

Why Do You Know That About Me?—Ethical Implications of Pervasive AR and Face Recognition

KUSHANI PERERA, University of Otago, Aotearoa New Zealand

NADIA PANTIDI, Victoria University of Wellington, Aotearoa New Zealand

HOLGER REGENBRECHT, University of Otago, Aotearoa New Zealand



Fig. 1. Pervasive AR-based face recognition in everyday social interactions introduces behavioural and ethical implications. Left: Conceptualisation of the experience. Middle: Real study experience. Right: Facial recognition-based information delivery as seen through the glasses. (Disclaimer: The images presented in this paper are staged for illustrative purposes and do not feature real participants. The individuals pictured in this and all figures have consented to the use of their images in this publication.)

Smart, pervasive Augmented Reality (AR) glasses are making their way out of the research labs. Many big tech companies are working on developing these promising next-generation interaction devices, and apps and services around them. When integrated with emerging face recognition technologies (FRT), Pervasive AR glasses can become powerful everyday tools. However, little is known about their acceptance, perceptual, and ethical ramifications. To address this, we developed a Pervasive AR technology probe with functional FRT and conducted an empirical study with 54 participants in a public environment. We collected interview data regarding perceived ethics about combining Pervasive AR with FRT. We developed five dominant themes informing the potential concerns and characteristics. Based on those findings, we propose to develop future Pervasive AR systems around principles of symmetry and consent—what we call a Kantian approach. We hope that our research will inform the design and development of near-future smart glasses.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**.

Additional Key Words and Phrases: Pervasive Augmented Reality, Ethics, Technology Probe, Empirical Study

Authors' Contact Information: Kushani Perera, University of Otago, Dunedin, Aotearoa New Zealand, kushani.perera@postgrad.otago.ac.nz; Nadia Pantidi, Victoria University of Wellington, Wellington, Aotearoa New Zealand, nadia.pantidi@vuw.ac.nz; Holger Regenbrecht, University of Otago, Dunedin, Aotearoa New Zealand, holger.regenbrecht@otago.ac.nz.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

Manuscript submitted to ACM

Manuscript submitted to ACM

ACM Reference Format:

Kushani Perera, Nadia Pantidi, and Holger Regenbrecht. 2026. Why Do You Know That About Me?—Ethical Implications of Pervasive AR and Face Recognition. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain*. ACM, New York, NY, USA, 27 pages. <https://doi.org/10.1145/3772318.3790645>

1 Introduction

Augmented Reality (AR) is rapidly transitioning from a single-application display technology [6, 84] to a context-aware, omnipresent technology [24, 67, 68]—Pervasive Augmented Reality. Pervasive AR delivers continuous AR visualisations embedded in our environment via inconspicuous devices that increasingly resemble regular spectacles [33]. Context-aware Pervasive AR systems deliver tailored content by drawing on prompts from always-on sensors, such as cameras and microphones, taking into account the user’s surroundings—including locations, objects, and bystanders¹ [86]—to construct the user’s context at any given moment.

We believe that Pervasive AR will leverage real-time face recognition to scrape information about a user’s bystanders. Real-time face recognition has shown promise in several applications, including assisting those with difficulty memorising faces [88], emotion analysis for neurodivergent individuals [9, 51], recognition support for visually impaired persons [48, 95], and social contexts such as recalling individuals the user has previously met [69].

As AR gains popularity, driven by substantial investments from major technology companies such as Meta, Apple, Google, and Snap Inc [28], there is growing interest in implementing face recognition and reverse-image search on these devices [53], which has been shown to be technically straightforward [34]. Several existing services² already allow reverse face searches to obtain any publicly available information on individuals.

However, few studies have explored the social, behavioural, and ethical implications of deploying such a powerful technology. Early initiatives like Google Glass and the Aria project serve as cautionary examples [5], escalating the need for designing future Pervasive AR systems responsibly by proactively addressing their ethical implications.

Looking ahead, we anticipate that access to face-recognition-based information will be tiered, similar to the levels offered by existing reverse-image search engines (see Figure 2). In such systems, a user’s subscription level would determine both their access to others’ FRT-based information³ and their ability to control how they are presented to other pervasive face recognition⁴ users. To explore the ethical implications of such a system, we conducted a 54-participant study, exposing participants to a functional pervasive face recognition probe in an ecologically valid environment—a public café—thus simulating a near-real Pervasive AR scenario. Moreover, the face recognition model was trained on participants’ photo data and deployed via subtle AR devices—Brilliant Frame (see Figure 3)—to mitigate previously reported participant concerns regarding the suitability of AR devices’ form factor for everyday use [61, 62]. To further ensure ecological validity, we scraped information about participants using PimEyes to be served upon recognition.

In this study, we attempt to identify **the ethical and social implications of receiving information about others in social interactions via pervasive face recognition**. To this end, we have defined the following two research questions, focusing the scope of our exploration:

¹In this study, a “bystander” refers to any individual present in the same environment as a Pervasive AR user. Bystanders could be other Pervasive AR users or non-users.

²e.g., www.pimeyes.com, www.socialcatfish.com, and www.facecheck.id

³Personal information retrieved by recognising an individual via FRT.

⁴We refer to Pervasive AR with face recognition capabilities as “pervasive face recognition”.

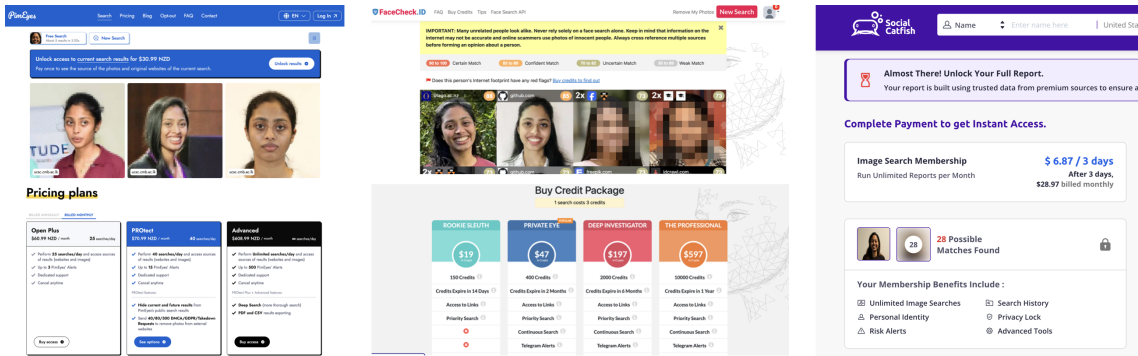


Fig. 2. This image depicts the access levels (subscription models) offered by three existing reverse image search engines. PimEyes (left), FaceCheck ID (middle), and Social Catfish (right). Each engine requires a paid subscription to access the original sources of the identified search results.

RQ1: What are the ethical and social implications of role-based exchange of pervasive face recognition-based information?

RQ2: What are the ethical and social implications of receiving scraped information via pervasive face recognition?

We intend to inform the responsible design of future Pervasive AR systems through our qualitative findings on users’ and non-users’ perceptions of pervasive face recognition.

2 Related Work

Pervasive AR is a ubiquitous technology that continuously augments a user’s environment with tailored information Grubert et al. [33]. Its context-aware capabilities are enabled by always-on sensors that constantly scan the user’s surroundings. We anticipate that by scanning a user’s environment, the system will recognise bystanders via FRT and deliver the obtained information to the user. In exploring this scenario, Perera et al. [63] reported several ethical and behavioural implications, creating numerous opportunities to further explore how ethical the use of pervasive face recognition systems would be—especially in a public space, as opposed to in a lab environment. We aim to explore the ethical and behavioural implications of pervasive face recognition systems in public settings, particularly when these systems deliver role-based scraped information to users about bystanders.

2.1 Applications of Real-Time FRT

FRT is most commonly deployed in security and law enforcement contexts. For example, they are widely used in airport security for border control and monitoring the arrival and departure of persons of interest [4], and at various stages of U.S. law enforcement procedures [30]. During COVID-19, FRT was also adopted for access control in public spaces [58]. In China, large-scale surveillance systems such as “Skynet”, equipped with over 560 million cameras, demonstrate the extensive public deployment of FRT [65].

Beyond security and surveillance, FRT has also been applied in assistive technologies, including support for visually impaired individuals [48, 95], those with difficulty recognising or remembering faces [88], and individuals who struggle with interpreting emotions [9, 51]. Additionally, FRT has been proposed as a social aid, such as Billingham and Starner [11]’s vision of a wearable that recognises a handshake, triggering FRT to display contextual details about the other person.

Worstell [92] highlights the appeal of FRT-enabled AR glasses that can discreetly remind users of names and past meetings. Rhodes [69] emphasises the value of context-aware wearables for task recall, noting the role of face recognition in supporting social interactions. Similarly, Utsumi et al. [87] present a mobile real-time FRT system to enhance human memory, while Singletary and Starner [74] also explore how FRT can improve memory and detect social interactions with minimal device interruptions.

In more recent developments of FRT delivered via AR devices, Meta announced that they intend to deliver face recognition capabilities through their upcoming AR device [53]. Hill [34] reports on two undergraduates who implemented a real-time face recognition apparatus via a Ray-Ban Meta device. This application, in particular, sparked a very important dialogue on FRT on wearables and their social implications. These undergraduates streamed the live video captured via their Ray-Ban Meta glasses to a face detection algorithm, which, when a face is detected, would search for that person using PimEyes. Then, using those results, they further scour through databases such as the voter registration database to identify the person’s name, address, and phone number. Finally, all these sources are compiled through a large language model (LLM) and are returned to the user⁵. In 2015, Wassom [89] noted that though major companies so far have shown restraint in employing FRT, it is too attractive a feature to be expected to be kept off of AR devices for long. Furthermore, they claim that facial privacy would be more effective when delivered by the market than by law enforcement. Similarly, Learned-Miller et al. [47] claims FRT-enabled surveillance could threaten a person’s right to anonymity in public, disrupting freedom of speech and discouraging association with certain individuals. Besides the privacy and data collection issues, the inaccuracies of the existing face recognition algorithms also pose a major threat to those being detected.

Albiero and Bowyer [2], through their experimental findings, state that female faces are inherently more similar to each other than male faces, causing face recognition systems to present a gender bias. Similarly, existing face recognition models have shown certain levels of racial bias [16]. When used in law enforcement applications, these biases then put the burden of proving oneself on potentially innocent people [50].

Furthermore, FRT adoption trends vary across countries and even within states, largely depending on existing privacy and data protection legislation [4, 27]. For example, the EU’s Artificial Intelligence regulation distinguishes between public use of FRT and applications such as device unlocking. Under Article 9 of the GDPR, biometric data may only be used for identification with an individual’s explicit consent, unless specific exemptions apply (e.g., law enforcement).

While privacy and data collection are surely of great concern, when used on unassuming bystanders, FRT poses ethical conundrums, too [34]. We assume pervasive and social FRT would alter our interactions as they deliver unprompted information about those we meet. Bar et al. [8] claim that first impressions are formed within the first minute of meeting someone. Accordingly, we anticipate that the information a user receives about another individual will inherently contribute to that individual’s self-presentation and influence the impression formation process. Thus, in the following section, we briefly review existing literature on self-presentation and first impressions.

