacy enhancing technology is increasing as we move closer to everyday consumer AR headsets [116, 144].

Consequently, PET designs are needed that both assume the constant, continual presence of cameras and other sensing, and consider the need for bystander agency regarding not just how they are sensed, but in what way this sensed data is processed and acted upon. It is insufficient to know the camera is active - the bystander requires agency over how an application uses that camera in relation to them, and their 'exposed' personal data. Our findings suggest this agency must come in terms of incorporating bystander consent into the act of generating this processed data; and in providing bystanders with specific awareness of activities pertaining to them.

6.1 Bystander Consent to AR Activities

Regarding consent, the challenge is in defining interactions and mechanisms for obtaining or revoking consent automatically based on some heuristic (e.g. taking into account the activity, immediate context, relationship, willingness, contextual integrity); or facilitating informed consent.

6.1.1 Contextual Permission and Automatic Determination of Consent. For inferring or otherwise automatically determining consent, there are two clear approaches. The first is for the AR system to employ context-aware privacy measures (e.g. [20]) to determine if the usage of bystander data is permissible. However, there are significant challenges in determining by what measures/contextual information an AR activity is deemed permissible.

Social/cultural norms around AR usage are yet to emerge, nor is there a complete understanding of how context in the ubiquitous computing sense (activity, location, social connections, etc.) should be considered - a particular problem for contextual integrity approaches given “many systems and devices span multiple contexts” [16]. Or will the bar instead be set by what is legally permissible? In the rush to market, it is probable the latter will be the priority for many AR headset platforms. Absent entrenched or transferable privacy norms around this technology, and without sufficient awareness of what these devices are doing, or agency over their activity, bystanders may find violations of privacy even in benign circumstances.

6.1.2 Personal Privacy Preferences. The second approach is to give bystanders some means of asynchronously conveying their preferences towards activity types, by defining privacy preferences, to enhance or support any context-based privacy implementation. However, there are practical issues to accomplishing this. How might we summarise the breadth of activities possible in a way users can meaningfully, and easily, record their preferences? And given those preferences, how do we convey them to the AR device and ensure and trust they are taken into account? For example, preferences could be retrieved based on wearable tokens [115], or via platform-level biometric ID and cloud-backed preferences and knowledge regarding social connections. This may sound like an ideal solution for companies like Meta/Facebook, already owners of vast amounts of personal/social data that are pursuing AR-driven technology - but arguably would itself introduce significant privacy concerns. A counter approach could be to develop P2P architectures for conveying said information. Such an endeavour is a significant technical challenge, but feasible given existing P2P technologies such as WiFi direct coupled with, say, optical tracking of infrared LEDs conveying a temporary user ID for messaging - again, provided it could avoid introducing new privacy risks. Regardless, we posit further research is required into how we can facilitate privacy-enhancing back-channel communications between AR users and bystanders.

6.1.3 Informed Consent. The alternative is bystanders have sufficient awareness of AR activities they can actively and manually provide, or withdraw, informed consent as the activity occurs (e.g. via gesture [69]). Such an approach depends heavily on what information is provided to the bystander however, moderated by whatever AR activity awareness cues are shared by the AR device. Consequently, achieving informed consent may be difficult [62]. What is it the bystander is consenting to? Is it usage of particular sensing, or the capture of particular data? Or is it the activity utilises the data? Based on our results, we would argue the latter - bystanders must be able to consent not just to what data is captured, but *how* that data is used and retained within the scope of the AR device (assuming it can be sufficiently restricted).

Complicating such matters is the added workload and expectations placed on bystanders to actively manage their consent - something bystanders may be reticent to do. Consideration should be given to consent PETs that achieve the underlying aim of protecting bystander privacy (and preventing resultant harms) in a way that is usable, feasible, and is likely to see adoption (e.g. prioritising protecting bystanders in the most important contexts). Whilst ambitious, an integrative combination of contextual permissions, taking into account shared privacy preferences if available, with the option for informed withdrawal or granting of consent, could offer a largely automated means by which bystanders could still have a degree of agency over how they are sensed/augmented.

6.2 Bystander Awareness of AR Activities

For activity awareness, we can see immediate parallels with research such as FaceDisplay [49], where a future AR headset might incorporate a transparent display over the lens of the camera, providing a visual sense of what the user is doing [111]. Such feedback could be mandated at a platform level, reporting activities based on the requested APIs an application uses (e.g. conveying to a bystander when volumetric capture or biometric ID are occurring). But even such measures are imperfect. For example, conveying this insight to a visually impaired bystander might require auditory feedback. There are also risks of habituation and the ‘crying wolf’ problem if we

constantly warn users about potentially 'risky' AR activities. Bullying users into protecting their privacy will not work as shown by usable security researchers [130] – It is important to employ usable and user-centered measures when designing approaches for raising awareness and obtaining consent from AR bystanders [8]. Consequently, how we design these mechanisms in ways that are accessible, usable, and avoid visual or multi-sensory clutter and information overload, remains an open question.

6.3 Open Challenges in AR Bystander Privacy

6.3.1 Awareness and Consent Versus State/Societal Needs and Legitimate Interests. There are also clearly legitimate scenarios, or at least legally justifiable, reasons for conducting bystander AR activities. For example, the medically-oriented AR headset with non-contact physiological sensing raised by some respondents. This will unlikely ever be a feasible medical-grade reality, but nonetheless raises questions regarding when legitimate interests and justifications exist, beyond the rationale of supporting accessibility needs we briefly examined. Consumer AR offers a potential boon to states/societies pre-disposed to distributed or mass surveillance [144]. To what extent would an AR user for example be comfortable with their device being co-opted by the police or local authorities temporarily to sense bystanders or gain situational awareness about an on-going incident? In such cases, questions of bystander awareness and consent are perhaps supplanted by the questions regarding AR user awareness and consent, and the degree to which an individual should retain agency over their wearable technology and the data it generates.

6.3.2 The Line in the Sand: The Counter-Argument to Awareness and Consent. There are, however, arguments against supporting awareness and consent. For example, what of vulnerable people or digital 'have-nots' that are not able to specify their preferences towards consent, or understand what they are being made aware of? Moreover, are the most risky activities we have discussed thus far even necessary to consider? For example, if we as a society ban, or otherwise legislate or limit the capability for at-a-distance physiological sensing (e.g. to medical AR headsets only) through responsible innovation, then there is no need to institute such protective measures. However, such an argument assumes we can as a society internationally agree a common and allowable set of functionalities for consumer AR, and uphold this standard such that e.g., Apple, Facebook/Meta, Microsoft, Google, Xiaomi all abide by these requirements in all territories. Such an argument also precludes the influence of platforms and other parties in persuading society a given capability should be allowable (e.g. see Meta's discussions regarding AR facial ID [53]). Whilst the line in the sand argument is tempting in its simplicity, easing the path for consumer deployment and adoption, our belief is if we as a community accept this as our first line of defence against potentially abusive bystander-AR interactions, we will see a slow, consensual erosion of bystander privacy that may be difficult to arrest.

