



Facial Recognition: Understanding Privacy Concerns and Attitudes Across Increasingly Diverse Deployment Scenarios

Shikun Zhang, Yuanyuan Feng, and
Norman Sadeh, *Carnegie Mellon University*

<https://www.usenix.org/conference/soups2021/presentation/zhang-shikun>

This paper is included in the Proceedings of the
Seventeenth Symposium on Usable Privacy and Security.

August 9–10, 2021

978-1-939133-25-0

Open access to the Proceedings of the
Seventeenth Symposium on Usable Privacy
and Security is sponsored by



Facial Recognition: Understanding Privacy Concerns and Attitudes Across Increasingly Diverse Deployment Scenarios

Shikun Zhang
Carnegie Mellon University
Pittsburgh, PA, USA
shikunz@andrew.cmu.edu

Yuanyuan Feng
Carnegie Mellon University
Pittsburgh, PA, USA
yuanyuanfeng@cmu.edu

Norman Sadeh
Carnegie Mellon University
Pittsburgh, PA, USA
sadeh@cs.cmu.edu

Abstract

The rapid growth of facial recognition technology across ever more diverse contexts calls for a better understanding of how people feel about these deployments — whether they see value in them or are concerned about their privacy, and to what extent they have generally grown accustomed to them. We present a qualitative analysis of data gathered as part of a 10-day experience sampling study with 123 participants who were presented with realistic deployment scenarios of facial recognition as they went about their daily lives. Responses capturing their attitudes towards these deployments were collected both in situ and through daily evening surveys, in which participants were asked to reflect on their experiences and reactions. Ten follow-up interviews were conducted to further triangulate the data from the study. Our results highlight both the perceived benefits and concerns people express when faced with different facial recognition deployment scenarios. Participants reported concerns about the accuracy of the technology, including possible bias in its analysis, privacy concerns about the type of information being collected or inferred, and more generally, the dragnet effect resulting from the widespread deployment. Based on our findings, we discuss strategies and guidelines for informing the deployment of facial recognition, particularly focusing on ensuring that people are given adequate levels of transparency and control.

1 Introduction

We live in a world full of cameras, from traditional closed-circuit televisions to the latest motion-sensing wireless IP

cameras. According to a report by IHS Markit, a total of over one billion cameras are expected to be deployed worldwide by 2021 [25]. Existing security and surveillance cameras can be easily augmented with facial recognition, a type of artificial intelligence (AI)-enabled video analytics technology that has become increasingly accurate with recent advances in deep learning and computer vision [43]. The U.S. Government Accountability Office (GAO) broadly defines facial recognition technology as computer applications that (1) detect faces in an image or video, (2) estimate a person’s demographic characteristics (e.g., age, race, gender) (3) verify a person’s identity by accepting or denying the claimed identity, and (4) identify an individual by matching an image of them to a database of known people [105]. Extensions of facial recognition also include facial expression recognition [87], mood detection, scene detection (e.g., identifying petty crime [85]), and more. In this paper, we adopt this broader definition of facial recognition.

In recent years, facial recognition has been widely deployed in public places, such as airports for security and surveillance purposes [42, 103], department stores for automatic detection of known shoplifters, rental car companies for self-checkout [37, 74]. While facial recognition technology can contribute to security, productivity, convenience, and more, its broad deployment also gives rise to serious privacy concerns [97]. These concerns have prompted increased scrutiny from both privacy advocates and regulators [24, 30, 60]. Recent studies have also reported limitations and flaws of facial recognition technology, including unsatisfactory levels of accuracy as well as bias towards underrepresented demographic groups and members of the LGBTQ+ community [5, 47, 59, 81]. Both policymakers and researchers have also expressed concerns about abusive uses of the technology, e.g., non-consensual surveillance [48, 49].

Our research focuses on the perceptions and attitudes of people (or “data subjects”) whose presence and activities can be captured by facial recognition technologies. This paper describes the results of an exploratory qualitative analysis of responses gathered as part of a 10-day experience sam-

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2021.
August 8–10, 2021, Virtual Conference.

pling study. The study involved asking participants to install a study app on their regular smartphones and using the app to present them with a range of realistic facial recognition scenarios at venues they visited during their everyday activities. The app was used to collect their reactions to these different scenarios. Data collected *in situ* was supplemented with additional information collected as part of a daily evening survey, in which participants were asked to review each of the deployment scenarios presented to them during the day and answer additional questions. Moreover, we analyzed 123 participants' post-survey responses and also interviewed 10 of them. This paper is the sequel to another publication on this study, where we presented a quantitative analysis of participants' privacy preferences and expectations in responses to these scenarios [115].¹ Through an in-depth analysis of the qualitative data collected from this study, we aim to develop a more holistic understanding of people's perception of the benefits and concerns associated with the diverse deployment scenarios considered in this study. In particular, we further contextualize participants' perception of privacy risks associated with facial recognition and explore their concerns about the limitations and bias found in some of these systems.

This article's contributions fall under three broad categories:

- We present an in-depth qualitative analysis of lay people's perceptions towards facial recognition. Our qualitative dataset contains both interview data and 123 participants' free-text responses over a 10-day period, which provides a comprehensive view of individuals' perceptions towards facial recognition.
- To our knowledge, this is the first qualitative study that uses carefully designed and realistic facial recognition deployment scenarios and differentiates between diverse attributes of the technologies (e.g., purpose, venue, type of analysis, data sharing) and to do so *in situ*, as participants went about their regular everyday activities.
- Based on our results, we propose guidelines and design recommendations for trustworthy deployments of facial recognition technology.

2 Related Work

2.1 Facial Recognition and Algorithmic Bias

Facial Recognition (FR) and its wide range of applications have been a prevailing research topic for decades. Traditional FR methods are mostly feature-based and are limited in their discriminant power [6, 14, 50, 67, 113]. Recent deep learning approaches have significantly boosted the performance

¹ See also [116] for results exploring the use of machine learning models to help predict people's privacy preferences (i.e., opt-in/opt-out preferences for different scenarios) and alleviate the user burden of exercising privacy choices.

of facial recognition models [26, 76, 90, 104], enabling it to approach and surpass human performance on FR benchmarks [45, 55, 58]. Despite the impressive progress, there are still many problems with FR. A series of reports on testing commercial facial recognition software conducted by the National Institute of Standards and Technology (NIST) revealed that software accuracy variations and potential bias existed for different demographic groups [4, 43, 44]. Several studies have tried to quantify the demographic biases of some of these deep learning models [17, 56, 81], with sources of bias attributed to unrepresentative data distributions in training sets [56] and to the use of certain optimization methods [93]. Besides issues related to accuracy and bias, prior studies have also questioned the effectiveness of emotion detection, which falls under the broad definition of facial recognition, exposing problems with classifiers trained on artificial displays of emotions failing to capture people's true inner states [12, 69]. These limitations can in turn lead to the mistreatment of certain demographic groups, exposing them to higher individual or societal risks, or impeding their access to some services [111]. Even though these problems have been acknowledged by the computing community [5] and legal scholars [53, 112], no research has been conducted to understand people's awareness and perception of these limitations and the risks they entail. Our work aims to fill this gap.