2.2 Self-Presentation

In the case of pervasive face recognition, the unprompted information the devices provide about a person will count as the first impression of that person, thus reducing the control a person has to mould their presentation to others. Goffman [32] state that individuals “stage” themselves—self-presentation—to be perceived a certain way by the “audience”. Kim

⁵www.youtube.com/watch?v=S6pYBEYRRaE



Fig. 3. This image depicts the evolution of AR Head Mounted Devices' (HMD) form factor over the years. Microsoft HoloLens 1 (left), Snap Spectacles 2021 (middle), and Brilliant Frame (right).

and Baek [39] claim that selective self-presentation involves sharing mostly positive information about oneself to be perceived more favourably by others; Goffman [32] refers to this as *the idealised self-presentation*.

A considerable body of work investigating self-presentation via social media and avatars already exists [7, 21, 35]. DeVito et al. [21] discuss different attributes of social media platforms (e.g., types and ephemerality of content) that affect a user's self-presentation. Similarly, Hollenbaugh [35] note that users' self-presentations are based on: social media affordances, anonymity, persistence, and visibility. Litt [49] identifies different elements that determine the *imagined audience* an individual bases their online self-presentation on, stating that most users consider their imagined audience to be their most contacted peers.

Freeman and Maloney [29] examine self-presentation in social VR, showing the importance of expressing one's cultural identities and using voice modulators for authentic self-presentations. Similarly, Kytö and McGookin [46] explore how self-curated representations can be useful at all stages of a conversation. Moreover, Chung et al. [17] explore the interactions between AR users and non-users, claiming that non-wearers deemed the loss of control over their self-presentations when wearers would project augmentations on non-wearers, a critical issue.

In this work, we aim to forecast the ethical implications of pervasive face recognition based on the behavioural and social changes it brings about. Therefore, in the next section, we briefly discuss existing frameworks that enable such mappings.

2.3 Ethical Frameworks

Morals are self-defined guides that distinguish what is right and wrong, influenced by one's socio-cultural background, upbringing, and values instilled (as per the Aristotelian view of morals). Alternatively, ethics is a universal set of rules rooted in morals that society believes to be a framework for acceptable conduct [70]. In the last decade, more attention has been paid to ethics in emerging technologies [15, 55, 81]. Especially with a focus on ethics related to Artificial Intelligence (AI) [3, 42, 56, 75, 80].

Following a thorough comparison of existing frameworks, such as Wright [93]'s framework for the ethical impact assessment of information technology and Brey [15]'s anticipatory technology ethics, we adopt Stahl et al. [81]'s Ethical Issues of Emerging ICT Applications (ETICA) framework to address our research questions. Stahl and Eke [80] note that there are significant overlaps among these three frameworks. ETICA examines issues related to explicit morality and moral intuitions [79] through a pluralist approach that encompasses both ethical and social concerns. Most importantly, the ETICA framework was developed by analysing the ethical implications of 11 core technologies, one of

which is AR. Moreover, Stahl et al. [81] provide guiding questions that enable us to map issues identified in different technological contexts to specific categories of ethical implications. The categories are, 1. Conceptual issues and ethical theories—addressing issues relating to the conceptual clarity regarding the technology, 2. Impact on individuals—addressing issues that affect an individual’s rights and their well-being, 3. Consequences for society—encompassing concerns that affect a group or society as a whole, 4. Uncertainty of outcomes—investigating unforeseen uses and behaviours of technologies and their implications, 5. Perceptions of technology—encompassing consequences of technology displaying anthropomorphic behaviours, autonomy, and its ability to replace humans, and 6. Role of humans—addressing issues arising from the technology’s influence on how individuals view themselves, others, and their interactions.

2.4 Research Gap and Summary

Pervasive AR is an imminent ubiquitous technology that will continuously deliver tailored augmentations to its user, including information about their bystanders scraped from the internet. This will bring about numerous social and behavioural changes with ethical implications. Although there exists a vast body of work debating the social acceptability and the ethics of FRT, there is only limited work that explicitly explores FRT via Pervasive AR (pervasive face recognition). We are by no means attempting to justify the facial data collection of systems such as the Aria project by Facebook [25], but instead we intend to explore further the potential boundaries that may need to be imposed and the contexts in which facial recognition will deliver a value proposition to users as opposed to being a threat to their privacy.

Although Perera et al. [63] explored the ethical implications of pervasive face recognition, the study was conducted in a controlled lab environment and did not employ face recognition that delivered scraped information. However, we recognise that pervasive face recognition needs to be explored in ecologically valid conditions. Thus, we conducted a study in a public setting, a café with bystanders present, using a technology probe with functional FRT trained on participants’ photo data, as opposed to merely emulating its functions. Additionally, the information about participants was scraped from the internet using an existing reverse image search service, PimEyes, as opposed to participants submitting their own information (cf. Perera et al. [63]), which would have created a false sense of security. While Perera et al. [63] stated that participants conveyed concerns about non-users being recognised via pervasive face recognition, this was not explicitly explored. This is a crucial investigation as non-users are always present in real environments and must be considered when shaping the ethical and responsible design of Pervasive AR (cf. Stefanidi et al. [82]).

3 A Study on Pervasive Face Recognition

Pervasive face recognition has now become a near-future scenario [34, 53] and has already been empirically examined [63]. However, the study described by Perera et al. [63] was limited to a controlled lab environment. In contrast, we exposed participants to ecologically valid conditions—specifically, a real public café setting, utilising a technology probe [36] with functional FRT that delivers scraped information. The use of a functional technology probe has been previously used to explore similar concepts [23, 72]. Accordingly, this study was designed to examine the ethical and social implications of: 1. receiving information about others in social interactions via Pervasive face-recognition; 2. pervasive face recognition on different user roles; and 3. pervasive face recognition based on the type of information shared and received (curated information versus scraped information). To address this aim, participants were assigned to different roles (paid users, free users, and non-users).

Table 1. User roles and role-based control over self-presentation and access to information about recognised users. *Information scraped in advance from the internet. **Self-curated information by participant.

Control	Role			
	Non-User	Primary User		
		Free	Paid	
Self-Presentation	Disabled	Disabled	Enabled	
Access to information about:	Non-Users	None	Scraped*	
	Free Users	None	Scraped*	
	Paid Users	None	Curated** + Scraped*	

3.1 Study Design

This study was approved by the University of Otago Human Ethics Committee. All participants took part in the study across two sessions and were exposed to the following conditions:

1 *Kantian Condition*

This condition served as the baseline for the study. Participants only saw each other’s names through the glasses, thus creating a *symmetric* exchange of information. The design for the Kantian condition is inspired by Emmanuel Kant’s principles, the categorical imperative and universalisability. Categorical imperatives are commands that all must universally and unconditionally follow, and universalisability states that individuals should only act in ways which could be applied to all without contradiction, thus *symmetrical* [37].

2 *In-the-Wild Condition*

During this exposure condition, participants saw role-based information about individuals who were recognised through their glasses. Different roles had different levels of access to different types of information (see Table 1), creating an *asymmetric* exchange of information.

As part of the study design, we assumed that users’ access to pervasive face recognition-based information will be tiered much like paid access versus free access in most reverse image search systems available (see Figure 2). We further opted to tier the control of one’s own presentation via the face recognition system, as this was an aspect of control that participants felt strongly about, as noted in Perera et al. [63]. Therefore, we established the variable **Role** with sub-levels as follows:

1 **Primary User**: Represents pervasive face recognition users.

1.1 **Free User**: Represents pervasive face recognition users who have access to fewer privileges of the system, similar to a free version.

1.2 **Paid User**: Represents a class of primary users who have elevated access. Similar to advanced access obtained by subscribing to a paid service as opposed to the free version of the same service.

2 **Non-User**: Secondary Actors [17]. Participants without pervasive face recognition systems.

The primary user roles are distinguished based on the level of **Control** they possess. This control attribute of primary user roles is similar to the dependent variable *control* explored by Perera et al. [63]. The control variable has two levels:

1 **Control of self-presentation**

2 **Control of access to information**

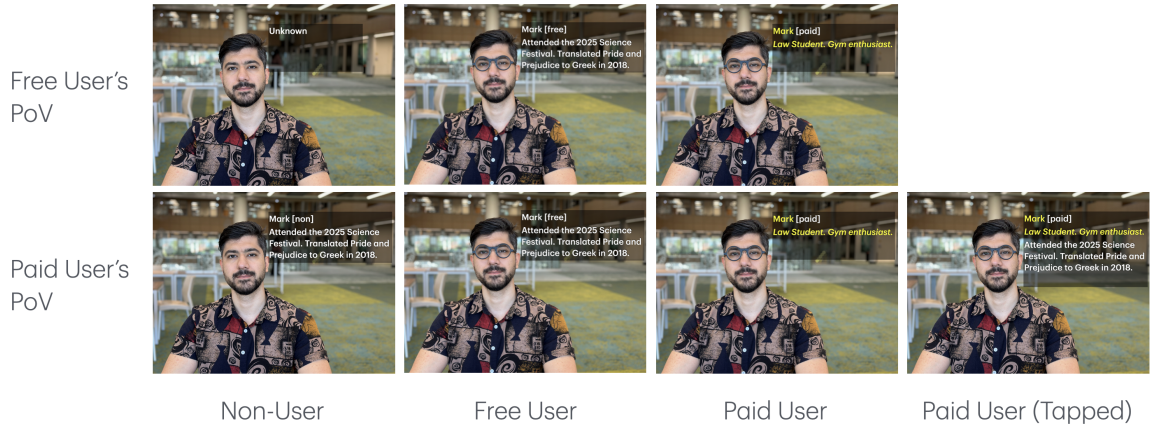


Fig. 4. This image depicts the level of information that different user roles had access to when interacting with other user roles. Curated information is shown in yellow font, and scraped information is shown in white font. (PoV: Point of View)

These two variables affect each type of primary user differently. Paid users have complete control of their own self-presentation and more access to others' information, compared to free users (see Table 1). A user-controlled self-presentation is referred to as *curated information*. Curated information about a primary user is therefore written up by the primary user themselves, similar to social media presentations, which are often idealised self-presentations [32]. In our study, only paid users are allowed to project such curated self-presentations of themselves. Meanwhile, free users' presentations are created with publicly available information about them, referred to as *scraped information*.

Seeing as how the different levels of the variable, *role*, are mutually exclusive to others with different roles being present in the environment, this variable is applied between-subject (see Figure 4). This simulation of different roles created an asymmetry of information. The findings from Perera et al. [63] suggested that participants favoured an idealised symmetric exchange of information. Thus, as a baseline, all participants were exposed to the *Kantian* condition in pairs (similar to the study design of Perera et al. [63]), prior to the *In-the-Wild* condition.

The *Kantian* condition served as an opportunity for participants to familiarise themselves with the AR glasses, especially for those assigned to the role of non-users. Ensuring this experience for non-users was important to provide all participants at least a vague idea of what others might be seeing via pervasive face recognition.

3.2 Procedure

Recruitment—Participants filled in a survey that allowed them to make a “profile” for themselves. This profile was a word description of themselves limited to a maximum of 80 characters (i.e., *curated information*). We asked the participants how they would describe themselves to someone they are meeting for the first time, with the following example, “*Computer Science student from Otago University. Interested in Mixed Reality.*”

Session 1—Participants first attended a brief 5-minute session. During this session, we explained the purpose of the study and captured two 1-minute videos of each participant's face from different angles and under different lighting conditions, both with and without the AR glasses. This data was subsequently used to construct the training dataset for the face recognition model.

Session 2—Was attended by six participants at a time; two participants each as paid users, free users, and non-users (see Figure 5). All participants were onboarded by completing a demographics questionnaire and being introduced to the concepts of Pervasive AR, FRT, and the aim of the study. They were then exposed to two conditions.