6.3.3 Contextual Integrity, Everyday AR and IoT. More broadly, further consideration must be given to how contextual integrity and other privacy frameworks map to everyday AR, and privacy concerns raised around bystander data in particular. As discussed, when considering contextual integrity and information flow, AR poses a number of challenges:

- The social context is effectively unbounded - flow can occur anywhere, anytime between anyone;
- The type of information transferred is entirely under the control of the AR user, not the bystander i.e. the "sender" does not have any agency in sending this data, the recipient has complete control;
- The attributes/type of data is malleable, based on further and longitudinal processing and inference applied to the raw sensor data;
- There are few transmission principles applied to this flow of information (outside of existing legal protections), as the bystander information is freely available to any device with a camera or microphone and sufficient software processing capability;

- And the occurrence of this information flow, and the resultant output of this activity, are entirely opaque and hidden from the bystander.

It should be noted that many of these points are shared with other smart technologies and pervasive sensing. For example, IoT [87], mobile device sensing [31] and other wearable sensing [118] all have the potential to expose similar data risks. However, everyday AR also poses additional concerns in *how* this data is actively used/presented in relation to the bystander, such as augmenting bystanders to reveal sensed insights in real-time to the AR user. Moreover, the likely impending adoption of everyday AR pushed for by industry hastens the need for PETs that address these points. Our focus has been on giving bystanders agency over activities pertaining to them, but there are other promising routes, for example instigating privacy policies to restrict what data and activities are available at a platform level for the most important or risky contexts, such as limiting what an AR headset can do in public spaces, to mitigate against misuse and abuse of bystander data [135].

6.3.4 Designing Privacy-Respecting 'Requisite' AR Headset Sensing. There is also the possibility AR headsets could be designed that entirely physically preclude supporting any of the activities discussed herein. Consider an AR headset without what we argue are currently 'requisite' sensors, cameras and microphones. Instead, other forms of sensing could be used to enable localization/positional tracking (e.g. using advanced sensor-fused forms of Global Navigation Satellite Systems [94]). Such a device would have obvious limitations (e.g. potentially lacking the ability to be visually/aurally context aware), but would still be able to provide exocentric spatial AR experiences, and context-aware augmentations (insofar as the context could be derived from mapping data, for example). Whilst such a device would be largely against the interests of the companies developing AR technology, we suggest there is merit in considering the extent we *need* cameras/microphones in everyday, wearable technology, given the potential systemic risks they introduce.

6.3.5 Beyond Awareness and Consent - Trust, Fairness, Accountability and Morality. Even if our proposals for consent and awareness PETs were adopted and transparent, this does not mean users would accept it. For example, users would need to have trust - both in the system's effectiveness (e.g. see the dangers of algorithmic bias [18] and the numerous examples of failures here [83]), and that the system itself would/could not be abused (by the user, application, the platform, governments, etc.). Such a system would need to be seen as being fair, as who would purchase an AR headset where the majority of its sophisticated capabilities remained locked behind onerous privacy protections? And finally there is the moral aspect of these AR activities, and the potentially abusive intent behind them. Is it enough for a headset to disclose it is augmenting your appearance or volumetrically capturing you, if said augmentation renders the bystander nude, or their likeness is later used in a violent VR game? This suggests the extraordinary scope of abuse potentially unlocked by AR technology requires extraordinary measures to safeguard future bystanders, and necessitates responsible innovation [110] from AR platforms and their application developers [12].

7 CONCLUSION

The anticipated everyday usage of AR headsets will pose significant challenges to the privacy not just of users, but of bystanders in physical range of AR headset sensing. AR headsets will have the capacity to sense a breadth of bystander data - biometric characteristics, actions and behaviour, physiological data, identity/appearance - and then store or otherwise further process this data towards a variety of ends, such as augmented appearance, augmented intelligence, etc. Existing Privacy-Enhancing Technologies (PETs) often safeguard against these risks at a low level (e.g., instituting camera access controls). However, we argue such PETs are incompatible with the need for always-on sensing given AR headsets' intended everyday use. Through an online survey, we examined bystanders' awareness of, and concerns regarding, potentially privacy infringing AR activities. In particular, we found a strong need to support mechanisms for consent that can consider the social relationship between user

and bystander, and a varying need to support awareness for particularly invasive activities where consent was not previously granted. Reflecting on these findings, we discuss how we might provide architectures to support this varying need for bystander awareness and consent, where such mechanisms might fail or be circumvented, and where privacy alone may not be sufficient to describe the harms bystanders are exposed to. Our work reinforces the view dystopian visions of distributed surveillance and AR-enabled panopticons must be addressed if we are to avoid further social rejection of this powerful technology in the future.

ACKNOWLEDGMENTS

This research is supported by [REPHRAIN: The National Research Centre on Privacy, Harm Reduction and Adversarial Influence Online](#), under UKRI grant: EP/V011189/1.