2.2 Attitudes towards Facial Recognition

A few prior studies have examined people's attitudes towards facial recognition through surveys [21, 97, 100, 101]. The Pew Research Center conducted a nationally representative survey on Americans' awareness and acceptance of facial recognition. They found that Americans in general trust law enforcement to use facial recognition responsibly more than technology companies and advertisers and that these attitudes also vary across demographic groups [97]. Another study further analyzed the Pew survey data and focused on gendered perceptions of workplace surveillance. This study found that women were less likely to accept the use of facial recognition in the workplace [100]. The Center for Data Innovation also conducted a national online poll through Google Surveys and found that fewer Americans think the government should limit the use of facial recognition [21]. A few interview studies have focused on specific functionalities of facial recognition [8] and the impact of facial recognition technology on marginalized demographic groups [47]. Hamidi et al. found transgender individuals have overwhelmingly negative attitudes towards recognition algorithms that automatically detect gender [47]. Andalibi et al. discussed users' attitudes towards emotion recognition technology, including perceptions of individual and societal risks [8]. Our work, which does not focus on the relationship between demographics and attitudes, sheds light on people's concerns about facial recognition across a variety of scenarios without targeting

any particular demographic group. Facial recognition has also attracted the attention of law scholars who have closely examined the legal and ethical issues of the emerging facial recognition through a slew of law review articles [53, 72, 112]. Our work complements these legal reviews by presenting and analyzing data collected from our study participants, relying on their own accounts of perceived threats and benefits associated with these deployments in realistic contexts experienced as part of their regular everyday activities.

2.3 Privacy Challenges of Facial Recognition

Facial recognition technology can be used to capture a variety of sensitive information about people, from biometric data (e.g., facial features and body pose) [38, 90, 104] to information about people's activities (e.g., where they are, whom they are with, and what they do) [38, 114] all the way to their emotions (e.g., attentive, depressed, and surprised) [64]. While people may notice some cameras, they have no way of knowing how captured footage is being processed (e.g., what types of algorithms might be run and for what purpose) and what happens to the data being captured (e.g., whom it is shared with and for how long data might be retained). The loss of privacy resulting from the deployment of this technology has been a common thread in the literature [19, 70, 78, 79]. Researchers have examined technical solutions to safeguard user data [31, 32, 34, 78, 89], including algorithms to avoid being tracked by facial recognition [94, 95], and systems to enable real-time opt-out of facial recognition systems [27, 28, 88]. But how to increase transparency around data privacy remains an unsolved issue [22, 82, 83].

In this paper, we explore three research questions:

- RQ1: What are users' attitudes towards facial recognition technology, and why?
- RQ2: What are some benefits and concerns people associate with facial recognition deployment scenarios?
- RQ3: What recommendations can we develop for the trustworthy deployment of facial recognition?

3 Methodology

3.1 Study Design

Prior work shows that context plays a critical role in influencing people's privacy attitudes and decisions [75]. In order to solicit realistic participant feedback, we designed an experience sampling study to collect people's responses to a variety of facial recognition deployments (or "scenarios") in the context of their regular everyday activities. The experience sampling method [51] has been successfully used in many real-life studies [20, 39, 54, 61, 84, 106, 107], enhancing the ecological validity of the results [13, 92].

In the 10-day experience sampling study, we presented participants with facial recognition scenarios that were likely to happen at places they visited as part of their daily activities. For example, when a participant visited a gym, they may be presented with a scenario where facial recognition was used to track their attendance. The scenarios included in the study were informed by an extensive survey of news articles about real-world deployments of facial recognition in a variety of contexts, i.e., identification of known criminals [2, 23, 40, 57], petty crime detection [85], operation optimization by businesses [71, 77, 86], demographic-based advertising [9, 35, 98], advertising based on reactions [15, 18, 91], engagement detection [63, 68, 110], ID/loyalty card replacement [10, 33, 73, 96], attendance tracking [3, 11, 41], health-related predictions [7, 66, 80], productivity predictions [29, 62], and medical diagnoses [1, 36, 46, 65].

3.2 Study Procedures

The 10-day study was carried out in the following steps. First, eligible participants who completed the consent forms could download the in-house study app from the Google Play Store. Second, while participants went about their regular daily activities, the study app collected the GPS location of their smartphones. As participants visited places for which we had plausible scenarios, the app would send them a push notification, prompting them to complete a short survey on a facial recognition scenario pertaining to their location. Third, at the end of each day, participants also received an email in the evening to answer a daily summary web survey ("evening review"). This web survey showed participants the places they visited when they received notifications, probed reasons for their in-situ answers, and asked additional questions. See Appendix 7.4 for screenshots of the app and an example of the evening review. Fourth, after completing 10 days of evening reviews, participants answered a post-survey where they provided open-ended text responses about their attitudes on facial recognition technology and their perceived beneficial and concerning contexts where facial recognition was applied. Fifth, we conducted semi-structured interviews with 10 participants over online video conferencing software (e.g., Skype, Google Hangouts) after they have completed the study. The full text of the post-survey, the scenarios presented during the study, and the interview scripts can be found in the Appendix.

3.3 Recruitment and Participants

We recruited participants from both online and offline channels. Our recruitment messages were posted on a variety of online platforms, including local online forums (i.e., Craigslist and Reddit), a university-based research platform, and a promotional Facebook advertisement. We also put up flyers on bus stops and local community bulletin boards. A short screening survey was used to determine participants' eligibility

(aged 18 or older, able to speak English, using an Android smartphone with a data plan). We also collected demographic information such as age, gender, and occupation in the screening survey. We avoided convenience samples of university students and collected data from a diverse pool of participants. A total of 164 participants downloaded our study app, and 123 of them completed our 10-day study and the post-survey. The demographics of the 123 participants is shown in Table 1 and 2. We sent out 17 invitations to participants who showed interest in participating in the follow-up interview and conducted online interviews with 10 participants who responded. This study was approved by our university's IRB and the human research protection office of the funding agency.

3.4 Qualitative Dataset

In this work, we focused on analyzing the qualitative dataset collected from the 10-day experience sampling study. The dataset includes 2,562 entries of text responses from participants' daily summaries, 1,230 entries of text responses in the post-survey, and 10 interview transcripts. In order to answer the research questions, it is crucial that the qualitative data collected reflects participants' attitudes towards facial recognition. Since we adopted an experience sampling method presenting realistic scenarios of facial recognition to participants over 10 days, we believe the data collected following these contextual cues would capture participants' perceptions and attitudes. We did not report other quantitative data collected from the experience sampling study since they are not the focus of this paper.

3.5 Interview Data Analysis

The interviews ranged from 26 to 40 minutes (mean=33) and were fully transcribed. A total of 326 minutes of transcripts were analyzed. One author first read and familiarized herself with all the transcripts. She then applied thematic analysis [16] to open code the transcripts. The second author met with the first author regularly to iterate on the themes.

3.6 Content Analysis of Textual Responses

From the 10-day study, we collected 2,562 entries of text responses from participants' daily summaries and 1,230 entries from the post-survey. In the post-survey, there were 10 open-ended questions. The first question was "What is the first thing that comes to your mind when you think about facial recognition technology?" We coded the sentiment (i.e., positive, negative, neutral, mixed) in each response.

We included two questions in the post-survey asking participants' perceived beneficial and concerning contexts to use facial recognition technology. We also asked questions eliciting participants' privacy concerns about facial recognition deployment scenarios. After reading the survey responses, we

realized many participants shared their attitudes and experiences with facial recognition deployment scenarios regardless of to which question they were responding. Since the daily summaries were also addressing similar issues, in our analysis, we broke down the boundaries between the data sources and conducted a content analysis [102] of all the participants' 3792 textual responses.

Two authors started from inductive coding [16] to extract codes that show participants' perceived benefits or concerns about facial recognition technology and developed a codebook. In total, we summarized 13 main codes with 32 subcodes about the benefits of facial recognition and 19 main codes with 40 subcodes about the concerns. In the end, we used a deductive coding approach, applying the codebook to the entire dataset. Two authors independently coded all data and met to resolve any discrepancies.