Kantian Condition (3 min): The participants introduced themselves to each other while wearing AR glasses that displayed each other's names. Following this session, they answered a questionnaire based on their experience.

In-the-Wild Condition (10 min): This condition was conducted in a café open to the public. Participants were told the following before the condition began:

- 1 There are two types of information displayed about others, a) information the detected person themselves has written up about them during recruitment (curated), and/or b) information scraped from the internet about a detected person (scraped).
- 2 Self-curated information will only be displayed about paid users (see Figure 4),
- 3 Scraped-information-based descriptions will be displayed about free users (see Figure 4),
- 4 The role of the person recognised is displayed next to their name on the AR glasses,
- 5 Only paid users see information about non-users, and that information is scraped (see Table 1),
- 6 Only paid users can see scraped information about other paid users in addition to the curated information, and the scraped information can be accessed by tapping on the glasses (see Figure 4), and
- 7 Free users see curated information about paid users, scraped information about other free users and no information about non-users.

The four participants in the roles of paid and free users were given AR glasses to wear for this condition, and all participants were advised to head down to the public café to find a suitable partner to attend a pub quiz with. Each participant was given a card with a pub quiz category (e.g., geography, pop-culture, sports, etc.). However, the task was only designed to encourage participants to interact with others in the environment.

Each session included 2–4 actors who posed as regular café patrons seeking to form pub-quiz teams. The actors were given a pub-quiz card and a \$10 café voucher and were instructed to, 1. initiate organic interactions, such as introducing themselves and discussing the pub quiz, 2. encourage participants to interact with other participants, 3. pretend not to know about the glasses (although they had tried them on and experienced face-recognition themselves prior to the sessions, as had the café employees), and 4. use the voucher to buy a primary user a coffee if the café is not too busy, thus triggering the recognition of café employees organically (see Figure 5). All café employees were informed of the study. Nine of the ten employees agreed to register with the face-recognition system either with their real details or fictional details which were generated using ChatGPT. All actors and registered café staff were treated as non-users, thus were recognised by paid users' glasses and displayed as "Unknown" on free users' glasses, much like any other non-user (see Figure 5 and Table 1). However, when participants interacted with persons who were not added to the system, such as the general public, the glasses displayed "Unknown".

The café employees and actors allowed us to simulate a realistic, pervasive face recognition-based future where bystanders are unpromptedly recognised by the system. Overall, there were five groups of people in the environment, 1. participants ($n = 6$); two participants per role, 2. café employees ($n = \sim 3$), 3. actors ($n = \sim 4$), 4. study facilitators ($n = 4$), and 5. the general public.

Following this condition, all participants were ushered back into a study room where they answered the final questionnaire. Following the completion of the questionnaire, all six participants sat with the researcher for a focus group interview. The interview lasted about 20–35 minutes. Finally, participants were thanked for participating in the



Fig. 5. Overview of the two-session study procedure. Session 1 involved recording participants with and without AR glasses for training data collection. Session 2 included the Kantian condition in a controlled setting and the In-the-Wild condition in a public café, followed by a concluding interview.

study and presented with a \$20 gift voucher as a token of appreciation. Once the study was completed, participants were debriefed via email, 1. about the actors and café employees being pre-registered in the system as opposed to being scanned and recognised in real-time, and 2. how the scraped information was manually collected by researchers instead of being retrieved from the internet in real-time.

3.3 Materials, Apparatus and Implementation

Interviews and Questionnaires—We were interested in the general themes of ethical implications from using pervasive face recognition in public social settings, from the perspectives of the different participant roles and their effect on the participants’ general behaviour in such situations. The guiding interview questions were adapted from the ETICA framework [81]. The questions provided us a scaffold to guide the interview rather than a rigid list of questions.

All participants filled out two types of questionnaires. Upon attending session 2, participants completed a demographics questionnaire about themselves, and following each condition, participants completed a questionnaire with 12 items regarding their experience in the respective condition. The questionnaire was answered on a 1–6 Likert-like scale. Eight questions were derived from the WEAR Scale⁶[38], and the remaining questions were developed by the researchers using Stahl et al. [81]’s guiding questions. The questionnaire served as a probe to direct participants towards our dimensions of interest during the interview. The questionnaire inquired about the system’s suitability for social interactions, its potential for privacy and exploitation violations, and was intended to prompt participants to reflect on these dimensions during the interview.

Training the Face Recognition Model—From in-lab testing, we determined that each person class requires approximately 150 images for accurate predictions to be made. Using FFmpeg⁷ we extracted three frames per second from the videos of participants shot during session 1. These frames were split into a training dataset and a validation dataset (80-20), and were used to train the locally hosted face recognition model. These images were then preprocessed using Multi-task Cascaded Convolutional Network (MTCNN)⁸ for face-detection-based image cropping to only include human faces, thus removing any other noise the face-recognition model might learn on [19, 94].

The preprocessed dataset was used to train a Dlib-based face recognition model with Python bindings⁹ [31, 40, 41]. We first used the Dlib face recognition model to extract each person’s face embeddings from images. Embeddings are 128-dimensional vectors that define a face’s unique features. Once all embeddings are saved as a pickle file¹⁰, they are

⁶Used to evaluate social acceptability of wearable technologies.

⁷FFmpeg is a library used for handling multimedia files. We use it to extract frames from video files [26].

⁸A deep-learning based framework used for localising facial landmarks (such as eyes, nose, etc.), face detection and providing a bounding box around the face.

⁹github.com/ageitgey/face_recognition

¹⁰A Python-specific format for storing Python objects as a serialised byte stream. Pickle files are commonly used for storing machine learning models [64].

used to compare detected face embeddings with the previously extracted embeddings in real time. The model then outputs the class name that best matches the detected face based on these comparisons.

Once the model recognised and output a name, it was used to query a database containing the user-role, curated information, and scraped information for each registered individual. Using a predefined database allowed us to obtain participants' consent in advance, rather than scraping the web in real-time based on images captured through the glasses during the experience, which would have subjected non-consenting members of the public to scanning. Prior to each session, a reverse image search of participants was conducted using PimEyes, and a description of each participant based on publicly available information as queried by the PimEyes search result (i.e., *scraped information*) was generated and stored in the local database.

Information scraping with PimEyes was performed only for study participants by submitting a photo taken during session 1 to the platform. The research team held a paid PimEyes subscription, whose privacy policy¹¹ states that uploaded images are used solely to provide lookup services and are never shared or sold to third parties. The policy further specifies that all images and search data are deleted within 30 days of account deletion. Accordingly, the research team deleted the PimEyes account promptly upon completing the final study session. This study underwent a rigorous ethics committee review, and participants were informed prior to session 1 about the intended use of their photographs, with the option to withdraw. The use of real scraped data was essential to ensure ecological validity and to inform and raise awareness among potential future Pervasive AR users about the functionalities of the technology.

Technology Probe—The probe used for the study was deployed on Brilliant Frame AR glasses. These glasses do not perform any processing onboard, but communicate with a host device via Low-Energy Bluetooth (BLE) to receive information to display. Due to the BLE radius limiting how far users can move, we opted to have a smartphone as the host device instead of the server itself. We implemented a mobile application¹² that handles the Bluetooth connections with the Frame AR glasses. The mobile application streamed images from the glasses to a server hosted on a private wireless network, which returned the detected names and corresponding information.

Carrying the mobile phone around allowed users to walk around the café using the FRT-enabled AR glasses. The Frame-based probe was also appealing due to the design of the glasses, as they resembled regular spectacles more than the other AR glasses in the market, such as Snap Spectacles¹³, which are more similar to sunglasses. Participant roles were predefined on the server side, and the corresponding information was delivered in JSON format based on these role mappings. The API request from the mobile application defined the requester's role as a query parameter. This parameter was used to determine the type of information to be served (e.g., `/process?user_id=free`). The received information was then converted to plain text and displayed on the Frames. The display only changed if the response was different from the previously displayed response. Furthermore, if a known person was detected, we delayed further scanning by three seconds to avoid slight movements returning a face-undetected response.

3.4 Data Collection and Analysis

The main data for this study were collected from interviews. The interviews were recorded, transcribed and pseudonymised using otter.ai. The primary author verified the transcripts for clarity against the original recordings. We followed the inductive thematic analysis process prescribed by Thomas [85] and referred to Braun and Clarke [14], and Clarke and Braun [18] for further clarifications when needed.

¹¹pimeyes.com/en/privacy-policy

¹²We built this application based on the codebase provided here: github.com/CitizenOneX/frame_vision_api_impl

¹³www.spectacles.com/

We conducted the thematic analysis with the following steps: 1. The primary author read and re-read the transcripts in their entirety for familiarisation. 2. The primary author then generated the initial set of codes. 3. A second author independently generated a list of codes by studying a set of transcripts. 4. The two sets of resulting codes from parallel coding were merged to create a single list of 70 codes. 5. The resulting categories were checked for clarity by the third author. 6. The codes were iteratively categorised to develop themes and discussed among authors. Furthermore, we have chosen not to quantify our qualitative findings when reporting to avoid unjustifiable generalisations by the readers, as the interviews were semi-structured group discussions [54].

4 Results

This study was conducted with 54 participants (32F, 21M, 1D) following ethics approval. 20 participants were assigned to the role of paid users, and 17 participants each to free users and non-users. The participants had a mean age of 25.4 and consisted of 45 university students and nine professionals. Participants were from 37 distinct disciplines, with the two main groups being Computer Science ($n = 8$) and Health Science ($n = 5$). 35 participants responded that they had no prior experience with AR, 18 responded they had *some* AR experience, and 1 responded they had *much* AR experience.

Although this study was primarily qualitative, participants completed a questionnaire following each condition. The questionnaire was designed by the researchers to serve as a probing tool for the subsequent interviews and, as such, was not validated. Therefore, we give limited significance to the quantitative results. However, we briefly report our findings so as not to disregard our participants' efforts. Overall, we were unable to find a significant difference between scores for the different user roles ($p = 0.665$) with the mixed ANOVA test. We were able to find a significant difference in scores for the Kantian condition and the *In-the-Wild* condition with a small effect size ($p = 0.0001$, $d = 0.38$).

We analysed the interview data using an iterative inductive thematic analysis process, through which we developed five themes that capture users' perceptions of pervasive face recognition in real-world contexts. These themes provide us with valuable insights into answering our research questions regarding the ethical and social implications of pervasive face recognition, especially in terms of roles and the use of scraped information versus curated information. These themes reflect users' perceptions of curating self-presentation, the asymmetries introduced by information and subscription models, their sense of control in using the system, the contrasts with traditional lookup methods, and anticipated future social norms.

4.1 To Curate or Not to Curate

This theme reports participants' reflections on the importance of managing their self-presentations. Self-presentation is a dyadic relationship consisting of two elements. 1. Those who are presenting themselves (actors) in a certain way to elicit a specific impression from others (audience) and 2. those who construct a certain impression of an individual based on the presentation they put forth [32]. In this theme, we discuss both the elements of self-presentation and their influence on the formation of first impressions. We discuss dimensions such as how the exchange of pervasive face recognition-based information affects self-presentation and the resulting debate about the curation of one's presentation via Pervasive AR. The first dimension we report under this theme captures the problems participants identified regarding the scraped information exchange.