REFERENCES

- [1] 2021. Guide to the UK General Data Protection Regulation (UK GDPR). <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/> Publisher: ICO.
- [2] 2021. RightsCon: As AR/VR becomes a reality, it needs a human rights framework. <https://www.eff.org/event/rightscon-arvr-becomes-reality-it-needs-human-rights-framework>
- [3] Article 19. 2021. Emotion Recognition Technology Report. <https://www.article19.org/emotion-recognition-technology-report/>
- [4] Article 19. 2021. When bodies become data: Biometric technologies and free expression. <https://www.article19.org/biometric-technologies-privacy-data-free-expression/>
- [5] Melvin Abraham, Pejman Saeghe, Mark McGill, and Mohamed Khamis. 2022. Implications of XR on Privacy, Security and Behaviour: Insights from Experts. In *Nordic Human-Computer Interaction Conference (Aarhus, Denmark) (NordiCHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 30, 12 pages. <https://doi.org/10.1145/3546155.3546691>
- [6] Alessandro Acquisti, Ralph Gross, and Fred Stutzman. 2011. Faces of facebook: Privacy in the age of augmented reality. *BlackHat USA 2* (2011), 1–20.
- [7] Alessandro Acquisti, Ralph Gross, and Frederic D. Stutzman. 2014. Face recognition and privacy in the age of augmented reality. *Journal of Privacy and Confidentiality* 6, 2 (2014), 1. <https://doi.org/10/gh3w6c>
- [8] Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (Dec. 1999), 40–46. <https://doi.org/10.1145/322796.322806>
- [9] Intiaz Ahmad, Rosta Farzan, Apu Kapadia, and Adam J. Lee. 2020. Tangible Privacy: Towards User-Centric Sensor Designs for Bystander Privacy. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (Oct. 2020), 116:1–116:28. <https://doi.org/10/gmk8w9>
- [10] Tousif Ahmed, Apu Kapadia, Venkatesh Potluri, and Manohar Swaminathan. 2018. Up to a Limit? Privacy Concerns of Bystanders and Their Willingness to Share Additional Information with Visually Impaired Users of Assistive Technologies. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (Sept. 2018), 89:1–89:27. <https://doi.org/10/gmmd9b>
- [11] Rawan Alharbi, Mariam Tolba, Lucia C. Petito, Josiah Hester, and Nabil Alshurafa. 2019. To Mask or Not to Mask? Balancing Privacy with Visual Confirmation Utility in Activity-Oriented Wearable Cameras. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (Sept. 2019), 72:1–72:29. <https://doi.org/10/gmk8wg>
- [12] 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. <https://doi.org/10/gk6g7n> Publisher: Elsevier.
- [13] A. Barth, A. Datta, J.C. Mitchell, and H. Nissenbaum. 2006. Privacy and contextual integrity: framework and applications. In *2006 IEEE Symposium on Security and Privacy (S P'06)*. 15 pp.–198. <https://doi.org/10.1109/SP.2006.32>
- [14] Mitchell Baxter, Anna Bleakley, Justin Edwards, Leigh Clark, Benjamin R. Cowan, and Julie R. Williamson. 2021. "You, Move There!": Investigating the Impact of Feedback on Voice Control in Virtual Environments. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3469595.3469609>
- [15] BBC News. 2019. Upskirting now a crime after woman's campaign. <https://www.bbc.co.uk/news/uk-47902522>
- [16] Sebastian Benthall, Seda Gürses, and Helen Nissenbaum. 2017. Contextual Integrity through the Lens of Computer Science. *Found. Trends Priv. Secur.* 2, 1 (dec 2017), 1–69. <https://doi.org/10.1561/33000000016>
- [17] Jolie Bonner, Joseph O'Hagan, Florian Mathis, Jamie Ferguson, and Mohamed Khamis. 2021. Using Personal Data to Support Authentication: User Attitudes and Suitability. In *20th International Conference on Mobile and Ubiquitous Multimedia (Leuven, Belgium) (MUM 2021)*. Association for Computing Machinery, New York, NY, USA, 35–42. <https://doi.org/10.1145/3490632.3490644>
- [18] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Conference on Fairness, Accountability and Transparency*. PMLR, 77–91. <http://proceedings.mlr.press/v81/buolamwini18a.html> ISSN:

- 2640-3498.
- [19] Kent Bye. 2019. XR Ethics Manifesto. <https://www.slideshare.net/kentbye/xr-ethics-manifesto-updated-nov-2-2019>
- [20] Supriyo Chakraborty, Chenguang Shen, Kasturi Rangan Raghavan, Yasser Shoukry, Matt Millar, and Mani Srivastava. 2014. ipShield: A Framework For Enforcing Context-Aware Privacy. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. USENIX Association, Seattle, WA, 143–156. <https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/chakraborty>
- [21] Raja Chatila and John C. Havens. 2019. The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. In *Robotics and Well-Being*, Maria Isabel Aldinhas Ferreira, João Silva Sequeira, Gurvinder Singh Virk, Mohammad Osman Tokhi, and Endre E. Kadar (Eds.). Vol. 95. Springer International Publishing, Cham, 11–16. https://doi.org/10.1007/978-3-030-12524-0_2 Series Title: Intelligent Systems, Control and Automation: Science and Engineering.
- [22] Bobby Chesney and Danielle Citron. 2019. Deep fakes: A looming challenge for privacy, democracy, and national security. *Calif. L. Rev.* 107 (2019), 1753. Publisher: HeinOnline.
- [23] SungIk Cho, Seung-wook Kim, JongMin Lee, JeongHyeon Ahn, and JungHyun Han. 2020. Effects of volumetric capture avatars on social presence in immersive virtual environments. In *2020 IEEE conference on virtual reality and 3D user interfaces (VR)*. 26–34. <https://doi.org/10/gh2qr9> ISSN: 2642-5254.
- [24] Shreya Chopra and Frank Maurer. 2020. Evaluating User Preferences for Augmented Reality Interactions with the Internet of Things. In *Proceedings of the International Conference on Advanced Visual Interfaces (Salerno, Italy) (AVI '20)*. Association for Computing Machinery, New York, NY, USA, Article 20, 9 pages. <https://doi.org/10.1145/3399715.3399716>
- [25] Soumyadeb Chowdhury, Md Sadek Ferdous, and Joemon M. Jose. 2016. Bystander Privacy in Lifelogging. <https://doi.org/10/gmk729>
- [26] Strauss A. L. Corbin J. M. 1998. *Basics of qualitative research: techniques and procedures for developing grounded theory*. SAGE Publications, Inc.
- [27] Kirsten Crager and Anindya Maiti. 2017. Information leakage through mobile motion sensors: User awareness and concerns. In *Proceedings of the European Workshop on Usable Security (EuroUSEC)*.
- [28] Poppy Crum. 2019. Hearables: Here come the: Technology tucked inside your ears will augment your daily life. *IEEE Spectrum* 56, 5 (2019), 38–43. <https://doi.org/10/gh2qvw> Publisher: IEEE.
- [29] Claudia Cuador. 2016. From Street Photography to Face Recognition: Distinguishing between the Right to Be Seen and the Right to Be Recognized. *Nova L. Rev.* 41 (2016), 237. Publisher: HeinOnline.
- [30] Jaybie A. De Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2019. Security and Privacy Approaches in Mixed Reality: A Literature Survey. *Comput. Surveys* 52, 6 (Oct. 2019), 110:1–110:37. <https://doi.org/10/ghbfgq>
- [31] Paula Delgado-Santos, Giuseppe Stragapede, Ruben Tolosana, Richard Guest, Farzin Deravi, and Ruben Vera-Rodriguez. 2022. A Survey of Privacy Vulnerabilities of Mobile Device Sensors. *ACM Comput. Surv.* 54, 11s, Article 224 (sep 2022), 30 pages. <https://doi.org/10.1145/3510579>
- [32] 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*. ACM, Toronto Ontario Canada, 2377–2386. <https://doi.org/10/gh2sn5>
- [33] 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 (CHI '14)*. Association for Computing Machinery, New York, NY, USA, 2377–2386. <https://doi.org/10/gh2sn5>
- [34] Ellysse Dick. 2020. *How to Address Privacy Questions Raised by the Expansion of Augmented Reality in Public Spaces*. Technical Report. Information Technology and Innovation Foundation. <https://itif.org/publications/2020/12/14/how-address-privacy-questions-raised-expansion-augmented-reality-public>
- [35] Ellysse Dick. 2021. *Balancing User Privacy and Innovation in Augmented and Virtual Reality*. Technical Report. Information Technology and Innovation Foundation. <https://itif.org/publications/2021/03/04/balancing-user-privacy-and-innovation-augmented-and-virtual-reality>
- [36] Mariella Dimiccoli, Juan Marín, and Edison Thomaz. 2018. Mitigating Bystander Privacy Concerns in Egocentric Activity Recognition with Deep Learning and Intentional Image Degradation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 4 (Jan. 2018), 132:1–132:18. <https://doi.org/10/gmmfbg>
- [37] Pierre Dragicevic. 2015. HCI Statistics without p-values. (06 2015), 36.
- [38] Brian L. Due. 2015. The social construction of a Glasshole: Google Glass and multiactivity in social interaction. *Psychology Journal* 13, 2 (2015).
- [39] Chloe Eghtebas, Francisco Kiss, Marion Koelle, and Paweł Woźniak. 2021. Advantage and Misuse of Vision Augmentation – Exploring User Perceptions and Attitudes using a Zoom Prototype. In *Augmented Humans Conference 2021 (AHs'21)*. Association for Computing Machinery, New York, NY, USA, 77–85. <https://doi.org/10/gmk86w>
- [40] Lisa A Elkin, Matthew Kay, James J Higgins, and Jacob O Wobbrock. 2021. An aligned rank transform procedure for multifactor contrast tests. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 754–768.