4 Findings

In this section, we present findings from qualitative analysis of interview and textual response data collected from evening reviews of in-situ scenarios participants received. We first present findings on participants' attitudes towards facial recognition technology and the reasons behind their attitudes. We then show the perceived beneficial and concerning contexts of facial recognition usage. We also unveil participants' concerns about the use of facial recognition, with a particular focus on privacy-specific concerns, as they are among the most prominent themes. Finally, we flesh out participants' proposed actions in responses to these deployment scenarios.

4.1 Impressions of Facial Recognition

We first present findings on participants' sentiment towards facial recognition technology. This is based on our coding of sentiment in participants' responses to the first question in the post-survey: "What is the first thing that comes to your mind when you think about facial recognition technology?"

4.1.1 Participants tend to be more negative towards FR

We observed that participants tended to be more negative towards the use of facial recognition: 51 (42%) participants displayed negative impressions while only 13 (11%) expressed positive sentiments. The negative connotation mostly revolves around problems of the technology, like the infringement on their right to privacy. Those negative first impressions also echo entrenched perceptions on problematic usages and privacy risks of facial recognition that are revealed in our subsequent analysis in Section 4.4.2 and 4.5.

Among the 13 participants with positive impressions, most praised facial recognition's usefulness, like its ability to increase public safety and catch criminals. A few also mentioned the "*advancement in technology*" (P36, positive). We

Gender	%	Age	%	Education	%	Income	%	Marital Status	%
Female	57.7	18-24 years old	8.1	Some high school	.8	Less than \$25,000	14.6	Single, never married	50.4
Male	40.7	25-34 years old	54.5	High School	4.1	\$25,000 to \$34,999	14.6	Married	41.5
Other	1.6	35-44 years old	23.6	Some college	13.8	\$35,000 to \$49,999	9.8	Separated	1.6
		45-54 years old	8.1	Associate's degree	7.3	\$50,000 to \$74,999	22.0	Divorced	3.3
		55-64 years old	3.3	Bachelor's Degree	35.0	\$75,000 to \$99,999	14.6	Widowed	0.8
		65-74 years old	2.4	Master's Degree	23.6	\$100,000 to \$149,999	14.6	I prefer not to answer	2.4
				More than Master's Degree	12.8	\$150,000 to \$249,999	2.4		
				Other	1.6	I prefer not to answer	7.3		

Table 1: Survey participant demographics and respective %

Occupation	%	Occupation	%
Business, or sales	12.2	Legal	3.3
Administrative support	9.8	Other	3.3
Scientist	8.9	Graduate student	2.4
Service	8.1	Skilled labor	2.4
Education	8.1	Homemaker	2.4
Computer engineer or IT	7.3	Retired	2.4
Other salaried contractor	7.3	Government	1.6
Engineer in other fields	6.5	Prefer not to say	1.6
Medical	6.5	Art or writing	.8
Unemployed	4.1	College student	.8

Table 2: Occupations of survey participants and respective %

also noted a mixed perspective of facial recognition from 11 (9%) respondents: *“It’s invasive and big brother esque. It can provide good information for law enforcement but is easily abusable”* (P83, mixed). 48 participants (39%) indicated their neutral impressions typically by describing main use cases or depicting how facial recognition works: *“the ability of computers to see normal people in plain view and identify their identity. This can then be passed to another decision-making system for a distinct purpose: law enforcement, advertising, efficiency, etc.”* (P12, neutral).

4.1.2 Participant views may be influenced by media portrayals

A few concepts also emerged from these responses, mostly related to media portrayals of facial recognition. Some participants were reminded of what they have watched in the movies or crime shows relating to facial recognition: *“I think of face scanners and searches people do when looking for criminals in crime tv shows and movies”* (P42, neutral). Other respondents made references to a dystopian world, with many citing the concept of Big Brother from the book 1984 — *“Cyberpunk dystopias, “Big Brother,” and similar instances in fiction, satire, and socio-political discussion about invasion of privacy on the part of powerful political and economic entities”* (P39, negative). China was brought up 7 times as the example of a surveillance state, which was associated with more negative sentiments (5 out of 7) than neutral tones (2

out of 7). For example, P80 alluded to a negative use case, *“China and the way they micromanage their citizens lives,”* and P5 expressed a more neutral impression: *“I think of China because the only times I’ve seen it on the news, it was being used in China.”*

In summary, respondents expressed more negative views about facial recognition than positive ones. Many were wary about potential problems linked to the technology. Around a quarter of participants’ views were influenced by the media portrayal of facial recognition (e.g., news, movies, and books).

4.2 Beneficial and Concerning Contexts

We present findings on users’ perceived beneficial and concerning use of facial recognition. This is based on the deductive coding of textual responses to the questions asking participants to identify up to 5 contexts each where they found the use of facial recognition technology to be beneficial and concerning. On average, each participant identified 2.7 ± 1.4 beneficial contexts, and 3.0 ± 1.4 concerning contexts. A paired Wilcoxon signed-rank test showed that participants recorded significantly more concerning contexts than beneficial contexts ($Z = 2.65, p < 0.01, r = 0.24$).

The findings are organized based on the major codes in the codebook, as shown in Table 3. These codes were further categorized into two groups: purposes for using facial recognition and entities that use facial recognition. We first report beneficial and concerning purposes in this subsection.

4.2.1 Beneficial purposes: security, authentication, and commerce

The majority (104 out of 123) of participants reported that security is a beneficial context for facial recognition. Among those, 42% thought that facial recognition could increase public security in general, and 32% thought that it is beneficial to use facial recognition to identify and catch criminals. Another important context for security, raised by 20%, is to find missing individuals. For example, P26 mentioned that facial recognition could be helpful in *“locating missing/abducted children and adults.”* 13% of them also mentioned that facial

recognition could be beneficial to deter crime, as expressed by P27 *“in public, especially in isolated places like parking garages, to help preserve women’s safety.”* Another context for facial recognition that 51 participants (42%) identified as beneficial is authentication. About half of them (24 out of 51) stated that facial recognition could be used to replace IDs and confirm identity. 31% mentioned that it could be used to log in devices and/or replace passwords. A quarter maintained that facial recognition could be useful to grant access in secured locations, which P46 described as *“helping identify people in high-security areas.”* 14% considered authentication in stores via facial recognition as a way to replace membership or reward cards to be beneficial as well. A sizable minority (27 out of 123 — 22%) of participants also saw merits in leveraging facial recognition in commercial settings; using facial recognition to improve services and tailor customer experiences was deemed beneficial by about half of those 27 participants, for example, in contexts like *“relocating people between the crowded check-out areas”* (P63) and *“customization of service based on who you are and known preferences”* (P55). Others considered marketing and tailored advertisement of potential benefit, like in *“retail scenarios (catered advertising)”* (P46) and *“providing information to retail companies about their customers”* (P111).

4.2.2 Concerning purposes: advertisement, profiling, and prediction

Most participants (64%) raised concerns about various purposes for which facial recognition is used. Specifically, 36 out of 123 (29%) participants found using facial recognition for advertisement troubling: P117 said, *“It can be used for marketing and branding purposes that are generally antagonistic.”* 18 participants were concerned about facial recognition used for profiling — *“using it to profile someone based on race or gender”* (P21). 17 respondents found *“when emotion recognition is in use”* to be concerning. 12 participants (10%) were specifically against their data being sold for profit *“random companies selling and profiting off of it”* (P40). 11 were worried about use cases of facial recognition that involves predicting or estimating intentions or behaviors — *“Any assessments that are psychologically based since there is a lot that could be wrongly inferred by only taking into account visual data”* (P12).

4.3 Beneficial and Concerning Entities

The right-hand side of Table 3 shows the percentages of participants who identified different entities (law/government, employers, etc.) as beneficial and/or concerning when they deploy facial recognition.