Scraped Information Concerns—Participants expressed that the information displayed about them on others' AR glasses affected their self-presentation, noting that the information available on the internet about a person is not necessarily what they would choose to share in a first meeting, and attributed this concern to the scraped nature of the information, which can project inaccurate or outdated presentations. Furthermore, they added that the scraped

information could reflect a version of themselves they no longer identify with, and that the continuous visibility of past mistakes through pervasive face recognition would make it difficult for individuals to reform and move on:

*Leila (Non)*¹⁴: “Some of the glasses draw information from the internet, and that may not necessarily represent what you want to share with other people.”

Most importantly, participants saw scraped information as a barrier to accurate self-presentation, as it influenced the first impressions they formed of others. They noted that such information often preceded actual interactions, which would have been crucial for impression formation otherwise. Besides the information itself, some participants—particularly those in a user role—indicated that the use of pervasive face recognition systems itself could alter the way they are presented to others in social interactions. Participants noted that consuming pervasive face recognition-based information would be considered rude to the extent that users of the system would be perceived as intrusive individuals or stalkers. Thus making users more self-conscious about using the system. Participants further suggested that such perceptions could be sufficiently negative to result in social exclusion or ostracism:

Roman (Free): “If I see someone wearing those glasses, I’m going to be behaving slightly differently, probably etching away.”

Curation of Self-Presentation and Reasons—This dimension reports participants’ reflections on actively curating their self-presentation through Pervasive AR systems, rather than leaving it to be determined by scraped information. Participants stated that, given the option to curate their own presentation, they would rather curate it themselves. They further claimed that the action of curating one’s own self-presentation via pervasive face recognition systems mirrors the ways in which people presently adapt their self-presentation to suit different social contexts.

Participants also identified ways in which curating their self-presentation through pervasive face recognition could add value to their social lives. They identified that it could become a medium through which they can portray their most authentic selves, bypassing current social constructs that limit their ability to do so (cf. Sim et al. [73]). Much like indicating pronouns on social media or in email signatures, this system enables individuals to project their true selves in real-time social interactions, aligning with Goffman [32]’s concept of the idealised self-presentation:

Alex (Paid): “[For] trans people like myself, they may be able to put their pronouns on there, and that would be really helpful.”

In contrast, rather than using the system to enhance the projection of their self-presentation, some participants suggested that those with paid privileges might instead use it to remain anonymous by only sharing very little or no information at all. This suggested use of curation would allow them to both maintain the present norms of interacting with somebody, with the explicit choice of whom one introduces themselves to, and it would also allow users to remain anonymous in social contexts that would otherwise project their information to other nearby pervasive face recognition users.

However, curating one’s own presentation and receiving curated information about others were viewed differently (cf. Sim et al. [73]). While curating one’s own presentation was deemed necessary and appealing for the reasons stated above, participants believed that users with such privileges would be too dishonest in how they portray themselves. In this regard, participants were at a crossroads. Adding further nuance to how curated presentations could be misused, participants suggested that it might enable forms of catfishing in real-world social contexts. Specifically, paid members

¹⁴Participants have been pseudonymised. The role of the participant is indicated within brackets as “Non” for non-users, “Free” for free-users, and “Paid” for paid-users.

of pervasive face recognition systems exploiting their privileges to deliberately mislead and manipulate other users, thereby creating an unsafe social environment:

Lucy (Paid): “[Curation] is encouraging dishonesty. So if you have the power to control exactly what it says up there, you could easily manipulate somebody [...] and that could be really dangerous.”

In this theme, we reported our findings pertaining to participants’ views on curating their self-presentation and implications of using scraped information to construct self-presentations, such as inaccurate first-impression formations. In the next theme, we report on participants’ reflections on the asymmetry that pervasive face recognition creates in terms of information, roles, tier-based privileges, and its influence on their behaviours.

4.2 Subscription-Induced Asymmetries

There were certain facets of asymmetry that emerged purely due to the study design, the subscription model of pervasive face recognition we exposed the participants to, the privileges different tiers were entitled to, such as the ability to curate a user’s self-presentation, and the glasses’ capabilities of performing FRT altogether. Participants commented that the asymmetry created by pervasive face recognition was concerning and that it affected their social interactions negatively.

Asymmetry from Role-Based Access to Information—Paid users received scraped information about non-users in the environment, including participants in non-user roles, actors, and café staff. In contrast, neither free users nor non-users had access to the same information, creating an asymmetry based on subscription roles. Participants found this imbalance uncomfortable, particularly when most individuals in the environment were not paid, such as users. They considered the ability of paid users to recognise non-users without reciprocal information especially problematic, as it could create a social power imbalance. Consequently, participants suggested that the system should only recognise other Pervasive AR users (cf. Perera et al. [63]):

Ross (Paid): “So I know who you are. You have no idea who I am, it just gets really creepy really quickly.”

This asymmetry was not only limited to paid users and non-users, but it also created friction between free users and paid users. Although free users had access to the glasses, they felt they were at a disadvantage, as they could only access information about other users. This disparity was evident to us when a free-user-participant approached us during the *In-the-Wild* session. A free user, having seen a paid user recognise several non-users, came to us to inform us that his glasses might be frozen because “*most people are coming up Unknown*”. In this case, we had to reiterate their role and send them back into the environment.

Asymmetry From the Ability to Curate Self-Presentation—In addition to the role-based access to information, participants were further concerned about self-presentation curation. While paid users projected a curated self-presentation, free users were subjected to a self-presentation that was scraped from the internet based on information from publicly available sources. Free users could not tell what would be displayed about them to other users, and believed that this gave a sense of undue privilege to paid users, putting free users in a vulnerable position. Paid users themselves recognised the unfair nature of this setup:

Alex (Paid): “Because there’s definitely a power imbalance there, especially when paid users can choose what other people see about them, there’s the manipulation of information, and who gets to see what, and that creates inequality.”

Asymmetry and Its Influence on Social Behaviours—As seen above, participants conveyed that such inequity could create divisions between those with access to information and those without, further straining social connections and deepening existing disparities. Most importantly, they believed this would reinforce income-based hierarchies, especially when those with greater financial means can exploit those with fewer resources simply by affording a paid subscription.

Participants highlighted several behavioural changes among the different types of users that were a result of the asymmetry. Several non-users expressed that they felt disadvantaged and unable to engage due to the inequities the system created. Especially, owing to the lack of information and the imbalance that was made increasingly apparent in their interactions with paid users. This non-user describes the sense of impediment as follows:

Rachel (Non): “I didn’t have the glasses, so I found it difficult to make conversations. People knew everything about me, and I couldn’t reciprocate.”

Another result of feeling disadvantaged was that it could lead to a sense of missing out and serve as the sole incentive to adopt the system. Individuals would opt to buy the technology and even the subscription just to ensure that they were not disadvantaged by their lack of access to the information. This sense of peer pressure is also exacerbated by existing socioeconomic disparities and the desire to conform. The final behavioural change we identified came from free users. Free users hinted at how they gravitated more towards other glass wearers solely because they would receive information about them. Thus, they are discouraged from engaging with non-users. This aspect further informs non-users’ notion of feeling disadvantaged or feeling left out. This could also worsen existing class segregations. This free user stated their behaviour during the *In-the-Wild* session as follows:

Idris (Free): “As a free user, I tended to talk to the people with glasses more, because it could actually scan them[...].”

Need for Symmetry and Compensating Behaviours—The final dimension we report under this theme emphasises participants’ call for symmetry. As a result of the outcomes reported above, participants highlighted the importance of symmetry to preserve existing social norms. As this participant states, asymmetry creates an unfair social dynamic:

Kenny (Non): “Ideally, if we were going to have this, then everyone gets access to the paid membership, there shouldn’t be a tier system because it’s an unfair aspect.”

Participants in user roles during the *In-the-Wild* condition, conveyed that they adopted certain compensating behaviours to ensure non-users felt more comfortable. While some participants made up for the asymmetry by volunteering more information about themselves to match what they were seeing about non-users they engaged with, others opted to only reveal parts of the information they were receiving about non-users.

Asymmetry is a commonly explored theme related to Pervasive AR. It is also referred to as the problem of the haves-and-the-have-nots [24, 62, 67]. This is one of the concerns we again identified in our context of pervasive face recognition in social interactions, as reported above. However, we focused our reporting specifically on our context of interest rather than the asymmetry arising from the general use of Pervasive AR. Thus, discussing the asymmetry arising as a consequence of the role-based access, the privileges different users were entitled to and the resulting social behavioural changes.

The next theme encompasses our findings on participants’ need for control in using pervasive face recognition systems and their perception of control that the subscription model provides.

4.3 Navigating Control

This theme discusses the multiple factors participants often highlighted concerning their sense of control during the experience. We outline several elements of control that participants considered important should pervasive face recognition systems become commonplace. The dimensions reported under this theme refer to the need for control in terms of one's own information and the sharing of it (consent), control over who can be identified via the glasses, and control over the type of information a user receives about another.

Managing Consent, Privacy, and Information Disclosure—

Although the subscription model was designed to depict several aspects of control, such as the ability to curate one's self-presentation and access to varying levels of information about others, participants noted that it did not necessarily ensure control and underscored the importance of consent. Participants expressed that the lack of explicit consent to be identified was problematic because the study was designed so that the system allowed a user to walk into a café and recognise everyone in their vicinity (cf. Hill [34]). Even participants in the role of users recognised the importance of obtaining consent before searching for a detected person on the internet and recognising them.

By depriving bystanders of their right to consent, the use of pervasive face recognition invades their privacy. Participants believed that scraping information violated others' right to privacy. Participants discussed the privacy implications of pervasive face recognition, especially the invasiveness of delivering publicly available information about others to pervasive face recognition users without those individuals' permission.

Furthermore, participants believed the lack of control over who they share their personal information with could create unsafe and dangerous situations for those who are detected, leaving individuals in vulnerable circumstances to be exploited by those who are receiving information about them. Thus, participants highlighted the potential for the information and the technology to be used for malicious and manipulative purposes. However, some participants also believed that the issue was more due to the novelty of the technology as opposed to it being truly an invasion of privacy. They further stated that once the technology is widely used and accepted by the majority, the public will be less concerned about these implications:

Leila (Non): "I think the more it becomes normalised, the privacy issue kind of starts to go away, because everyone's doing it."

Managing Who Is Recognised—Based on the dilemma of consent, participants proposed several measures to regulate who is recognised by the system. These measures would enable the recognition of specific individuals through pervasive face recognition, rather than recognition based solely on proximity. In particular, participants suggested that only people an individual already knows should be able to access information about that individual, rather than anyone nearby recognising them simply by looking. Thus, ensuring each individual has control over whom they share their information with. Similarly, participants identified assisting with remembering names of people they had already met to be the only acceptable use-case for pervasive face recognition in daily life (cf. Rhodes [69]), thereby allowing users to retain some control over who is recognised by the system. This sentiment is reflected in the following quote:

Jason (Paid): "I'm someone who's very bad at remembering people's names, so maybe that will be kind of helpful."

Furthermore, participants believed that if the system was not implemented in a way that protects individuals' right to share, then the general public would opt to wear face covers in public as a control measure to regain control and remain anonymous.

Managing Incoming Information—The final dimension of control that participants highlighted as important was the kind of information they received about others. Participants expressed that when one looks someone up intentionally, they pick the sources of information they refer to, and that is often information curated by the individual one is looking up. However, when the glasses perform face recognition, they often gather information from indirect sources. Thus, the user no longer has control over picking which sources they refer to. Additionally, participants believed that the inability to verify information themselves by consulting only trusted sources, as they would otherwise do, could introduce new problems to their social interactions. They worried that the technology might misidentify individuals and that relying on such system-generated results could negatively shape their interactions. This final section on users’ need for control serves as a useful bridge to the next part of the findings. In the next theme, we explicitly report our findings on participants’ perceptions of pervasive face recognition in comparison to performing reverse face search using existing wearable devices such as mobile phones.