- [41] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10/gf5d6v>
- [42] Liv Erickson. 2020. Exploring Digital Rights: Data Sovereignty in XR. <https://www.youtube.com/watch?v=H3tMiSzRHA0>
- [43] Cori Faklaris, Francesco Cafaro, Asa Blevins, Matthew A. O'Haver, and Neha Singhal. 2020. A Snapshot of Bystander Attitudes about Mobile Live-Streaming Video in Public Settings. *Informatics* 7, 2 (June 2020), 10. <https://doi.org/10/gmmfdh> Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- [44] Md Sadek Ferdous, Soumyadeb Chowdhury, and Joemon M. Jose. 2017. Analysing privacy in visual lifelogging. *Pervasive and Mobile Computing* 40 (Sept. 2017), 430–449. <https://doi.org/10/gbx5p3>
- [45] Mary Anne Franks. 2017. The Desert of the Unreal: Inequality in Virtual and Augmented Reality. *U.C.D. L. Rev.* 51 (Jan. 2017), 499. https://repository.law.miami.edu/fac_articles/539
- [46] Sofien Gannouni, Arwa Aledaily, Kais Belwafi, and Hatim Aboalsamh. 2021. Emotion detection using electroencephalography signals and a zero-time windowing-based epoch estimation and relevant electrode identification. *Scientific Reports* 11, 1 (March 2021), 7071. <https://doi.org/10/gk6hbg> Bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Computational biology and bioinformatics; Neuroscience Subject_term_id: computational-biology-and-bioinformatics; neuroscience.
- [47] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & security* 77 (2018), 226–261.
- [48] Uwe Gruenefeld, Abdallah El Ali, Wilko Heuten, and Susanne Boll. 2017. Visualizing out-of-view objects in head-mounted augmented reality. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*. Number 81. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3098279.3122124>
- [49] 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. <https://doi.org/10.1145/3173574.3173628>
- [50] Jassim Happa, Mashhuda Glencross, and Anthony Steed. 2019. Cyber Security Threats and Challenges in Collaborative Mixed-Reality. *Frontiers in ICT* 6 (2019). <https://doi.org/10/gh2pgv> Publisher: Frontiers.
- [51] Jassim Happa, Anthony Steed, and Mashhuda Glencross. 2021. Privacy-certification standards for extended-reality devices and services. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, Lisbon, Portugal, 397–398. <https://doi.org/10/gmfc7p>
- [52] David Harborth and Sebastian Pape. 2021. Investigating privacy concerns related to mobile augmented reality Apps – A vignette based online experiment. *Computers in Human Behavior* 122 (Sept. 2021), 106833. <https://doi.org/10.1016/j.chb.2021.106833>
- [53] Tim Hardwick. 2021. Facebook Weighing Up Legality of Facial Recognition in Upcoming Smart Glasses. <https://www.macrumors.com/2021/02/27/facebook-facial-recognition-smart-glasses-legal/>
- [54] Rakibul Hasan, David Crandall, Mario Fritz, and Apu Kapadia. 2020. Automatically Detecting Bystanders in Photos to Reduce Privacy Risks. In *2020 IEEE Symposium on Security and Privacy (SP)*. 318–335. <https://doi.org/10/gmbpxc> ISSN: 2375-1207.
- [55] Mariam Hassib, Hatem Abdelmoteleb, and Mohamed Khamis. 2020. Are my Apps Peeking? Comparing Nudging Mechanisms to Raise Awareness of Access to Mobile Front-facing Camera. In *19th International Conference on Mobile and Ubiquitous Multimedia*. ACM, Essen Germany, 186–190. <https://doi.org/10/gmmgch>
- [56] Brittan Heller. 2020. Reimagining Reality: Human Rights and Immersive Technology. *Carr Center Discussion Paper Series* 2020-008 (2020).
- [57] Brittan Heller. 2020. Watching Androids Dream of Electric Sheep: Immersive Technology, Biometric Psychography, and the Law. *Vanderbilt Journal of Entertainment & Technology Law* 23, 1 (Dec. 2020), 1. <https://scholarship.law.vanderbilt.edu/jetlaw/vol23/iss1/1>
- [58] Matt Honan. 2013. I, Glasshole: My Year With Google Glass. *Wired* (Dec. 2013). <https://www.wired.com/2013/12/glasshole/>
- [59] Sabrina Hoppe, Tobias Loetscher, Stephanie A. Morey, and Andreas Bulling. 2018. Eye Movements During Everyday Behavior Predict Personality Traits. *Frontiers in Human Neuroscience* 12 (2018), 105. <https://doi.org/10.3389/fnhum.2018.00105>
- [60] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. Association for Computing Machinery, New York, NY, USA, 571–582. <https://doi.org/10/cqdz>
- [61] Olivier Hugues, Philippe Fuchs, and Olivier Nannipieri. 2011. New augmented reality taxonomy: Technologies and features of augmented environment. In *Handbook of augmented reality*. Springer, 47–63.
- [62] Soheil Human and Florian Cech. 2021. A Human-Centric Perspective on Digital Consenting: The Case of GAFAM. In *Human Centred Intelligent Systems (Smart Innovation, Systems and Technologies)*, Alfred Zimmermann, Robert J. Howlett, and Lakhmi C. Jain (Eds.). Springer, Singapore, 139–159. <https://doi.org/10/gh6fc4>
- [63] Marcello Lenca. 2021. Do We Have a Right to Mental Privacy and Cognitive Liberty? <https://blogs.scientificamerican.com/observations/do-we-have-a-right-to-mental-privacy-and-cognitive-liberty/>