Purpose				Entity			
Beneficial		Concerning		Beneficial		Concerning	
Code	%	Code	%	Code	%	Code	%
Security	84.6	Ads	29.3	Law/Gov	14.6	Law/Gov	18.7
Authentication	41.5	Profiling	14.6	Public	11.4	Employer	17.1
Commercial	22.0	Emotion	13.8	Health	8.1	Business	15.4
Personal	9.8	Profit	9.6	Employer	5.7	Insurer	14.6
Other	8.1	Predictive	8.9	Myself	5.7	Health	7.3
		Security	5.7	Business	4.9		

Table 3: Codes from Content Analysis and the Percentages of Participants Who Mentioned Them

4.3.1 Weighing between beneficial versus concerning

It is interesting to observe that law enforcement/the government were deemed concerning and beneficial both by a sizable number of respondents, which is also similar in the case of health-related entities (e.g., hospitals and clinics). The neck-and-neck numbers seem to suggest that those entities entail both rather apparent pros and cons of using facial recognition. For example, *“law enforcement falsely accusing someone”* (P83) is rather concerning, while facial recognition aids *“law enforcement to track and apprehend criminals”* (P42) is clearly beneficial. On the other hand, significantly more participants considered businesses, employers, or health insurers’ use of facial recognition more concerning than beneficial. More participants see harm than benefit brought by facial recognition usages by these entities, as elaborated by P59, *“The data collected seems worth more to the company than any coupons could possibly be for me.”*

4.3.2 Attributes influencing attitudes towards entities

The interview data revealed in-depth deliberations participants had while weighing various entities obtaining their facial recognition information. **Trust** was one of the factors that can erase participants’ doubts about potentially questionable facial recognition usages. Two interviewees explained why they trust their employer or the government/law enforcement, therefore trusting their use of facial recognition. P55 explained, *“I trust my manager personally to have my own interests in heart...Right now, personally, I have a good relationship with my manager and with the company. So I am pretty comfortable with what they do, decide to do, and feel like that they are not going to use it against me.”* Believing in the democratic government, P57 maintained, *“The government supposedly is “by the people, for the people” as supposed to private corporations...So if it’s used by law enforcement, I am a bit more comfortable with that.”*

More evidence on trust being an influential factor also emerged in the answers from evening surveys: *“Because law enforcement and the government have a history of using data for purposes other than what they were intended for or what we were told it was for”* (P26), *“I don’t trust insurance companies to make fair decisions”* (P116), *“I trust the library mostly not to do anything bad with the video”* (P97), *“This is*

a large entity that I trust”(P51), and etc.

Besides trust, whether entities that deploy facial recognition have **control** over data subjects is another important attribute. Three interviewees were reserved about their employer or the government using this technology as those entities intrinsically have more **control** over them. In their views, facial recognition can be used against them by powerful entities, such as governments, employers, and big corporations, as expressed in the following quotes.

“I am used to people that advertise to me, trying to sell me something...I have more control over that relationship because I can always turn down buying something, even with coercive tactics that are manipulative. But with my boss or the government, I don’t have the power in that relationship at all. So it’s more information for them that they can use against me basically.” — P50

“I mean whoever’s behind it [facial recognition] has more data and information, what people need, what individual person wants, and how to best serve the people around, like get their product to the people. And also they have more control...over their customers.” — P52

Three interviewees were worried about advertisers’ or corporations’ usage that could decrease their sense of **autonomy**. Thanks to facial recognition technologies, businesses would leverage highly fine-grained and even real-time data to improve their marketing techniques. For example, P56 expressed her concern, *“With the ability to read your reactions and then be able to market responses specifically to you, you might be losing some free choice. Because they are able to pinpoint and push harder things they think are important to you, because you are reacting to them, they can get real-time reactions to products...They can start using terms that look like something and trick you into buying something.”* Such practices can be manipulative and encroach on people’s freedom.

4.4 Concerns About Facial Recognition

4.4.1 Participants were concerned about facial recognition even for anonymous demographic detection

Current facial recognition software enables different levels of identification: some can recognize the shape of faces and humans; some can detect specific demographic features; others can match faces to images of people stored in databases. Demographic detection has been used in contexts like targeted advertising and marketing [9, 35, 40, 98].

When designing the study, we initially conjectured that people would be more comfortable with anonymous demographic detection than personally identifiable detection. Nonetheless, 9% of participants expressed reservations about using anonymous demographic detection for advertising as they saw it as a form of profiling. P50 explicitly pointed out, *“I was also pretty concerned when the notifications popped up about predicting purchases based on racial classifications because*

that just seemed very racist to me. Just because someone is African American or Hispanic, you can’t predict what they are going to want to buy based on their race; that seems a really not very good policy.”

Others were really against gender-based advertising. For example, P50 mentioned, *“And gender, there is such a spectrum, just because you’re female, that doesn’t mean you are going to wanna wear makeup or buy pretty dresses. Same thing for guys. I just think lumping every person into a classification is over-generalized; you are going to miss people.”* Some participants questioned the efficacy of advertising based on gender and race, *“I wouldn’t think it will be very accurate, you could target something to me being white that would not at all relate to me still based on that one factor. But it may relate to a non-white person. I think it wouldn’t even be accurate. I think you need a lot more than race and gender to advertise to someone effectively”* (P106). This type of practices, even though beneficial at times, can also reinforce existing gender and cultural stereotypes — *“I understand that some ethnic groups might benefit from this (for instance, African American women need specific hair care products that aren’t always easy to find.) But I am concerned about the potential for misuse of this technology to discriminate. Also, people don’t always “look like” the racial or ethnic background with which they identify”* (P27).

Some participants, including some parents, were leery of age-based advertising, especially worrying about kids being susceptible to those practices. *“Things are marketed to kids nowadays, and kids can buy things on apps without their parents even knowing...I don’t think they should be marketed towards kids necessarily”* (P50). We also observed reservations from participants who were afraid of being labeled as a specific demographic group, such as religious groups. P53 said, *“I think it is kind of dangerous to pinpoint one person as part of a group vs. just the individual. So I think the times I was most concerned during the research was when I would go to someplace that was religious[ly] affiliated or like a non-profit organization. If there was a video of me and my friends maybe at a church or at a Jewish organization. Does that put us more in danger if we are associated with that group? I feel like there is this danger of having a label placed on you, and if the wrong person gets that information, and that could be a catalyst for violence.”* P89 summarized her feelings towards demographic-based facial recognition, *“I do think it will divide us more if they are targeting specifically based on what you look like, not even necessarily your profile and who you are...I think it just gives an overall weird and gross feeling, especially in today’s society where it comes up a lot.”*

4.4.2 Participants were worried about incorrect detection and interpretation

About a third of the participants reported their concerns about the accuracy of facial recognition during the study. Some

were worried about the technology not accurate enough and could make “mistakes in the face recognition (twins, relatives)” (P65). One interviewee P107 shared his firsthand experience with inaccurate facial recognition in details, “I don’t know how accurate they would be based on stuff that I have tested out before. Like even with having a beard, it throws off a lot of things that try to guess things. Actually, at work, just for fun, one of the guys had it. It is for visually impaired people who are blind. It scans anything and tells you what it is. It scans faces and got a lot of people like “39 male,” and it would be really close, but when it comes to me, it would say 40 where I am 25. It would say frowning even though I am smiling because of it tracking the mustache...if they are trying to pick up people with negative emotions for security purposes, maybe it could be pretty wrong.” Others also echoed their doubts about the accuracy of emotion detection, like P68 “I don’t see how it (emotion analysis) could be that accurate unless you are monitoring what I am saying too. Like I said, I went through a breakup that week, and sometimes I was not in a good mood no matter where I was, no matter how good the food was. How are they supposed to know? It just seemed like it was an unnecessary addition that wouldn’t end up being very accurate.”