4.4 From Searching to Seeing: Contrasting Traditional Lookup and Pervasive Face Recognition

We used PimEyes to generate the information that users received about others (scraped information). Therefore, we were particularly interested in finding out how participants perceived receiving this information via pervasive face recognition in comparison to looking someone up using the same tools. They noted both similarities and differences between the two functions.

“The Same as Googling Someone”—Some participants felt that receiving information about others via pervasive face recognition was similar to looking someone up, as the same type of information is obtained. Consequently, they viewed scraping information about others as reasonable, arguing that presenting publicly available information is not problematic since it already exists. The following quote clearly portrays this view:

Lucy (Paid): “It’s information that’s there anyway. So if I wanted to find something out about someone, then I could go and search it up.”

From Active Search to Passive Exposure—In contrast, the majority of the participants highlighted how receiving pervasive face recognition-based information was, in fact, different from looking someone up based on several factors. Firstly, they conveyed that it is not the information itself that is concerning, but rather the unpromptedness and lack of intent. They further highlighted how the action of taking someone’s picture to look them up, as opposed to receiving the same information via pervasive face recognition, would forego the user’s choice to perform the search:

George (Paid): “The difference is intent, isn’t it? If you get your phone out, you’re purposely making a choice there to look someone up, and hopefully with their knowledge of it as well, because it’s quite an obvious thing to do. But here, you don’t really have any choice.”

The example quote above highlights another important difference participants noticed, which is the discreteness with which the action of looking someone up is executed via pervasive face recognition. Participants believed the action of taking someone’s picture would inherently allow those individuals to object when they notice being photographed. However, with Pervasive AR, the cameras are always-on and are more inconspicuous, as the camera would naturally be at the position a camera would need to be to photograph an individual’s face.

The final difference participants highlighted was the speed of information delivery. They noted that manually looking someone up takes time—you must open an app, upload a picture, and follow the steps—whereas pervasive

face recognition completes these steps automatically and instantly, providing information within seconds and further emphasising the unprompted nature of the delivery.

In this theme, we compiled our participants' views on how pervasive face recognition differs from the traditional act of looking someone up using the same tools. In the following final theme, we report our findings on behavioural changes participants foresee for a future of widespread use of pervasive face recognition.

4.5 Future Norms

Participants believed that pervasive face recognition would alter our behaviours both positively and negatively, with the regular use of these systems.

Shifts in How Interactions are Initiated—They conveyed that users (especially paid users) immediately showed altered social behaviours during the *In-the-Wild* condition. Where they would approach non-users, directly reading the information they receive about them from the glasses. Thus, exhibiting distorted behaviours in initiating engagements that deviate from currently acceptable norms. This participant, who was in a non-user role, described it as:

Ava (Non): "As a non user, the users were lacking in social interaction, they'd come up to me and they're like, 'Oh, you just won this award' because they were obviously seeing [a] Google blurb, and then they didn't offer any information about [themselves], so they know my name, what I do, and then I have to go, cool. What's your name?"

Further substantiating this claim, some observed that the glass-wearers appeared distracted or less engaged, focusing more on the information delivered through the glasses than on the interaction itself. Participants also believed that constant access to information about everyone could make people more cautious or private both in public and online, potentially leading to social anxiety and isolation. Furthermore, participants suggested that information scraping could prompt new behaviours, such as individuals becoming more conscious of what they post online, solely to ensure that their self-presentation via pervasive face recognition is acceptable. Both positives and negatives of this scenario were debated as seen below:

Smith (Non): "I think it could definitely result in people becoming more private and definitely more careful and more secluded."

Lili (Free): "Maybe it wouldn't be bad, because people would be more conscious about what they are putting on social network[s]."

Filtered Interactions—Participants expressed concerns about using the information that pervasive face recognition users see about others to modulate interactions, thus screening people. Participants further conveyed that the information displayed about someone would reinforce existing prejudices in screening people out. Which was argued as both a bad thing and a good thing. Either done for personal safety or purely due to prejudice, thus hurting social interactions and connections.

Furthermore, participants believed this action of filtering people out in real life would lead to limited interactions and ultimately lead to a more isolated future. When participants only interact with others who have similar views or seem alike based on the information they read, it will lead to information silos, which are both isolating and lacking important exposure to varying points of view. However, participants also noted that filtering interactions based on the information received via pervasive face recognition will be deemed useful in specific contexts, such as supporting individuals who are socially awkward or neurodivergent. By accessing information about others, these individuals could engage more confidently in interactions and filter social encounters in ways that would allow them to remain within their comfort

zones. Participants further commented on the tool being useful in environments that are designed for specific levels of engagement, such as conferences, where everyone present would use the device for efficient networking among one another. For example:

Ophelia (Free): “The only context that I could see them being useful and not problematic ethically would be in specific scenarios where you wear them for a certain amount of time and everyone’s agreed to it, and you write up what information you want people to see, [such as at] a conference.”

5 Discussion and Recommendations

To recap our motivation for this research, and its importance and timeliness: Just this year, several companies have announced new AR devices with AI capabilities, such as the Brilliant Halo¹⁵, Snap Spectacles¹⁶, Meta Hypernova¹⁷, and Google Gemini Glasses¹⁸. These new AR glasses are becoming increasingly more inconspicuous, showing growing promise in being integrated with users’ everyday lives. Almost three decades ago, Rhodes [69] conceptualised the ‘Remembrance Agent’, an AR application that remembers the names of individuals a user has met. Now, Brilliant has brought this concept to life with AR glasses equipped with an always-on multilingual AI agent that will do exactly that¹⁹. Another recent instance of FRT being paired with AR came from a group of undergraduates who created a pipeline between Ray-Ban Meta glasses, PimEyes and voter registration databases to then generate a profile of each person detected via a large language model (LLM) [34]. When an invasive tool like this can be implemented in mere days by anyone, we have to thoroughly investigate the social implications it will have on people.

In our study, we attempted to simulate a similar environment to gather both users’ and non-users’ positions on this technology. We believe that it is through anticipatory measures that we can best ensure the responsible design of Pervasive AR systems. We identified five themes that contribute towards the knowledge base of such social and ethical implications of delivering pervasive face recognition-based information. All five themes indicated that the use of Pervasive AR with FRT capabilities will alter our behaviours, thus leading to social consequences due to asymmetry, control, self-presentation, and the adoption of new behaviours. Accordingly, we utilise Stahl et al. [81]’s ETICA framework to systematically map the identified concerns to ethical and social implications that inform our research questions.

RQ1: What are the ethical and social implications of role-based exchange of pervasive face recognition-based information?—Although we implemented role-based access to mitigate users’ perceived lack of control in pervasive face recognition [63] and to reflect existing reverse face-search engines, our findings suggest that it instead intensified ethical and social concerns related to Stahl et al. [81]’s sub-category *Digital Divides* under *Consequences to Society*. Specifically, it created asymmetries in information access, fostered exclusion, and ultimately reinforced the digital divide—rather than alleviating it.

It is inevitable that a tailored technology such as Pervasive AR would result in inequalities and informational asymmetries [24, 63, 67]. However, as clearly reported in our theme *Subscription-Induced Asymmetries*, role-based access to information created asymmetries in both the information participants received and the information they disclosed, creating a sense of power imbalance, indicating a novel social and ethical implication that we believe has not

¹⁵brilliant.xyz/products/halo

¹⁶newsroom.snap.com/launch-specs-2026

¹⁷www.roadtovr.com/meta-connect-2025-schedule-smart-glasses-metaverse/

¹⁸blog.google/products/android/android-xr-gemini-glasses-headsets/

¹⁹www.linkedin.com/posts/bobak-tavangar-29445012_i-dont-remember-his-name-what-did-activity-7360135233867800577-Tci0

been explored in this context before. Moreover, by positioning non-users at a disadvantage without the technology, such access exacerbates existing socioeconomic inequalities by favouring those who can afford paid subscriptions. Nevertheless, as our participants indicated, this does not imply that paid users would necessarily be comfortable in this pervasive face recognition future, as non-users expressed intentions to distance themselves from pervasive face recognition users, thereby fostering stigma towards them—this hints at the re-emergence of the Google Glass “gl***holes” [22].

RQ2: What are the ethical and social implications of receiving scraped information via pervasive face recognition?—Collectively, the implications discussed below align with the social and ethical challenges described by Stahl et al. [81] under the category—*Impact on Individuals*, especially the sub-categories of *Privacy*, *Autonomy*, *Identity*, and *The Treatment of Humans*. Our theme, *To Curate or Not to Curate*, captured several concerns participants raised about scraped information, especially concerning the threat scraped information poses to self-presentation, as it shifts control from the individual to an internet-based dossier about themselves. Thus, scraped information-based self-presentations risk producing portraits of individuals that are potentially inaccurate, outdated, and entirely beyond one’s control (*Autonomy*). Building on this, *Navigating Control* highlights issues related to consent (*Treatment of Humans*), privacy (*Privacy*), and agency (*Autonomy*)—not only in how one’s own information is disclosed, but also in being forced to receive scraped information about others. Ultimately, such scraped information-based representations not only undermine one’s ability to present oneself as they choose to, but also risk resurfacing past indiscretions (*Identity*). As Solove [76] argue, although the data already exists, the creation of a dossier that creates a portrait of a person is a privacy violation. Moreover, they state that the disclosure of truthful and accurate information about an individual that will affect how they are judged by others, and the increased access to such public information about a person, also constitute privacy violations [77, 78].

Besides the ethical implications discussed above that directly answer our two research questions, Stahl et al. [81]’s guiding questions revealed several more participant concerns that reflect broader ethical and social consequences. Under *Uncertainty of Outcomes*—that is, issues arising from unforeseen misuse or unintended consequences—the following concerns participants raised pose significant ethical and social challenges: 1. potential manipulations based on received information or dishonest curations of self-presentation which would lead to social anxiety and altered online behaviours, and 2. distrust arising from misinformation and misidentification of individuals via pervasive face recognition. The final category of ethical concerns we report on is *Perception of Technology*—these ethical and social consequences relate to autonomous behaviours of pervasive face recognition systems, such as their unprompted delivery of information about bystanders and their inconspicuous scanning and recognition of bystanders without requiring explicit user input.

The concern of inconspicuous scanning by Pervasive AR systems is widely discussed, particularly regarding the privacy violations it can trigger (cf. Koelle et al. [44, 45], Wolf et al. [91]). In this work, however, we focus specifically on the privacy issues unique to pervasive face recognition and the delivery of scraped information, rather than the well-established concerns associated with always-on cameras. These include the social implications of continuous data collection and sousveillance [52, 62, 68, 71, 89]—as well as the impact on bystanders, including violations of their privacy (cf. Koelle et al. [43, 44], Zhao et al. [96]). Prior work has also shown that existing privacy indicators are largely ineffective in mitigating such concerns (cf. Bhardwaj et al. [10], Koelle et al. [45]). Additional risks of always-on cameras include, privacy and security vulnerabilities arising from data attacks [13, 60]. Finally, researchers have proposed several design recommendations to address these issues, including automatically closing shutters in sensitive contexts [83], actively notifying nearby bystanders [57], notifying them and seeking their permission [20], as well as displaying the user’s real-time view to people in the vicinity [45].