- [64] 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 (New Orleans, LA, USA) (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 6, 8 pages. <https://doi.org/10.1145/3491101.3516396>
- [65] Matthew Kay, Lisa A. Elkin, James J. Higgins, and Jacob O. Wobbrock. 2021. ARTool: Aligned Rank Transform. <https://CRAN.R-project.org/package=ARTool>
- [66] Matthew Kay, Lisa A. Elkin, and Jacob O. Wobbrock. 2021. Contrast tests with ART. <https://cran.r-project.org/web/packages/ARTool/vignettes/art-contrasts.html>
- [67] Danielle Keats Citron. 2018. Sexual privacy. *Yale Lj* 128 (2018), 1870. Publisher: HeinOnline.
- [68] Mohamed Khamis and Florian Alt. 2021. Privacy and Security in Augmentation Technologies. In *Technology-Augmented Perception and Cognition*, Tilman Dingler and Evangelos Niforatos (Eds.). Springer International Publishing, Cham, 257–279. https://doi.org/10.1007/978-3-030-30457-7_8
- [69] 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 (NordiCHI '18)*. Association for Computing Machinery, New York, NY, USA, 473–481. <https://doi.org/10/gjbvvp>
- [70] Marion Koelle, Wilko Heuten, and Susanne Boll. 2017. Are you hiding it? usage habits of lifelogging camera wearers. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10/gmk87j>
- [71] Marion Koelle, Edgar Rose, and Susanne Boll. 2019. Ubiquitous Intelligent Cameras—Between Legal Nightmare and Social Empowerment. *IEEE MultiMedia* 26, 2 (April 2019), 76–86. <https://doi.org/10/gmk7rh> Conference Name: IEEE MultiMedia.
- [72] 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 (CHI EA '19)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10/gmk82v>
- [73] Marion Koelle, Katrin Wolf, and Susanne Boll. 2018. Beyond LED Status Lights - Design Requirements of Privacy Notices for Body-worn Cameras. In *Proceedings of the Twelfth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '18)*. Association for Computing Machinery, New York, NY, USA, 177–187. <https://doi.org/10/gmk7cz>
- [74] Jacob Leon Kröger, Leon Gellrich, Sebastian Pape, Saba Rebecca Brause, and Stefan Ullrich. 2022. Personal information inference from voice recordings: User awareness and privacy concerns. *Proceedings on Privacy Enhancing Technologies* 2022, 1 (2022), 6–27.
- [75] Katharina Kromholz, Adrian Dabrowski, Matthew Smith, and Edgar Weippl. 2017. Exploring Design Directions for Wearable Privacy. <https://publications.cispa.saarland/2808/>
- [76] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking. In *Privacy and Identity Management. Data for Better Living: AI and Privacy: 14th IFIP WG 9.2, 9.6/11.7, 11.6/SIG 9.2.2 International Summer School, Windisch, Switzerland, August 19–23, 2019, Revised Selected Papers*, Michael Friedewald, Melek Önen, Eva Lievens, Stephan Krenn, and Samuel Fricker (Eds.). Springer International Publishing, Cham, 226–241. https://doi.org/10.1007/978-3-030-42504-3_15
- [77] Mikko Kytö, Ilyena Hirschy-Douglas, and David McGookin. 2021. From Strangers to Friends: Augmenting Face-to-face Interactions with Faceted Digital Self-Presentations. In *Augmented Humans Conference 2021 (AHs'21)*. Association for Computing Machinery, New York, NY, USA, 192–203. <https://doi.org/10/gmk7c9>
- [78] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–31.
- [79] 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)*. 392–408. <https://doi.org/10/gmbp4q> ISSN: 2375-1207.
- [80] Kyungjun Lee, Daisuke Sato, Saki Asakawa, Hernisa Kacorri, and Chieko Asakawa. 2020. Pedestrian Detection with Wearable Cameras for the Blind: A Two-way Perspective. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376398>
- [81] Mark A. Lemley and Eugene Volokh. 2017. Law, virtual reality, and augmented reality. *U. Pa. L. Rev.* 166 (2017), 1051. https://scholarship.law.upenn.edu/penn_law_review/vol166/iss5/1/ Publisher: HeinOnline.
- [82] Jackson Lingane. 2021. FPF Report: Mitigate the Privacy Risks of AR & VR Tech. <https://fpf.org/blog/fpf-report-outlines-opportunities-to-mitigate-the-privacy-risks-of-ar-vr-technologies/>
- [83] Ryan Mac. 2021. Facebook Apologizes After A.I. Puts 'Primates' Label on Video of Black Men. *The New York Times* (Sept. 2021). <https://www.nytimes.com/2021/09/03/technology/facebook-ai-race-primates.html>
- [84] makobouzu. 2021. voice-canceling sound + semantic segmentation. <https://twitter.com/amako0609/status/1381454621990735873>
- [85] Vincent Manancourt and Mark Scott. 2020. Facebook earmarks €302M for privacy fines. <https://www.politico.eu/article/facebook-earmarks-e302m-for-privacy-fines/>