In addition to questioning how accurate facial recognition can be, some participants also argued that seemingly suspicious behavior, when viewed out of context, can be misinterpreted by those systems, potentially resulting in grave consequences. For example, P53 described one such scenario in her friend’s life that could be misconstrued, “I think a lot of the times like my friend she locked herself out of her apartment this past weekend, so she tried to jump in through her window. So if a recognition program saw that, they might think that it is a thief or criminal or whatever. And that is not the case. She is not breaking into her own house. It needs to be able to interpret scenarios correctly. It needs to be able to have a context for them. Not just to assume that something looks like a criminal act is a criminal act.” Similarly, P68 gave another example, “I think it could misinterpret scenarios, it could misinterpret the guy trying to break into his own car to get his keys out, or the boyfriend putting his hand in the girlfriend’s pocket.” An interviewee P57 was worried about such inaccuracies leading to deadly consequences — “because if someone was marked for shoplifting and they didn’t do, that could cost a lot of trouble, in some scenarios that could cost someone’s life.”

4.4.3 Participants were concerned about racial and other biases introduced by facial recognition

One-tenth of our participants reported being concerned about potential bias in the facial recognition systems, especially about the deep implications it might have on minority groups. Many were worried that racial bias in these algorithms could exacerbate the entrenched bias and infringe upon the rights

of those impacted groups. Two interviewees’ elaborated accounts provide us with more insights: P68 stated, “Any system I’ve seen has inevitably been used only to profile people of color and the LGBTQ+ community. I think even if we have this surveillance, somebody is like, “Oh, it is just gonna automatically detect petty crimes.” The reality is that it will still be looking harder at a black person and their actions to see if that is a petty crime than it could with a white person. I still think at the end of the day, a human is gonna analyze the data. I think you still have a lot of misidentification where people of color and LGBTQ+ community members are going to be scrutinized more strongly, not given the benefit of the doubt that white people are.” Similarly, P53 noted, “I wouldn’t want a program like that to decide that for example, a black man equals thief or even to give a warning sign to a program to flag that because that is not the case. So I think that is the danger of having that type of use for facial recognition. I think it can too easily be biased, intentionally or unintentionally. The person programming it might think that they might have statistics to back up the demographics of thieves or demographics of criminals, but I don’t think that is a good way of deciding who is or who is not a criminal.”

4.5 Perceived Privacy Risks of FR

Privacy is repeatedly brought up as a key concern by our study participants. Around 70% of participants voiced privacy concerns during the study. In this section, we summarize the major themes around perceived privacy risks of facial recognition, in light of concepts from established privacy frameworks (i.e., Solove’s “Taxonomy of Privacy” [99] and Westin’s states of privacy from *Privacy and Freedom* [109]).

4.5.1 Violation of solitude

The feeling of surveillance prevails A third of our respondents found surveillance through facial recognition to be concerning. Surveillance can exert adverse psychological effects like discomfort and anxiety on subjects. For example, P68 pointed out that “I had this paranoia that I would be judged based on every action I took at work without the full context.” Similarly, P29 stated that “always being watched and analyzed which in itself is scary.” Moreover, surveillance is also harmful due to its infringement on people’s freedom to act. P89 contextualized this concern — “There is a feeling of freedom as I enter the library where I participate in a Spanish speaking group on Wednesday morning...in the small classroom where we speak, I would feel rather self-conscious if I were videoed.” This infringement upon freedom can also possibly lead to inhibition and behavior alteration, as P84 noted “I’d always have to be concerned about how my actions might be perceived on camera,” and in P20’s view, “I want to know where all of the cameras are, so I can always be aware and I can always be on guard and vigilant. So if

something happens, I can be ready to defend myself or defend the findings.” Surveillance can also have a chilling effect on civil and political engagement. For instance, P117 pointed out that facial recognition *“is used to identify anti-fascists and peaceful protesters”*, and P39 found *“any and all efforts at using such technology against political dissenters”* to be concerning.

Deprived of the right to be let alone Warren and Brandeis first articulated privacy as the “right to be let alone” [108]. Privacy risks also lie in the probing action itself which perturbs this right, making “the person being questioned feel uncomfortable” as noted by Solove [99]. Two-fifths of our participants regarded some deployment scenarios of facial recognition as unwarranted and prying. For instance, P68 manifested their concern, *“It is the idea of somebody being able to surveil my life and know my business...Even though on sight it’s something different through a camera, that knowing somebody is interested in the data, and wants it, and is just getting it for free. Something about it really bothers me.”* Some participants responded to data collection of facial recognition rather abruptly, *“It’s none of anyone’s business, as long as I’m obeying the law, where I am and what I’m doing”* (P114). Some participants reported that facial recognition is intrusive into one’s life, and they cannot be let alone under the presence of facial recognition. For example, P83 mentioned that they are *“unable to hide from people”*, and P104 noted, *“I feel like I’m being stalked by the man, the powers that be, wealthy corporations.”* Others regarded facial recognition as disruptions to their daily activities: P69 mentioned, *“Don’t want to be filmed eating,”* and P62 commented on their experience in stores, *“It’s like being stared at in the face by someone while I’m just trying to shop.”*

4.5.2 Unwanted exposure and violation of anonymity

Not able to stay anonymous 17% of participants stressed the importance of anonymity and scrutinized how facial recognition enabled the identification of normal people in plain view. P63 gave examples of circumstances when people may want to stay anonymous, *“Probably if you go to some kind of clinics, like sexual health clinics, or food pantry.”* P12 voiced their concerns about facial recognition used for advertising, *“If it is generating tailored advertising then it implies it is tracking my shopping habits and linking it to my face.”* P55 elaborated a situation when he wants to remain anonymous, *“I don’t do any sort of very secretive things. The only possible scenarios are if I was trying to...plan a surprise birthday party for my wife, some notification got sent to both of us of where I was, and then she figures that out...There is a mixed scenario of people who are doing slightly illegitimate things but are legal to do, like having affairs with people on their partners, they would definitely not like stuff like that.”* Identification, a method to connect people to collected data, is hard to avoid

as the deployment of facial recognition technologies becomes widespread.

Unwanted exposure to others This issue involves “exposing to others of certain physical and emotional attributes about a person,” which often “creates embarrassment and humiliation” as defined by Solove [99]. 22% of our participants pointed out that it is easy to reveal emotions under contexts of facial recognition involving emotion recognition. For example, P68 described her personal experience, *“I went through a breakup that week. I was really emotional a lot of the time. I do not want my health insurance, my employer, my parents getting updates like ‘hey, she’s trying to get through the pain while she is working today.’”* P50 commented on a facial recognition scenario that occurred at the vet they went to, *“People experience deep personal emotions at the vet.”* Some respondents were cautious about carrying out private actions. For example, P89 elaborated, *“I might be caught at the gym entering and adjusting a bra strap, etc.,”* and *“doing something like picking your nose, something like that, not doing something against the law, but something you don’t want others to see.”* Such unwanted exposure in public spaces might not have been feasible without facial recognition technologies.

4.5.3 Non-consensual and insecure disclosure

Secondary use without consent This refers to the privacy issue of data collected for additional purposes without data subjects’ knowledge or consent. In the context of facial recognition, this problem is exacerbated because of the lack of ways to properly convey data practices to subjects other than using signs that say “face recognition security cameras in use.” Given the sensitive nature of facial recognition data, around a quarter of our participants reported concerns about unauthorized secondary use. Many respondents questioned whether companies would retain data for intended use only, as P12 described, *“As I’m doing this study more, I think it’s my trust in their ability to safe keep the data and only for that use. I would doubt their compliance even if I do want them to get the competitive advantage by the use of video surveillance.”* P89 also hoped for regulations to prevent secondary use, *“If there were laws in place that they could never ever use it for anything else like they couldn’t sell it to marketing companies.”* P106 provided a concrete example of secondary use with regards to workplaces using facial recognition to track attendance, *“I think if used to replace a time card is fine, but I could see it being abused by overbearing managers.”* A few participants expressed concerns about their data being sold, which can also be regarded as a secondary use.