5.1 A Kantian Approach

Having identified these numerous ethical and social challenges, we anticipate that asymmetry and the lack of informed consent will be the most crucial. We would boldly argue that a Kantian implementation of pervasive face recognition systems will be considerably more acceptable than a role-based implementation (cf. Perera et al. [63]). In a broader sense, the quality and quantity of information a pervasive face recognition user is willing to disclose would determine the quality and quantity of information that user can receive from others. This would also implement a mechanism of implicit consent. Such an implementation, grounded in Kantian principles of universalisability—that a moral judgment or rule must be applicable to everyone and in every relevantly similar situation—would ensure that all individuals have access to the same level of information rather than being led by socioeconomic standings. It will ensure a transition from a world of public anonymity to one of symmetrical real-life social networking. Such an approach would also allow individuals to retain a sense of autonomy and agency over their self-presentation, rather than being thrust into an intrusive reality. Particularly, ensuring symmetry for users will, 1. ensure that users are in control of how much they share, thus also ensuring control over their own self-presentation, and 2. in addition, empower the presentation of users’ true selves, such as their pronouns (cf. Bonner et al. [12], O’Hagan et al. [59]). In the same vein, 3. users will have control of first impressions, 4. users’ agency to maintain their right to anonymity will be enhanced, 5. it will support voluntary exclusion for non-users without placing them at a disadvantage, and 6. bypass issues of unfavourable social dynamics arising from role-based privileges such as disparities.

As briefly mentioned above, we also advocate for seeking explicit consent from individuals, as difficult as it might be to implement this. Bystanders lacking informed consent to share their information with pervasive face recognition users purely on the grounds of the information being publicly available pose a critical ethical issue (see Stahl et al. [81]’s *Impact on Individuals - Treatment of Humans*). One should always reserve the right to control how they are portrayed and how their data is used²⁰. In subscribing to seeking informed consent, not only will the design address issues around consent, but also address 1. privacy concerns relating to the processing and dissemination of personal data, 2. concerns arising from users’ self-perceptions of feeling like intruders as the information is now willingly shared, 3. concerns surrounding the use of scraped information, such as its validity, accuracy, and suitability in specific contexts, 4. individuals’ sense of control in social interactions [90], and to an extent, 5. potential concerns arising from misinformation.

Although we concur with Acquisti et al. [1] that opt-ins are an ineffective tool, given that the data required to re-identify an individual already exists, we argue that technology designers should exercise restraint when sharing such information that conflicts with bystanders’ preferences and should instead seek explicit consent from individuals. Further informing this dialogue, O’Hagan et al. [59] found that, among the eleven AR activities they presented, participants rated personal identification as one of the top three intrusions, with participants overwhelmingly preferring everyday AR designs adopting an opt-out-by-default mechanism, with the possibility of seeking consent, rather than opt-in-by-default mechanisms. Windl et al. [90] adopts a user-centric approach to consent in Pervasive AR, allowing designers to follow a rigorous framework to determine the optimal consent mechanism for various environments and social scenarios. Alternatively, Rajaram et al. [66] offer several negotiable options that achieve a compromise between the required sensing capabilities and the privacy requirements of all stakeholders, rather than a one-size-fits-all approach, providing a starting point for operationalising consent-based implementation of Pervasive AR systems.

²⁰www.edps.europa.eu/data-protection/data-protection#Data-Protection

5.2 Limitations and Future Work

Our design propositions create space for future work exploring practical mechanisms for operationalising symmetry and informed consent in pervasive face recognition. We must stress that Pervasive AR-based research needs to be conducted empirically, to both capture experience-based reflections, but also to proactively raise awareness of what it means to live in a world with Pervasive AR. The early inclusion of potential future users in the dialogue is crucial to ensuring they feel heard and respected and that this powerful technology is accepted. It is also important to recognise that bystanders are a significant group of stakeholders in Pervasive AR research. Thus, they should be heard not only in terms of privacy violations but also in other, more specific contexts that explore social, behavioural, and ethical implications.

Nevertheless, there are several factors that could have impacted our results. Due to cancellations and no-shows, each group of roles were made up of uneven numbers (non-users: 17, free: 17, paid: 20). Given that this is a mutually exclusive study experience, we struggled to balance these numbers. However, we were still able to collect substantial qualitative data that provided valuable insights into the social and ethical implications of pervasive face recognition systems in the wild. Despite our best efforts, we were unable to recruit more professionals as opposed to university students, owing to our location, which should be a valid consideration for future studies.

Although participants seemed to believe that the scraping of information was being done in real-time (as opposed to being pre-scraped via PimEyes and stored in a local database), their involvement in session 1 to construct the training database may have signalled that the process was simulated. At the same time, the use of actors proved to be rather effective at creating a sense of real-time scraping as they were instructed to pose as if they were unaware of the AR glasses scanning them or having attended a prior session. The technology probe initially had a delay of 3 s in recognising individuals, which we managed to reduce to 1.8 s. Nevertheless, a few participants noticed this delay. This latency resulted from the time required to query the server. In principle, if scraping were performed truly in real time—for instance, by integrating an existing reverse-face-search engine—this delay could have been minimised. However, such an approach was not suitable for a public setting where regular café-goers were present.

Moreover, in this study, we situated our participants in a casual café setting, with the pop-quiz task serving as an incentive to interact. However, in more specific and charged contexts such as political rallies, a similar study would yield different results. For example, as noted in the related work, Learned-Miller et al. [47] anticipate that FRT-enabled surveillance could discourage people from protesting due to a loss of anonymity, thereby affecting freedom of speech. Similarly, in different cultural settings, such as indigenous contexts, pervasive face recognition may be perceived differently than discussed in this study. Finally, we identify that although existing ethical frameworks provide sufficient scaffolding, we need more focused frameworks that exclusively address future ubiquitous technologies, such as Pervasive AR, with a keen focus on social acceptability that is not limited to a device's form factor. These limitations create interesting avenues for exploration in future work.

6 Conclusion

The impending everyday use of context-aware Pervasive AR raises pressing social and ethical challenges that affect both users and non-users. As evident in recent AR device launches, face recognition technology (FRT) is no longer speculative but a likely feature of AR systems. In this study, we empirically examined the social and ethical implications of pervasive face recognition for individuals with differently controlled levels of access and privilege to the technology. We identified five themes that capture stakeholders' perceptions of pervasive face recognition and its impact on social

life, including effects on social asymmetries and disparities, self-presentation, first impressions, sense of agency, and everyday behaviours. Drawing on these findings, we argue that future Pervasive AR systems incorporating FRT must be designed with Kantian principles of symmetry and consent at their core, in order to prevent ethical rejection of the technology as a whole.

Author Contributions

Kushani Perera: writing- original draft, conceptualisation, methodology, software, formal analysis, investigation, data curation and visualisation. **Nadia Pantidi:** writing- review & editing, conceptualisation, methodology, formal analysis, and supervision. **Holger Regenbrecht:** writing- review & editing, conceptualisation, methodology, formal analysis, investigation, funding acquisition, and supervision.

Acknowledgments

We sincerely thank Mr Grant Bowie for funding this study through a PhD scholarship. We are also grateful to Pulina Udapamunuwa for his unwavering support. We thank our study facilitators—Krishna, Alban, and Connor—as well as our actors for their invaluable help in successfully conducting the study. We further extend our gratitude to the staff at the Te Mātiti Café and our colleagues in the University of Otago HCI Group for their time, support, and cooperation.

References

- [1] Alessandro Acquisti, Ralph Gross, and Fred Stutzman. 2014. Face Recognition and Privacy in the Age of Augmented Reality. *Journal of Privacy and Confidentiality* 6 (12 2014), 1–20. Issue 2. <https://doi.org/10.29012/jpc.v6i2.638>
- [2] Vitor Albiero and Kevin W. Bowyer. 2020. Is Face Recognition Sexist? No, Gendered Hairstyles and Biology Are. In *Proceedings of the 31st British Machine Vision Conference (BMVC)*. British Machine Vision Association (BMVA), Durham, UK, 1–13. <https://arxiv.org/abs/2008.06989>
- [3] Colin Allen, Iva Smit, and Wendell Wallach. 2005. Artificial morality: Top-down, bottom-up, and hybrid approaches. *Ethics and Information Technology* 7 (9 2005), 149–155. Issue 3. <https://doi.org/10.1007/S10676-006-0004-4/METRICS>
- [4] Sofia Andrade. 2023. Clear Wants to Scan Your Face at Airports. Privacy Experts Are Worried. *The Washington Post*, Dec. 20, 2023. [Online]. Available: www.washingtonpost.com/travel/2023/12/20/clear-facial-recognition-technology-airport-security/.
- [5] 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/10.1016/j.jrt.2021.100010>
- [6] Ronald T Azuma. 1997. A Survey of Augmented Reality. *Presence: Teleoperators and Virtual Environments* 6 (1997), 355–385. Issue 4. <https://doi.org/10.1162/pres.1997.6.4.355>
- [7] Chantal Bacev-Giles and Reeshma Haji. 2017. Online first impressions: Person perception in social media profiles. *Computers in Human Behavior* 75 (10 2017), 50–57. <https://doi.org/10.1016/J.CHB.2017.04.056>
- [8] Moshe Bar, Maital Neta, and Heather Linz. 2006. Very first impressions. *Emotion* 6, 2 (2006), 269.
- [9] Esmā Mansouri Benssassi, Juan Carlos Gomez, Louanne E. Boyd, Gillian R. Hayes, and Juan Ye. 2018. Wearable Assistive Technologies for Autism: Opportunities and Challenges. *IEEE Pervasive Computing* 17 (4 2018), 11–21. Issue 2. <https://doi.org/10.1109/MPRV.2018.022511239>
- [10] Divyanshu Bhardwaj, Alexander Ponticello, Shreya Tomar, Adrian Dabrowski, and Katharina Krombholz. 2024. In Focus, Out of Privacy: The Wearer’s Perspective on the Privacy Dilemma of Camera Glasses. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI ’24). Association for Computing Machinery, New York, NY, USA, Article 577, 18 pages. <https://doi.org/10.1145/3613904.3642242>
- [11] M. Billinghurst and T. Starner. 1999. Wearable devices: new ways to manage information. *Computer* 32, 1 (1999), 57–64. <https://doi.org/10.1109/2.738305>
- [12] Jolie Bonner, Florian Mathis, Joseph O’Hagan, and Mark McGill. 2023. When Filters Escape the Smartphone: Exploring Acceptance and Concerns Regarding Augmented Expression of Social Identity for Everyday AR. In *Proceedings of the 29th ACM Symposium on Virtual Reality Software and Technology* (Christchurch, New Zealand) (VRST ’23). Association for Computing Machinery, New York, NY, USA, Article 14, 14 pages. <https://doi.org/10.1145/3611659.3615707>
- [13] Martina Brachmann, Gregoire Phillips, Utku Gülen, and Valentin Tudor. 2023. Toward Privacy-Preserving Localization and Mapping in eXtended Reality: A Privacy Threat Model. In *2023 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. IEEE, Gothenburg, Sweden, 635–640. <https://doi.org/10.1109/EuCNC/6GSummit58263.2023.10188227>
- [14] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3 (2006), 77–101. Issue 2. <https://doi.org/10.1191/1478088706QP0630A>