- [86] Steve Mann. 2013. Veilance and reciprocal transparency: Surveillance versus sousveillance, AR glass, lifelogging, and wearable computing. In *2013 IEEE International Symposium on Technology and Society (ISTAS): Social Implications of Wearable Computing and Augmented Reality in Everyday Life*. 1–12. <https://doi.org/10/gmk7hr> ISSN: 2158-3412.
- [87] Karola Marky, Nina Gerber, Michelle Gabriela Pelzer, Mohamed Khamis, and Max Mühlhäuser. 2022. “You offer privacy like you offer tea”: Investigating Mechanisms for Improving Guest Privacy in IoT-Equipped Households. *Proceedings on Privacy Enhancing Technologies* 4 (2022), 400–420. <https://doi.org/10.56553/popets-2022-0115>
- [88] Karola Marky, Sarah Prange, Florian Krell, Max Mühlhäuser, and Florian Alt. 2020. “You just can’t know about everything”: Privacy Perceptions of Smart Home Visitors. In *19th International Conference on Mobile and Ubiquitous Multimedia*. ACM, Essen Germany, 83–95. <https://doi.org/10/gmk7df>
- [89] Karola Marky, Alexandra Voit, Alina Stöver, Kai Kunze, Svenja Schröder, and Max Mühlhäuser. 2020. “I don’t know how to protect myself”: Understanding Privacy Perceptions Resulting from the Presence of Bystanders in Smart Environments. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society (NordCHI ’20)*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10/gmk7dg>
- [90] Arunesh Mathur, Jonathan Mayer, and Mihir Kshirsagar. 2021. What Makes a Dark Pattern... Dark? Design Attributes, Normative Considerations, and Measurement Methods. *arXiv:2101.04843 [cs]* (Jan. 2021). <https://doi.org/10.1145/3411764.3445610> arXiv: 2101.04843.
- [91] Daniel McDuff and Christophe Hurter. 2018. InPhysible: Camouflage Against Video-Based Physiological Measurement. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 5784–5789. <https://doi.org/10/gmmfcw> ISSN: 1558-4615.
- [92] Mark McGill. 2021. The IEEE Global Initiative on Ethics of Extended Reality (XR) Report—Extended Reality (XR) and the Erosion of Anonymity and Privacy. *Extended Reality (XR) and the Erosion of Anonymity and Privacy - White Paper* (Nov. 2021), 1–24. Conference Name: Extended Reality (XR) and the Erosion of Anonymity and Privacy - White Paper.
- [93] Mark McGill, Stephen Brewster, David McGookin, and Graham Wilson. 2020. Acoustic Transparency and the Changing Soundscape of Auditory Mixed Reality. (April 2020). <https://doi.org/10.1145/3313831.3376702>
- [94] 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. *26th ACM Symposium on Virtual Reality Software and Technology* (Nov. 2020), 1–10. <https://doi.org/10/ghwq2>
- [95] Maryam Mehrnezhad, Ehsan Toreini, Siamak F Shahandashti, and Feng Hao. 2018. Stealing PINs via mobile sensors: actual risk versus user perception. *International Journal of Information Security* 17, 3 (2018), 291–313.
- [96] Franziska Meirose, Sven Schultze, Sebastian Kuehlewind, Marion Koelle, Larbi Abdenebaoui, and Susanne Boll. 2018. Towards Respectful Smart Glasses through Conversation Detection. (2018). <https://doi.org/10/gmk87r> Accepted: 2018-07-09T08:40:22Z Publisher: Gesellschaft für Informatik e.V..
- [97] Monica Athnasious. 2021. Creepshots: what are they and why are they still happening? Victims share their experiences. <https://screenshot-media.com/politics/human-rights/what-is-a-creepshot/>
- [98] Alec G. Moore, Ryan P. McMahan, Hailiang Dong, and Nicholas Ruozzi. 2021. Personal Identifiability of User Tracking Data During VR Training. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 556–557. <https://doi.org/10/gk6gwx>
- [99] Emilio Mordini and Holly Ashton. 2012. The Transparent Body: Medical Information, Physical Privacy and Respect for Body Integrity. In *Second Generation Biometrics: The Ethical, Legal and Social Context*, Emilio Mordini and Dimitros Tzovaras (Eds.). Springer Netherlands, Dordrecht, 257–283. https://doi.org/10.1007/978-94-007-3892-8_12
- [100] N. A. Moreham. 2014. BEYOND INFORMATION: PHYSICAL PRIVACY IN ENGLISH LAW. *The Cambridge Law Journal* 73, 2 (July 2014), 350–377. <https://doi.org/10/gpgcqW> Publisher: Cambridge University Press.
- [101] Anne Nassauer and Nicolas M. Legewie. 2021. Video Data Analysis: A Methodological Frame for a Novel Research Trend. *Sociological Methods & Research* 50, 1 (Feb. 2021), 135–174. <https://doi.org/10/gfwsfz> Publisher: SAGE Publications Inc.
- [102] BBC News. 2015. Google apologises for Photos app’s racist blunder. <https://www.bbc.co.uk/news/technology-33347866>. Accessed: 2021-09-08.
- [103] BBC News. 2021. Facebook apology as AI labels black men primates. <https://www.bbc.co.uk/news/technology-58462511>. Accessed: 2021-09-08.
- [104] Helen Nissenbaum. 2009. Privacy in context. In *Privacy in Context*. Stanford University Press.
- [105] Joseph O’Hagan, Mohamed Khamis, Mark McGill, and Julie R. Williamson. 2022. Exploring Attitudes Towards Increasing User Awareness of Reality From Within Virtual Reality. In *ACM International Conference on Interactive Media Experiences (Aveiro, JB, Portugal) (IMX ’22)*. Association for Computing Machinery, New York, NY, USA, 151–160. <https://doi.org/10.1145/3505284.3529971>
- [106] Joseph O’Hagan, Mohamed Khamis, and Julie R. Williamson. 2021. Surveying Consumer Understanding & Sentiment Of VR. In *Proceedings of the International Workshop on Immersive Mixed and Virtual Environment Systems (MMVE ’21)* (Istanbul, Turkey) (MMVE ’21). Association for Computing Machinery, New York, NY, USA, 14–20. <https://doi.org/10.1145/3458307.3460965>