Fear of data leakage and abuse About one-third of our participants expressed their concerns about their facial recognition data being hacked or abused. Because it is almost im-

practical to relinquish biometric data when compromised, the security of facial recognition data is ever more pressing. Many of the participants reported that they do not trust data collectors' ability to safeguard their data. For example, P122 noted, *"I don't think data security is a strong priority for these companies, and when they do have data leaks, they don't care because it doesn't affect them, and the punishment is not enough to incentivize them to change their practices,"* which parallels the concerns of P54 about identified frivolous activities being leaked, *"Frivolities that end up being insecure, like entertainment or stores."* Also, the fear of insecurity can induce privacy risks by placing people to whom it pertains in a vulnerable state, as corroborated by P122, *"It's very troubling to think of how this info could be used by bad actors."*

4.5.4 Inaccurate dissemination and violation of reserve

Dissemination of inaccurate or misleading information

Around one-third of our participants were concerned about the dissemination of inaccurate or misleading information [99]. This issue is also mostly linked to the inaccuracies of facial recognition as presented in Section 4.4.2. Our participants were concerned about being falsely identified or judged out of context. For example, P46 noted, *"Bad luck or timing could lead law enforcement to be suspicious of an innocent citizen."* P11 referred to their experiences when shopping in stores, *"I would really not like supposedly meaningful data to be recorded if I happened to smile remembering something while walking down the condom aisle."* Distortion can be detrimental, as illustrated by P59, *"Reputational damage could occur if someone is falsely accused of a crime."*

Decisional interference Solove defined this as the intrusion on private decisional making, especially by the government [99]. In our study, participants mostly focused on the unwarranted influence on their purchasing autonomy by private companies with the help of facial recognition. This is also discussed in Section 4.3.2. In addition, P89 lamented, *"It's machines taking over and my freedom circumvented."* P122 echoed this thought, *"I do not want to have this information used against me or used to try and subvert my thinking."*

4.6 Proposed Actions and Responses

Our qualitative data also reveals participants' reported desire to take action when encountering facial recognition in their everyday life. They also express a desire for transparency and indicate they would like to be notified about nearby deployments of facial recognition technology. At the same time, their notification preferences vary with some participants expressing concerns about potentially overly disruptive notifications.

4.6.1 Participants want transparency and control over the collection of their data

About 30% of participants expressed strong views about the need for entities collecting sensitive facial recognition data to notify them and to actively obtain consent from them before data collection. For example, P50 commented, *"I think if they are going to record our image, they should have to notify you before they do anything with it like if they are going to use it for a specific purpose, we should be able to know what they are using it for, and we should be able to say 'yes, that's fine,' or 'no, it's not. Delete my stuff from your system.'"* While most participants agreed about the need to obtain consent, they did not provide consistent answers with regard to the frequencies of such notifications. Some participants wanted to be notified every time when such data collection is taking place, as illustrated by the quote from P56, *"I think it is important to know when you are in areas where data is being collected, passive consent really disturbs me. I know it happens all the time when I am on my phone or computer, and it is really hard to know what data is being collected, what it is being used for, etc....So, if I have my preference, I would want to know every time someone is engaging in this practice,"* whereas others were wary of repeated reminders and preferred less frequent notices, as P17 elaborated, *"I frequent this establishment pretty often, so a constant reminder would annoy me. It would be nice to be reminded every now and then in case I simply forget."* These results suggest a need for customizable notification functionality where different individuals can select from a number of notification options.

4.6.2 Participants find existing notice mechanisms inadequate

While the majority of participants wanted to be informed about facial recognition in use, our follow-up interviews disclosed the specific ways how some participants found the existing notice mechanisms inadequate. For instance, P68 described how they missed the existing signs in physical spaces that were supposed to notify them about the presence of cameras, *"There will be places where I would want to be notified every time, and then I look over, and see a sign that I have just passed by a dozen times, and realize I am being notified."* When probed about what is a good way to give them notice or obtain their consent, some interviewees reported that no existing mechanisms would achieve the goal, as P53 said, *"I think that [obtaining consent] is hard...It is hard because you cannot pass a form when you walk into a restaurant or a store, it cannot be formal...I guess trying to do it remotely like through the Internet or your phone would be the easiest."* Specifically, P50 expressed their desire to provide consent based on different purposes of facial recognition, *"It would depend on what they were using it for. If it was just like someone committed a crime, and they needed FR for that, then that's fine. Maybe if it's to replace a swipe card or a membership cards, that would*

be okay, but if it's for tracking my purchases, or tracking my attendance, emotions." The information on the purposes for which facial recognition is deployed is not available to data subjects in the majority of current deployments. Also, it is also hard to design notice mechanisms with the desired level of intrusiveness, as P89 elaborated, *"I would not want to think about it at all times, so I want it to be subtle whatever the notification is, but also not so subtle that you don't know that it is happening ever,"* which highlights the problem of privacy as a secondary goal.

4.6.3 Some participants fear being overwhelmed by frequent notifications

While most participants report that they want to be notified, more than half are also weary of too frequent notifications. In particular, some participants realized during the course of the study that the number of notifications they would receive might become a nuisance if they request to be notified each time they get within range of facial recognition technology. For instance, P53 described her thought process, *"When I first started, I was saying once in a while, and then I realized that would be really annoying to get multiple notifications."* About half (55 out of 123) of participants reported that they were unlikely to avoid places that deploy facial recognition technology, even if they indicated being concerned about these deployments, revealing a general sense of resignation. For instance, P11 underscored, *"There is nothing I can do about it, and this is the only accessible grocery around my workplace, so I don't have an alternative."* A similar sense of helplessness and resignation was expressed by P67: *"I give up. Spy on me. What can I do about it? I'm old. I'll be dead soon."*

At the same time, not all participants reported concern. We also observed a small number of participants who did not care about the usage of facial recognition in general, referencing the "nothing to hide" argument. For instance, P55 elaborated, *"I am not likely to be so concerned about it, because I don't do any sort of very secretive things... There are more legitimate reasons why people would want to value their privacy more than I do, but I am not sure how much of the population that would really affect."*

5 Discussion

5.1 Limitations

We would like to first remind the reader that the results presented in this paper focus on a qualitative analysis of data collected as part of our study. A sister publication presented earlier this year provides a quantitative analysis of additional data collected as part of the same experience sampling study [115]. We invite the reader to look at it for additional details about our study protocol and to develop a more comprehensive view of our findings.

We acknowledge that, while ideally, we would have liked to collect data from a representative cross-section of the general public, study participants were recruited from the population of a mid-sized city in the United States (Pittsburgh). Our sample is skewed towards somewhat younger and more educated participants, which might have biased some of our findings. Accordingly, we do not claim that our results are representative of the general population. In addition, our analysis results rely on participants' self-reported qualitative data, which may not necessarily match their actual behaviors.

While describing study scenarios, we strove to maintain a balanced narrative without overly emphasizing benefits or potential risks associated with different deployments. We acknowledge that on occasions, our phrasing might inadvertently have primed participants in one direction or the other.

Finally, our participants generally expressed somewhat negative views of various facial recognition deployment scenarios. This could, in part, be a reflection of the fact that they did not actually experience true interactions with these deployment scenarios and, as a result, may not have had a chance to appreciate what they consider as benefits associated with some of these scenarios (e.g., marketing scenarios).