- [15] Philip A.E. Brey. 2012. Anticipating ethical issues in emerging IT. *Ethics and Information Technology* 14 (12 2012), 305–317. Issue 4. <https://doi.org/10.1007/S10676-012-9293-Y/TABLES/1>
- [16] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. , 77–91 pages. <https://proceedings.mlr.press/v81/buolamwini18a.html>
- [17] Ji Won Chung, Xiyu Jenny Fu, Zachary Deocadiz-Smith, Malte F Jung, and Jeff Huang. 2023. Negotiating Dyadic Interactions through the Lens of Augmented Reality Glasses. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) (DIS '23). Association for Computing Machinery, New York, NY, USA, 493–508. <https://doi.org/10.1145/3563657.3595967>
- [18] Victoria Clarke and Virginia Braun. 2017. Thematic analysis. *The Journal of Positive Psychology* 12 (5 2017), 297–298. Issue 3. <https://doi.org/10.1080/17439760.2016.1262613>
- [19] Iván de Vicente. 2018. MTCNN face detector for TensorFlow. <https://github.com/ipazc/mtcnn>. Accessed: 2025-04-23.
- [20] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 2377–2386. <https://doi.org/10.1145/2556288.2557352>
- [21] Michael Ann DeVito, Jeremy Birnholtz, and Jeffery T. Hancock. 2017. Platforms, People, and Perception: Using Affordances to Understand Self-Presentation on Social Media. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17). Association for Computing Machinery, New York, NY, USA, 740–754. <https://doi.org/10.1145/2998181.2998192>
- [22] Brian Lystgaard Due. 2016. The social construction of a Glasshole: Google Glass and multiactivity in social interaction. *Psychology Journal* 13, 2-3 (2016), 149–178.
- [23] Lucy E Dunne, Halley Profita, Clint Zeagler, James Clawson, Scott Gilliland, Ellen Yi-Luen Do, and Jim Budd. 2014. The social comfort of wearable technology and gestural interaction, In 2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. 2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society NA, NA, 4159–4162. <https://doi.org/10.1109/EMBC.2014.6944540>
- [24] Chloe Eghtebas, Gudrun Klinker, Susanne Boll, and Marion Koelle. 2023. Co-Speculating on Dark Scenarios and Unintended Consequences of a Ubiquitous(ly) Augmented Reality. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (, Pittsburgh, PA, USA,) (DIS '23). Association for Computing Machinery, New York, NY, USA, 2392–2407. <https://doi.org/10.1145/3563657.3596073>
- [25] Jakob Engel, Kiran Somasundaram, Michael Goesele, Albert Sun, Alexander Gamino, Andrew Turner, Arjang Talattof, Arnie Yuan, Bilal Souti, Brigid Meredith, Cheng Peng, Chris Sweeney, Cole Wilson, Dan Barnes, Daniel DeTone, David Caruso, Derek Valleroy, Dinesh Ginjupalli, Duncan Frost, Edward Miller, Elias Mueggler, Evgeniy Oleinik, Fan Zhang, Guruprasad Somasundaram, Gustavo Solaira, Harry Lanaras, Henry Howard-Jenkins, Huixuan Tang, Hyo Jin Kim, Jaime Rivera, Ji Luo, Jing Dong, Julian Straub, Kevin Bailey, Kevin Ekenhoff, Lingni Ma, Luis Pesqueira, Mark Schwesinger, Maurizio Monge, Nan Yang, Nick Charron, Nikhil Raina, Omkar Parkhi, Peter Borschowa, Pierre Moulon, Prince Gupta, Raul Mur-Artal, Robbie Pennington, Sachin Kulkarni, Sagar Miglani, Santosh Gondi, Saransh Solanki, Sean Diener, Shangyi Cheng, Simon Green, Steve Saarinen, Suvam Patra, Tassos Mourikis, Thomas Whelan, Tripti Singh, Vasileios Balntas, Vijay Baiyya, Wilson Dreeves, Xiaqing Pan, Yang Lou, Yipu Zhao, Yusuf Mansour, Yuyang Zou, Zhaoyang Lv, Zijian Wang, Mingfei Yan, Carl Ren, Renzo De Nardi, and Richard Newcombe. 2023. Project Aria: A New Tool for Egocentric Multi-Modal AI Research. arXiv:2308.13561 [cs.HC] <https://arxiv.org/abs/2308.13561>
- [26] FFmpeg Developers. 2024. FFmpeg Documentation: ffmpeg Tool. FFmpeg. [Online]. Available: <https://www.ffmpeg.org/ffmpeg.html>.
- [27] Geoffrey A. Fowler. 2022. TSA is Adding Face Recognition at Big Airports. Here’s How to Opt Out. The Washington Post, Dec. 2, 2022. [Online]. Available: www.washingtonpost.com/technology/2022/12/02/tsa-security-face-recognition/.
- [28] Amy Francombe. 2025. Meta Has Already Won The Smart Glasses Race. Wired, Aug. 25, 2025. [Online]. Available: <https://www.wired.com/story/meta-has-already-won-the-smart-glasses-race/>.
- [29] Guo Freeman and Divine Maloney. 2021. Body, Avatar, and Me: The Presentation and Perception of Self in Social Virtual Reality. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 239 (Jan. 2021), 27 pages. <https://doi.org/10.1145/3432938>
- [30] Claire Garvie, Alvaro Bedoya, and Jonathan Frankle. 2016. Perpetual Line Up - Unregulated Police Face Recognition in America. <http://www.perpetuallineup.org/>
- [31] Adam Geitgey. 2017. face_recognition. https://github.com/ageitgey/face_recognition. Accessed: 2025-04-23.
- [32] Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. Doubleday, New York.
- [33] Jens Grubert, Tobias Langlotz, Stefanie Zollmann, and Holger Regenbrecht. 2017. Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 23 (2017), 1706–1724. Issue 6. <https://doi.org/10.1109/TVCG.2016.2543720>
- [34] Kashmir Hill. 2024. Two Students Created Face Recognition Glasses. It Wasn’t Hard. - The New York Times. https://www.nytimes.com/2024/10/24/technology/facial-recognition-glasses-privacy-harvard.html?unlocked_article_code=1.U04.v3IY.7Whhe942GgqL&smid=em-share
- [35] Erin E Hollenbaugh. 2021. Self-presentation in social media: Review and research opportunities. *Review of communication research* 9 (2021), 80.
- [36] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (CHI '03). Association for Computing Machinery, New York, NY, USA, 17–24. <https://doi.org/10.1145/642611.642616>

- [37] Robert Johnson and Adam Cureton. 2024. Kant's Moral Philosophy. In *The Stanford Encyclopedia of Philosophy (Fall 2024 Edition)*, Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Stanford, CA, USA. <https://plato.stanford.edu/archives/fall2024/entries/kant-moral/>
- [38] Norene Kelly and Stephen Gilbert. 2016. The WEAR Scale: Developing a Measure of the Social Acceptability of a Wearable Device. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) (*CHI EA '16*). Association for Computing Machinery, New York, NY, USA, 2864–2871. <https://doi.org/10.1145/2851581.2892331>
- [39] Yoonkyung Kim and Young Min Baek. 2014. When is selective self-presentation effective? An investigation of the moderation effects of "self-esteem" and "social trust". *Cyberpsychology, Behavior, and Social Networking* 17 (11 2014), 697–701. Issue 11. <https://doi.org/10.1089/CYBER.2014.0321/ASSET/IMAGES/LARGE/FIGURE1.JPEG>
- [40] Davis E. King. 2009. Dlib C++ Library. <https://github.com/davisking/dlib>. Accessed: 2025-04-23.
- [41] Davis E. King. 2009. Dlib-ml: A Machine Learning Toolkit. *Journal of Machine Learning Research* 10, 60 (2009), 1755–1758. <http://jmlr.org/papers/v10/king09a.html>
- [42] Alistair Knott, Dino Pedreschi, Toshiya Jitsuzumi, Susan Leavy, David Eyers, Tapabrata Chakraborti, Andrew Trotman, Sundar Sundareswaran, Ricardo Baeza-Yates, Przemyslaw Biecek, Adrian Weller, Paul D. Teal, Subhadip Basu, Mehmet Haklidir, Virginia Morini, Stuart Russell, and Yoshua Bengio. 2024. AI content detection in the emerging information ecosystem: new obligations for media and tech companies. *Ethics and Information Technology* 26 (12 2024), 1–14. Issue 4. <https://doi.org/10.1007/S10676-024-09795-1/METRICS>
- [43] Marion Koelle, Matthias Kranz, and Andreas Möller. 2015. Don't look at me that way! Understanding User Attitudes Towards Data Glasses Usage. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Copenhagen, Denmark) (*MobileHCI '15*). Association for Computing Machinery, New York, NY, USA, 362–372. <https://doi.org/10.1145/2785830.2785842>
- [44] 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* (Glasgow, Scotland Uk) (*CHI EA '19*). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3290607.3312837>
- [45] 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* (Stockholm, Sweden) (*TEI '18*). Association for Computing Machinery, New York, NY, USA, 177–187. <https://doi.org/10.1145/3173225.3173234>
- [46] Mikko Kytö and David McGookin. 2017. Augmenting Multi-Party Face-to-Face Interactions Amongst Strangers with User Generated Content. *Computer Supported Cooperative Work: CSCW: An International Journal* 26 (12 2017), 527–562. Issue 4-6. <https://doi.org/10.1007/S10606-017-9281-1>
- [47] Erik Learned-Miller, Joy Buolamwini, Vicente Ordóñez, and Jamie Morgenstern. 2020. Facial Recognition Technologies In The WILD. The Algorithmic Justice League, [Online]. Available: www.ajl.org/federal-office-call.
- [48] 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* (Honolulu, HI, USA) (*CHI '20*). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376398>
- [49] Eden Litt. 2012. Knock, Knock. Who's There? The Imagined Audience. *Journal of Broadcasting & Electronic Media* 56 (7 2012), 330–345. Issue 3. <https://doi.org/10.1080/08838151.2012.705195>
- [50] Jennifer Lynch. 2020. Face off: Law enforcement use of face recognition technology. Available at SSRN 3909038 NA, NA (2020), 1–39.
- [51] Miriam Madsen, Rana el Kaliouby, Matthew Goodwin, and Rosalind Picard. 2008. Technology for just-in-time in-situ learning of facial affect for persons diagnosed with an autism spectrum disorder. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility* (Halifax, Nova Scotia, Canada) (*Assets '08*). Association for Computing Machinery, New York, NY, USA, 19–26. <https://doi.org/10.1145/1414471.1414477>
- [52] Steve Mann. 2004. "Sousveillance": inverse surveillance in multimedia imaging. In *Proceedings of the 12th Annual ACM International Conference on Multimedia* (New York, NY, USA) (*MULTIMEDIA '04*). Association for Computing Machinery, New York, NY, USA, 620–627. <https://doi.org/10.1145/1027527.1027673>
- [53] Cecily Mauran. 2025. Meta forges ahead with facial recognition for its AI glasses. <https://mashable.com/article/meta-facial-recognition-ai-glasses-privacy-concerns>
- [54] Joseph A. Maxwell. 2010. Using Numbers in Qualitative Research. <http://dx.doi.org/10.1177/1077800410364740> 16 (4 2010), 475–482. Issue 6. <https://doi.org/10.1177/1077800410364740>
- [55] James H. Moor. 2005. Why we need better ethics for emerging technologies. *Ethics and Information Technology* 7 (9 2005), 111–119. Issue 3. <https://doi.org/10.1007/S10676-006-0008-0/METRICS>
- [56] Ivy Munoko, Helen L. Brown-Liburd, and Miklos Vasarhelyi. 2020. The Ethical Implications of Using Artificial Intelligence in Auditing. *Journal of Business Ethics* 167 (11 2020), 209–234. Issue 2. <https://doi.org/10.1007/S10551-019-04407-1/TABLES/6>
- [57] Syed Ibrahim Mustafa Shah Bukhari, Maha Sajid, Bo Ji, and Brendan David-John. 2025. Rethinking Privacy Indicators in Extended Reality: Multimodal Design for Situationally Impaired Bystanders. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*. IEEE, Daejeon, Republic of Korea, 265–272. <https://doi.org/10.1109/ISMAR-Adjunct68609.2025.00059>
- [58] Pauline Norstrom and Anekanta Consulting. 2021. Has Covid increased public faith in facial recognition? *Biometric Technology Today* 2021 (11 2021), 5. Issue 11. [https://doi.org/10.1016/S0969-4765\(21\)00121-1](https://doi.org/10.1016/S0969-4765(21)00121-1)