- [107] Joseph O'Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. 2022. Dataset and full survey for "Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent". <https://doi.org/10.5281/zenodo.7244156>
- [108] Joseph O'Hagan and Julie R. Williamson. 2020. Reality Aware VR Headsets. In *Proceedings of the 9TH ACM International Symposium on Pervasive Displays* (Manchester, United Kingdom) (*PerDis '20*). Association for Computing Machinery, New York, NY, USA, 9–17. <https://doi.org/10.1145/3393712.3395334>
- [109] 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* (Manchester, United Kingdom) (*PerDis '20*). Association for Computing Machinery, New York, NY, USA, 19–27. <https://doi.org/10.1145/3393712.3395339>
- [110] Richard Owen, Phil Macnaghten, and Jack Stilgoe. 2012. Responsible research and innovation: From science in society to science for society, with society. *Science and Public Policy* 39, 6 (12 2012), 751–760. <https://doi.org/10.1093/scipol/scs093> arXiv:<https://academic.oup.com/spp/article-pdf/39/6/751/4588094/scs093.pdf>
- [111] Joseph O'Hagan, Julie R Williamson, Mohamed Khamis, and Mark McGill. 2022. Exploring Manipulating In-VR Audio To Facilitate Verbal Interactions Between VR Users And Bystanders. In *International Conference on Advanced Visual Interfaces (AVI 2022)*.
- [112] Joseph O'Hagan, Julie R. Williamson, Mark McGill, and Mohamed Khamis. 2021. Safety, Power Imbalances, Ethics and Proxy Sex: Surveying In-The-Wild Interactions Between VR Users and Bystanders. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 211–220. <https://doi.org/10.1109/ISMAR52148.2021.00036>
- [113] Alfredo Perez, Sherali Zeadally, Luis Matos Garcia, Jaouad Mouloud, and Scott Griffith. 2018. FacePET: Enhancing Bystanders' Facial Privacy with Smart Wearables/Internet of Things. *Electronics* 7, 12 (Dec. 2018), 379. <https://doi.org/10/gmk837>
- [114] Alfredo J. Perez and Sherali Zeadally. 2018. Privacy Issues and Solutions for Consumer Wearables. *IT Professional* 20, 4 (July 2018), 46–56. <https://doi.org/10/f79t> Conference Name: IT Professional.
- [115] Alfredo J. Perez, Sherali Zeadally, Scott Griffith, Luis Y. Matos Garcia, and Jaouad A. Mouloud. 2020. A User Study of a Wearable System to Enhance Bystanders' Facial Privacy. *IoT* 1, 2 (Dec. 2020), 198–217. <https://doi.org/10/gmk83t> Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- [116] Mark Pesce. 2020. Augmented Reality and the Surveillance Society. <https://spectrum.ieee.org/computing/hardware/augmented-reality-and-the-surveillance-society>
- [117] Sarah Prange, Ahmed Shams, Robin Piening, Yomna Abdelrahman, and Florian Alt. 2021. PriView– Exploring Visualisations to Support Users' Privacy Awareness. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 69, 18 pages. <https://doi.org/10.1145/3411764.3445067>
- [118] Andrew Raij, Animikh Ghosh, Santosh Kumar, and Mani Srivastava. 2011. Privacy Risks Emerging from the Adoption of Innocuous Wearable Sensors in the Mobile Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (*CHI '11*). Association for Computing Machinery, New York, NY, USA, 11–20. <https://doi.org/10.1145/1978942.1978945>
- [119] Philipp A. Rauschnabel, Jun He, and Young K. Ro. 2018. Antecedents to the adoption of augmented reality smart glasses: A closer look at privacy risks. *Journal of Business Research* 92 (2018), 374–384. <https://doi.org/10.1016/j.jbusres.2018.08.008>
- [120] D.A. Reid, S. Samangoeei, C. Chen, M.S. Nixon, and A. Ross. 2013. Soft Biometrics for Surveillance: An Overview. In *Handbook of Statistics*. Vol. 31. Elsevier, 327–352. <https://doi.org/10.1016/B978-0-444-53859-8.00013-8>
- [121] Neil Richards and Woodrow Hartzog. 2016. Privacy's Trust Gap: A Review Book Review. *Yale Law Journal* 126, 4 (2016), 1180–1224. <https://heinonline.org/HOL/P?h=hein.journals/ylr126&i=1230>
- [122] Neil Richards and Woodrow Hartzog. 2017. Privacy's Trust Gap: A Review. *THE YALE LAW JOURNAL* (2017), 45.
- [123] Jan Ole Rixen, Teresa Hirzle, Mark Colley, Yannick Etzel, Enrico Rukzio, and Jan Gugenheimer. 2021. Exploring Augmented Visual Alterations in Interpersonal Communication. (2021).
- [124] Katitza Rodriguez and Kurt Opsahl. 2020. Augmented Reality Must Have Augmented Privacy. <https://www.eff.org/deeplinks/2020/10/augmented-reality-must-have-augmented-privacy>
- [125] Franziska Roesner, Tamara Denning, Bryce Clayton Newell, Tadayoshi Kohno, and Ryan Calo. 2014. Augmented reality: hard problems of law and policy. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp '14 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 1283–1288. <https://doi.org/10/gmmd9x>
- [126] Franziska Roesner and Tadayoshi Kohno. 2021. Security and Privacy for Augmented Reality: Our 10-Year Retrospective. (Aug. 2021), 5.
- [127] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (2014), 88–96. <https://doi.org/10/gh2pf6> Publisher: ACM New York, NY, USA.
- [128] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J. Wang. 2014. World-Driven Access Control for Continuous Sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Scottsdale Arizona USA, 1169–1181. <https://doi.org/10/ggp6n4>
- [129] Adam Rogers. 2018. So Long, Glassholes: Wearables Aren't Science Projects Anymore. *Wired* (May 2018). <https://www.wired.com/story/google-glass-predicted-the-future/>