5.2 Combating Inaccuracy and Bias

While most of participants reported seeing benefits in facial recognition deployments such as security and authentication scenarios, their reported attitude towards many other scenarios was generally more negative. Part of their willingness to embrace the technology was dampened by concerns over accuracy and bias of facial recognition systems, echoing concerns voiced by marginalized interviewees in a prior study [47]. Our data suggest that these concerns extend to the more general population. Recent reports of people wrongly arrested due to faulty facial recognition algorithms likely contributed to reservations captured in our study [52] and also illustrate the severe consequences that deployment of this technology can have if deployed and relied upon without adequate safeguards. Minimally, technology should be evaluated for potential biases and minimal levels of accuracy, especially when deployed in support of particularly sensitive activities such as law enforcement. Their performance and limitations should be clearly communicated and taken into account. And decisions based on these algorithms should be meticulously cross-checked and manually vetted if we are to avoid more of these nightmarish scenarios.

5.3 Contextualizing Perceived Privacy Risks

Our analysis organized perceived privacy risks associated with facial recognition deployments around key dimensions identified in well-established privacy frameworks [99, 108, 109]. We were able to elicit more nuanced and contextualized privacy concerns than prior work [21, 97, 100] as shown in

Section 4.5. While legal arguments support people’s reasonable expectations of privacy in public places [53], our study provides strong evidence that these expectations are real and widespread and that some facial recognition deployment scenarios are perceived as overstepping the boundaries of personal solitude, making people feel deprived of “the(ir) right to be let alone” [108]. These concerns are further exacerbated by the sensitive nature of biometric data, the information that can be inferred from facial recognition data (e.g., location, activity, and mood), as well as risks of secondary use of this data and its security. These findings underscore the need for more transparency in notifying people about not just the deployment of facial recognition technology but also sufficient details for individuals to gauge their perceived privacy risks.

5.4 Designing Effective Notice and Choice

Our study confirms that privacy concerns are a major obstacle to acceptance of a variety of facial recognition scenarios [21, 22, 83], although these deployments are becoming increasingly widespread. Responses from our participants indicate a strong desire to be notified about different deployment scenarios and to have some control over the collection and analysis of their data. Current deployments generally fall short when it comes to effectively notifying people about the presence of facial recognition technologies, including details about the type of analysis they rely on and how results are being used and possibly shared. Also, current deployments generally fail to provide people with opt-in or opt-out choices.

How to effectively notify people and offer them adequate controls is not trivial. Entities deploying facial recognition should inform data subjects in a clear and noticeable manner. Today’s “this area under camera surveillance” signs do not provide them with enough information, such as type of analysis, the purpose for collection and analysis, sharing, etc. Privacy controls (e.g., opt-in and opt-out choices) should obviously include mechanisms to authenticate data subjects (to make sure they are whom they claim to be when they request to opt in or out of some practices), giving rise to privacy issues. With the possible exception of security-related deployments, which many view as generally beneficial, people should be offered some control over the collection and use of their footage — preferably in the form of opt-ins.

One solution involves requiring people to opt in by providing training data about their face [27, 28]. In this system, a privacy-aware infrastructure is used to notify people about the presence of nearby facial recognition deployments, including who has deployed the technology, what analysis is performed, and for how long the footage is retained. Users who do not opt in for facial recognition by default have their face (or possibly their entire body) obfuscated in real-time in the captured footage. Notifications about nearby facial recognition deployment are provided via a “Privacy Assistant” mobile app that users install on their smartphones. This infrastructure

has been deployed to support notice and choice for a variety of Internet of Things data collection processes — not just facial recognition [27, 88].

Our data highlight individuals’ diverse notification preferences, with some preferring to be systematically notified about FR deployments, while others only would prefer just occasional notices and reminders. The Internet of Things Privacy Infrastructure introduced by Das et al. offers users of its “Privacy Assistant” mobile app different settings they can configure to specify the types of data collection processes they want to be notified about as well as the frequency of these notifications (e.g., “only the first time,” “every time,” or “never”). These settings are consistent with results discussed in Section 4.6, which indicate that different participants have different notification preferences and that these preferences can also evolve. Further research is needed to determine what personalized settings are likely to work best and how to alleviate the user burden that might be entailed by opt-in or opt-out settings associated with a potentially large number of facial recognition deployments.

Finally, our study indicates that participants fear losing their autonomy when commercial entities can assemble and leverage near real-time facial recognition data, including their emotions, to tailor advertisements presented to them. Our participants also expressed reservations about the power this technology can bestow on already powerful entities such as their employers or law enforcement authorities. These results further emphasize the need for more effective notice and choice mechanisms if people are to become less fearful about the deployment of facial recognition.

6 Conclusion

Deployment of facial recognition technologies is already widespread and continuing to grow. While many people are familiar with typical video surveillance scenarios, most have little or no awareness of the increasingly diverse set of scenarios where this technology is being deployed. We analyzed data from a 10-day in-situ study where we collected information about people’s awareness and perceptions of a variety of facial recognition deployments they could realistically encounter as part of their everyday activities. Our data show that people’s privacy concerns are complex and depend on different attributes characterizing these deployment scenarios. Our analysis reveals serious concerns about the privacy impact of these technologies, including the lack of mechanisms to effectively notify people and give them some control over the collection, analysis, and use of their footage. Our data also suggest that people’s views about facial recognition technologies have been impacted by recent reports about the inconsistent accuracy and bias found in deployed systems. The qualitative analysis presented in this paper complements a quantitative analysis of data collected as part of the same study presented in a recent sister publication [115].

Acknowledgments

We thank Dr. Xu Wang and Zheng Yao for their input on this paper. We also thank Dr. Lujo Bauer, Dr. Lorrie Cranor, and Dr. Anupam Das for their feedback on the study. This research has been supported in part by DARPA and AFRL under agreement number FA8750-15-2-0277 and in part by NSF under grants from the National Science Foundation Secure and Trustworthy Computing program (CNS-15-13957, CNS-1801316). The US Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notice thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied of DARPA, AFRL, NSF or the US Government.

References

- [1] Augmented mental health: Revolutionary mental health care using emotion recognition. <https://www.augmentedmentalhealth.com/blog/augmented-mental-health-revolutionary-mental-health-care-using-emotion-recognition>, 2018.
- [2] Chinese man caught by facial recognition at pop concert. <https://www.bbc.com/news/world-asia-china-43751276>, 2018.
- [3] Facial recognition: School ID checks lead to GDPR fine. <https://www.bbc.com/news/technology-49489154>, 2019.
- [4] Facial recognition technology: Ensuring transparency in government use. <https://www.nist.gov/speech-testimony/facial-recognition-technology-ensuring-transparency-government-use>, 2019.
- [5] ACM U.S. Technology Policy Committee. Statement on principles and prerequisites for the development, evaluation and use of unbiased facial recognition technologies. <https://www.acm.org/binaries/content/assets/public-policy/ustpc-facial-recognition-tech-statement.pdf>, 2020.
- [6] Timo Ahonen, Abdenour Hadid, and Matti Pietikäinen. Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:2037–2041, 2006.
- [7] Marshall Allen. Health insurers are vacuuming up details about you — and it could raise your rates. <https://www.propublica.org/article/health-insurers-are-vacuuming-up-details-about-you-and-it-could-raise-your-rates>, 2018.
- [8] Nazanin Andalibi and Justin Buss. The human in emotion recognition on social media: Attitudes, outcomes, risks. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, pages 1–16, New York, NY, USA, 2020. ACM.
- [9] Association Press. Tesco’s plan to tailor adverts via facial recognition stokes privacy fears. <https://www.theguardian.com/business/2013/nov/03/privacy-tesco-scan-customers-faces>, 2013.
- [10] Rachel Bachman. Your gym’s tech wants to know you better. <https://www.wsj.com/articles/your-gyms-tech-wants-to-know-you-better-1497281915>, 2017.
- [11] Sarah Pulliam Bailey. Skipping church? Facial recognition software could be tracking you. <http://www.washingtonpost.com/news/acts-of-faith/wp/2015/07/24/skipping-church-facial-recognition-software-could-be-tracking-you/>, 2015.
- [12] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychol Sci Public Interest.*, 20(3):165–166, 2019.
- [13] Daniel J Beal. ESM 2.0: State of the art and future potential of experience sampling methods in organizational research. *Annual review of organizational psychology and organizational behavior*, 2(1):383–407, 2015.
- [14] Peter N. Belhumeur, João P. Hespanha, and David J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:711–720, 1997.
- [15] Bloomberg News. Mannequins collect data on shoppers via facial-recognition software. https://www.washingtonpost.com/business/economy/mannequins-collect-data-on-shoppers-via-facial-recognition-software/2012/11/22/0751b992-3425-11e2-9cfa-e41bac906cc9_story.html, 2012.
- [16] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
- [17] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. volume 81 of *Proceedings of Machine Learning Research*, pages 77–91, New York, NY, USA, 23–24 Feb 2018. PMLR.