- [59] Joseph O’Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. 2023. 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 (Jan. 2023), 35 pages. <https://doi.org/10.1145/3569501>
- [60] Octav Opaschi and Radu-Daniel Vatavu. 2020. Uncovering Practical Security and Privacy Threats for Connected Glasses with Embedded Video Cameras. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 4, Article 167 (Dec. 2020), 26 pages. <https://doi.org/10.1145/3432700>
- [61] Kushani Perera, Tobias Langlotz, Nadia Pantidi, and Holger Regenbrecht. 2024. Exploring Eye Visibility and Mutual Gaze in Augmented Reality Glasses. *ACM International Conference Proceeding Series* 1 (12 2024). Issue 1. <https://doi.org/10.1145/3726986.3726999>
- [62] Kushani Perera, Tobias Langlotz, Nadia Pantidi, and Holger Regenbrecht. 2024. What you see is (not necessarily) what I see—Pervasive AR for Public Displays, In Proceedings of the 36th Australian Computer-Human Interaction Conference. *ACM International Conference Proceeding Series* NA, NA. <https://doi.org/10.1145/3726986.3727005>
- [63] Kushani Perera, Holger Regenbrecht, Nadia Pantidi, and Tobias Langlotz. 2025. Don’t Look at Me Like That—How AR Face Recognition Changes Our Social Behaviour. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, Daejeon, South Korea, 1–11. <https://doi.org/10.1109/ISMAR67309.2025.00057>
- [64] Python Software Foundation. 2024. pickle — Python Object Serialization. Python 3 Documentation, [Online]. Available: <https://docs.python.org/3/library/pickle.html>.
- [65] Xiao Qiang. 2019. The Road to Digital Unfreedom: President Xi’s Surveillance State. *Journal of Democracy* 30 (1 2019), 53–67. Issue 1. <https://doi.org/10.1353/JOD.2019.0004>
- [66] Shwetha Rajaram, Jiasi Chen, and Michael Nebeling. 2025. Privacy Equilibrium: Balancing Privacy Needs in Dynamic Multi-User Augmented Reality Scenarios. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST ’25)*. Association for Computing Machinery, New York, NY, USA, Article 151, 24 pages. <https://doi.org/10.1145/3746059.3747783>
- [67] Holger Regenbrecht, Alistair Knott, Jennifer Ferreira, and Nadia Pantidi. 2024. To See and be Seen—Perceived Ethics and Acceptability of Pervasive Augmented Reality. *IEEE Access* 12 (2024), 32618–32636. <https://doi.org/10.1109/ACCESS.2024.3366228>
- [68] Holger Regenbrecht, Sander Zwanenburg, and Tobias Langlotz. 2022. Pervasive Augmented Reality—Technology and Ethics. *IEEE Pervasive Computing* 21 (2022), 84–91. Issue 3. <https://doi.org/10.1109/MPRV.2022.3152993>
- [69] Bradley J. Rhodes. 1997. The wearable remembrance agent: A system for augmented memory. *Personal and Ubiquitous Computing* 1 (1997), 218–224. Issue 4. <https://doi.org/10.1007/BF01682024/METRICS>
- [70] David Robinson, Chris Garratt, and David Robinson. 2008. *Introducing ethics* (fourth ed.). Icon Books, New York, NY, USA, 175 pages. Includes index Previous ed.: published as *Ethics for beginners*. 1996 Formerly CIP UK.
- [71] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57 (2014), 88–96. Issue 4. <https://doi.org/10.1145/2580723.2580730>
- [72] Valentin Schwind and Niels Henze. 2020. Anticipated User Stereotypes Systematically Affect the Social Acceptability of Mobile Devices. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) (*NordiCHI ’20*). Association for Computing Machinery, New York, NY, USA, Article 13, 12 pages. <https://doi.org/10.1145/3419249.3420113>
- [73] Mattea Sim, Basia Radka, Emi Yoshikawa, Franziska Roesner, Kurt Hugenberg, and Tadayoshi Kohno. 2025. To Reveal or Conceal: Privacy and Marginalization in Avatars. *Proceedings on Privacy Enhancing Technologies* 2025 (4 2025), 363–381. Issue 2. <https://doi.org/10.56553/POPETS-2025-0066>
- [74] Brad Singletary and Thad Starner. 2001. Symbiotic Interfaces For Wearable Face Recognition. In *HCI2001 workshop on wearable computing, New Orleans, LA*. NA, NA, NA. <https://api.semanticscholar.org/CorpusID:5928115>
- [75] Alexander Skulmowski and Patricia Engel-Hermann. 2025. The ethics of erroneous AI-generated scientific figures. *Ethics and Information Technology* 27 (6 2025), 1–10. Issue 2. <https://doi.org/10.1007/S10676-025-09835-4/TABLES/2>
- [76] D.J. Solove. 2004. *The Digital Person: Technology and Privacy in the Information Age*. NYU Press, New York, NY, USA. <https://books.google.co.nz/books?id=CuugBwAAQBAJ>
- [77] Daniel J. Solove. 2006. A Taxonomy of Privacy. *University of Pennsylvania Law Review* 154, 3 (2006), 477–564. <https://doi.org/10.2307/40041279>
- [78] Daniel J. Solove. 2010. *Understanding Privacy*. Harvard University Press, Cambridge, MA.
- [79] Bernd Stahl. 2012. Morality, Ethics, and Reflection: A Categorization of Normative IS Research. *Journal of the Association of Information Systems* 13 (08 2012), 636–656. <https://doi.org/10.17705/1jais.00304>
- [80] Bernd Carsten Stahl and Damian Eke. 2024. The ethics of ChatGPT – Exploring the ethical issues of an emerging technology. *International Journal of Information Management* 74 (2024), 102700. <https://doi.org/10.1016/j.ijinfomgt.2023.102700>
- [81] Bernd Carsten Stahl, Job Timmermans, and Catherine Flick. 2017. Ethics of Emerging Information and Communication Technologies: On the implementation of responsible research and innovation. *Science and Public Policy* 44 (6 2017), 369–381. Issue 3. <https://doi.org/10.1093/SCIPOL/SCW069>
- [82] Helen Stefanidi, Jan-Hendrik Sünderkamp, Markus Tatzgern, Alina Itzlinger, and Alexander Meschtscherjakov. 2025. You’re making things AR-kward: Exploring Augmented Reality In-the-Wild. *Proc. ACM Hum.-Comput. Interact.* 9, 5, Article MHC1041 (Sept. 2025), 24 pages. <https://doi.org/10.1145/3743740>
- [83] 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* (Denver, Colorado) (*ETRA ’19*). Association for Computing Machinery, New York, NY, USA, Article 26, 10 pages. <https://doi.org/10.1145/3314111.3319913>

- [84] Ivan Sutherland. 2001. The Ultimate Display. *Proceedings of the IFIPS Congress 65(2):506-508*. New York: IFIP 2 (01 2001).
- [85] David R. Thomas. 2006. A General Inductive Approach for Analyzing Qualitative Evaluation Data. *American Journal of Evaluation* 27 (6 2006), 237–246. Issue 2. <https://doi.org/10.1177/1098214005283748>
- [86] Tanh Quang Tran, Tobias Langlotz, and Holger Regenbrecht. 2024. A Survey On Measuring Presence in Mixed Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (, Honolulu, HI, USA,) (*CHI '24*). Association for Computing Machinery, New York, NY, USA, Article 543, 38 pages. <https://doi.org/10.1145/3613904.3642383>
- [87] Yuzuko Utsumi, Yuya Kato, Kai Kunze, Masakazu Iwamura, and Koichi Kise. 2013. Who are you? A wearable face recognition system to support human memory. In *Proceedings of the 4th Augmented Human International Conference* (Stuttgart, Germany) (*AH '13*). Association for Computing Machinery, New York, NY, USA, 150–153. <https://doi.org/10.1145/2459236.2459262>
- [88] Xi Wang, Xi Zhao, Varun Prakash, Weidong Shi, and Omprakash Gnawali. 2013. Computerized-eyewear based face recognition system for improving social lives of prosopagnosics. In *Proceedings of the 7th International Conference on Pervasive Computing Technologies for Healthcare* (Venice, Italy) (*PervasiveHealth '13*). ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Brussels, BEL, 77–80. <https://doi.org/10.4108/icst.pervasivehealth.2013.252119>
- [89] Brian Wassom. 2014. *Augmented Reality Law, Privacy, and Ethics: Law, Society, and Emerging AR Technologies* (1st ed.). Syngress Publishing, Boston. <https://linkinghub.elsevier.com/retrieve/pii/C20130134137>
- [90] Maximiliane Windl, Petra Zsofia Laboda, and Sven Mayer. 2025. Designing Effective Consent Mechanisms for Spontaneous Interactions in Augmented Reality. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (*CHI '25*). Association for Computing Machinery, New York, NY, USA, Article 1225, 18 pages. <https://doi.org/10.1145/3706598.3713519>
- [91] Katrin Wolf, Albrecht Schmidt, Agon Bexheti, and Marc Langheinrich. 2014. Lifelogging: You're Wearing a Camera? *IEEE Pervasive Computing* 13 (2014), 8–12. Issue 3. <https://doi.org/10.1109/MPRV.2014.53>
- [92] Tim Worstall. 2013. The Killer Google Glass App That Google Won't Let You Have. *Forbes*, Nov. 20, 2013. [Online]. Available: www.forbes.com/sites/timworstall/2013/11/20/the-killer-google-glass-app-that-google-wont-let-you-have/?sh=6630a4f01299.
- [93] David Wright. 2010. A framework for the ethical impact assessment of information technology. *Ethics and Information Technology* 2010 13:3 13 (7 2010), 199–226. Issue 3. <https://doi.org/10.1007/S10676-010-9242-6>
- [94] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks. *IEEE Signal Processing Letters* 23, 10 (2016), 1499–1503. <https://doi.org/10.1109/LSP.2016.2603342>
- [95] Yuhang Zhao, Shaomei Wu, Lindsay Reynolds, and Shiri Azenkot. 2018. A Face Recognition Application for People with Visual Impairments: Understanding Use Beyond the Lab. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3173789>
- [96] Yuhang Zhao, Yaxing Yao, Jiayu Fu, and Nihan Zhou. 2023. "If sighted people know, i should be able to know": privacy perceptions of bystanders with visual impairments around camera-based technology. In *Proceedings of the 32nd USENIX Conference on Security Symposium* (Anaheim, CA, USA) (*SEC '23*). USENIX Association, USA, Article 261, 18 pages.