- [130] Angela Sasse. 2015. Scaring and Bullying People into Security Won't Work. *IEEE Security Privacy* 13, 3 (2015), 80–83. <https://doi.org/10.1109/MSP.2015.65>
- [131] Roman Schlegel, Apu Kapadia, and Adam J. Lee. 2011. Eyeing your exposure: quantifying and controlling information sharing for improved privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS '11)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10/fzcp8r>
- [132] Stefan Schneegass, Romina Poguntke, and Tonja Machulla. 2019. Understanding the Impact of Information Representation on Willingness to Share Information. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (*CHI '19*). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3290605.3300753>
- [133] Hanna Schraffenberger and Edwin Van der Heide. 2014. Everything augmented: On the real in augmented reality. *Journal of Science and Technology of the Arts* 6, 1 (2014), 17–29. <https://doi.org/10/gh2qvb>
- [134] Hanna Kathrin Schraffenberger. 2018. *Arguably augmented reality: relationships between the virtual and the real*. Ph.D. Dissertation. ISBN: 9789492679673 OCLC: 1082194193.
- [135] Gwen Shaffer. 2021. Applying a Contextual Integrity Framework to Privacy Policies for Smart Technologies. *Journal of Information Policy* 11 (2021), 222–265. <https://www.jstor.org/stable/10.5325/jinfopoli.11.2021.0222>
- [136] Dangdang Shao, Chenbin Liu, and Francis Tsow. 2021. Noncontact Physiological Measurement Using a Camera: A Technical Review and Future Directions. *ACS Sensors* 6, 2 (Feb. 2021), 321–334. <https://doi.org/10/gmk7ht> Publisher: American Chemical Society.
- [137] Jiayu Shu, Rui Zheng, and Pan Hui. 2017. Your Privacy Is in Your Hand: Interactive Visual Privacy Control with Tags and Gestures. In *Communication Systems and Networks (Lecture Notes in Computer Science)*, Nishanth Sastry and Sandip Chakraborty (Eds.). Springer International Publishing, Cham, 24–43. <https://doi.org/10/gmmfbk>
- [138] Jiayu Shu, Rui Zheng, and Pan Hui. 2018. Cardea: context-aware visual privacy protection for photo taking and sharing. In *Proceedings of the 9th ACM Multimedia Systems Conference*. Association for Computing Machinery, New York, NY, USA, 304–315. <https://doi.org/10.1145/3204949.3204973>
- [139] Dan Simmons. 2021. 13 Countries with GDPR-like Data Privacy Laws. <https://insights.comforte.com/13-countries-with-gdpr-like-data-privacy-laws>
- [140] Samarth Singhal, Carman Neustaedter, Thecla Schiphorst, Anthony Tang, Abhisekh Patra, and Rui Pan. 2016. You are Being Watched: Bystanders' Perspective on the Use of Camera Devices in Public Spaces. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. Association for Computing Machinery, New York, NY, USA, 3197–3203. <https://doi.org/10/gmmfbj>
- [141] Mel Slater, Cristina Gonzalez-Lienres, Patrick Haggard, Charlotte Vinkers, Rebecca Gregory-Clarke, Steve Jelley, Zillah Watson, Graham Breen, Raz Schwarz, William Steptoe, Dalila Szostak, Shivashankar Halan, Deborah Fox, and Jeremy Silver. 2020. The Ethics of Realism in Virtual and Augmented Reality. *Frontiers in Virtual Reality* 1 (2020). <https://doi.org/10/ggpvct> Publisher: Frontiers.
- [142] Snapchat. 2018. Lens Studio. <https://lensstudio.snapchat.com/>. Accessed: 2021-09-08.
- [143] Softpedia News. 2015. In Japan, Phone Camera Shutter Sounds Can't Be Muted to Prevent Upskirt Photography. <https://www.softpedia.com/blog/in-japan-phone-camera-shutter-sounds-can-t-be-muted-to-prevent-upskirt-photography-487446.shtml>
- [144] Titus Stahl. 2016. Indiscriminate mass surveillance and the public sphere. *Ethics and Information Technology* 18, 1 (March 2016), 33–39. <https://doi.org/10/ggsqqr>
- [145] Jay Stanley. 2019. The Dawn of Robot Surveillance: AI, Video Analytics, and Privacy. <https://www.aclu.org/report/dawn-robot-surveillance>
- [146] Julian Steil, Inken Hagedstedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications (ETRA '19)*. Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10/gf6dj3>
- [147] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019. PrivacEye: privacy-preserving head-mounted eye tracking using egocentric scene image and eye movement features. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM, Denver Colorado, 1–10. <https://doi.org/10/gf8ged>
- [148] Megan Sullaway. 2022. Lone Wolves and Wolf Packs: Revenge Porn, Cyber Mobs, and Creepshots. *Indoctrination to Hate: Recruitment Techniques of Hate Groups and how to Stop Them* (2022), 217.
- [149] Omer Tene and Jules Polonetsky. 2013. A theory of creepy: technology, privacy and shifting social norms. *Yale JL & Tech.* 16 (2013), 59. Publisher: HeinOnline.
- [150] Marc Teyssier, Marion Koelle, Paul Strohmeier, Bruno Fruchard, and Jürgen Steimle. 2021. Eyecam: Revealing Relations between Humans and Sensing Devices through an Anthropomorphic Webcam. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10/gksmb5>
- [151] Alexandra Thompson and Leigh Ellen Potter. 2019. Overlays and Goggles and Projections, Oh My! Exploring Public Perceptions of Augmented Reality Technologies. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction (OZCHI'19)*. Association for Computing Machinery, New York, NY, USA, 295–301. <https://doi.org/10/gmk7c6>

- [152] Alexandra Thompson and Leigh Ellen Potter. 2020. Defining AR: Public Perceptions of an Evolving Landscape. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10/gmk7c7>
- [153] Marc Tran. 2015. Combatting gender privilege and recognizing a woman's right to privacy in public spaces: Arguments to criminalize catcalling and creepshots. *Hastings Women's Lj* 26 (2015), 185.
- [154] Yilun Wang and Michal Kosinski. 2018. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of personality and social psychology* 114, 2 (2018), 246.
- [155] Susan Wiedenbeck, Jim Waters, Leonardo Sobrado, and Jean-Camille Birget. 2006. Design and evaluation of a shoulder-surfing resistant graphical password scheme. In *Proceedings of the working conference on Advanced visual interfaces (AVI '06)*. Association for Computing Machinery, New York, NY, USA, 177–184. <https://doi.org/10/bcrv2p>
- [156] Julie R. Williamson, Joseph O'Hagan, John Alexis Guerra-Gomez, John H Williamson, Pablo Cesar, and David A. Shamma. 2022. Digital Proxemics: Designing Social and Collaborative Interaction in Virtual Environments. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 423, 12 pages. <https://doi.org/10.1145/3491102.3517594>
- [157] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 143–146. <https://doi.org/10.1145/1978942.1978963>
- [158] Katrin Wolf, Karola Marky, and Markus Funk. 2018. We should start thinking about Privacy Implications of Sonic Input in Everyday Augmented Reality! (2018). <https://doi.org/10/gmk7sk> Accepted: 2018-08-18T10:44:01Z Publisher: Gesellschaft für Informatik e.V..
- [159] Niels Wouters, Ryan Kelly, Eduardo Velloso, Katrin Wolf, Hasan Shahid Ferdous, Joshua Newn, Zaher Joukhadar, and Frank Vetere. 2019. Biometric Mirror: Exploring Ethical Opinions towards Facial Analysis and Automated Decision-Making. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. Association for Computing Machinery, New York, NY, USA, 447–461. <https://doi.org/10/gjbt7h>
- [160] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. 2018. Generative Image Inpainting with Contextual Attention. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5505–5514. <https://doi.org/10.1109/CVPR.2018.00577>
- [161] Rafael Yuste, Jared Genser, and Stephanie Herrmann. 2021. It's Time for Neuro-Rights. *Horizons: Journal of International Relations and Sustainable Development* 18 (2021), 154–165. Publisher: Center for International Relations and Sustainable Development.
- [162] Nan-ning Zheng, Zi-yi Liu, Peng-ju Ren, Yong-qiang Ma, Shi-tao Chen, Si-yu Yu, Jian-ru Xue, Ba-dong Chen, and Fei-yue Wang. 2017. Hybrid-augmented intelligence: collaboration and cognition. *Frontiers of Information Technology & Electronic Engineering* 18, 2 (Feb. 2017), 153–179. <https://doi.org/10/gg6r35>