- [18] David Burrows. Facial expressions show Mars the adverts that will drive sales. <https://www.foodnavigator.com/Article/2017/03/23/Facial-expressions-show-Mars-the-adverts-that-will-drive-sales>, 2017.
- [19] Ramon Caceres and Adrian Friday. Ubicomp systems at 20: Progress, opportunities, and challenges. *IEEE Pervasive Computing*, 11(1):14–21, 2011.
- [20] Laura L. Carstensen et al. Emotional experience improves with age: Evidence based on over 10 years of experience sampling. *Psychology and Aging*, 26(1):21, 2011.
- [21] Daniel Castro and McLaughlin Michael. Survey: Few Americans want government to limit use of facial recognition technology, particularly for public safety or airport screening. <https://datainnovation.org/2019/01/survey-few-americans-want-government-to-limit-use-of-facial-recognition-technology-particularly-for-public-safety-or-airport-screening/>, 2019.
- [22] Richard Chow. The last mile for IoT privacy. *IEEE Security & Privacy*, 15(6):73–76, 2017.
- [23] Ben Conarck. Florida court: Prosecutors had no obligation to turn over facial recognition evidence. <https://www.jacksonville.com/news/20190123/florida-court-prosecutors-had-no-obligation-to-turn-over-facial-recognition-evidence>, 2019.
- [24] Kate Conger, Richard Fausset, and Serge F. Kovalski. San Francisco bans facial recognition technology. <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html>, 2019.
- [25] Elly Cosgrove. One billion surveillance cameras will be watching around the world in 2021, a new study says. <https://www.cnn.com/2019/12/06/one-billion-surveillance-cameras-will-be-watching-globally-in-2021.html>, 2019.
- [26] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-FCN: Object detection via region-based fully convolutional networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16, page 379–387, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [27] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. Personalized privacy assistants for the Internet of Things: Providing users with notice and choice. *IEEE Pervasive Computing*, 17(3):35–46, 2018.
- [28] Anupam Das et al. Assisting users in a world full of cameras: A privacy-aware infrastructure for computer vision applications. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1387–1396. IEEE, 2017.
- [29] Bobby J Davidson. How your business can benefit from facial recognition technology. <https://percentotech.com/how-your-business-can-benefit-from-facial-recognition-technology/>, 2019.
- [30] Dean DeChiaro. New York City eyes regulation of facial recognition technology. <https://www.rollcall.com/news/congress/new-york-city-eyes-regulation-of-facial-recognition-technology>, 2019.
- [31] Benchaa Djellali, Kheira Belarbi, Abdallah Chouarfia, and Pascal Lorenz. User authentication scheme preserving anonymity for ubiquitous devices. *Security and Communication Networks*, 8(17):3131–3141, 2015.
- [32] Yitao Duan and John Canny. Protecting user data in ubiquitous computing: Towards trustworthy environments. In *International Workshop on Privacy Enhancing Technologies*, pages 167–185. Springer, 2004.
- [33] Melanie Ehrenkranz. Burger joint teams up with surveillance giant to scan your face for loyalty points. <https://gizmodo.com/burger-joint-teams-up-with-surveillance-giant-to-scan-y-1821498988>, 2017.
- [34] Zekeriya Erkin et al. Privacy-preserving face recognition. In Ian Goldberg and Mikhail J. Atallah, editors, *Privacy Enhancing Technologies*, pages 235–253, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [35] Darrell Etherington. Baidu and KFC’s new smart restaurant suggests what to order based on your face. <https://techcrunch.com/2016/12/23/baidu-and-kfcs-new-smart-restaurant-suggests-what-to-order-based-on-your-face/>, 2016.
- [36] Ingrid Fadelli. Analyzing spoken language and 3-D facial expressions to measure depression severity. <https://techxplore.com/news/2018-11-spoken-language-d-facial-depression.html>, 2019.
- [37] Caitlin Fairchild. Hertz is now using facial recognition to check out cars. <https://www.nextgov.com/emerging-tech/2018/12/hertz-now-using-facial-recognition-check-out-cars/153479/>, 2018.
- [38] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.

- [39] Denzil Ferreira, Jorge Goncalves, Vassilis Kostakos, Louise Barkhuus, and Anind K. Dey. Contextual experience sampling of mobile application micro-usage. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI '14)*, pages 91–100, 2014.
- [40] Chris Frey. Revealed: how facial recognition has invaded shops—and your privacy. <https://www.theguardian.com/cities/2016/mar/03/revealed-facial-recognition-software-infiltrating-cities-saks-toronto>, 2016.
- [41] Sarah Fister Gale. Employers turn to biometric technology to track attendance. <https://www.workforce.com/news/employers-turn-to-biometric-technology-to-track-attendance>, 2013.
- [42] Shirin Ghaffary and Rani Molla. Here’s where the us government is using facial recognition technology to surveil Americans. <https://www.vox.com/recode/2019/7/18/20698307/facial-recognition-technology-us-government-fight-for-the-future>, 2019.
- [43] Patrick Grother, Mei Ngan, and Kayee Hanaoka. Ongoing face recognition vendor test (FRVT) part 2: Identification. <https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8238.pdf>, 2018.
- [44] Patrick Grother, Mei Ngan, and Kayee Hanaoka. Ongoing face recognition vendor test (FRVT) part 1: Verification. https://www.nist.gov/system/files/documents/2019/11/20/frvt_report_2019_11_19_0.pdf, 2019.
- [45] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. MS-Celeb-1M: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016.
- [46] Yaron Gurovich et al. Identifying facial phenotypes of genetic disorders using deep learning. *Nature Medicine*, 25(1):60–64, 2019.
- [47] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham. *Gender Recognition or Gender Reductionism? The Social Implications of Embedded Gender Recognition Systems*, page 1–13. ACM, New York, NY, USA, 2018.
- [48] Drew Harwell. Amazon extends ban on police use of its facial recognition technology indefinitely. <https://www.washingtonpost.com/technology/2021/05/18/amazon-facial-recognition-ban/>, 2021.
- [49] Drew Harwell. Senators seek limits on some facial-recognition use by police, energizing surveillance technology debate. <https://www.washingtonpost.com/technology/2021/04/21/data-surveillance-bill/>, 2021.
- [50] Xiaofei He, Shuicheng Yan, Yuxiao Hu, Partha Niyogi, and Hong-Jiang Zhang. Face recognition using Laplacianfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:328–340, 2005.
- [51] Joel M. Hektner, Jennifer A. Schmidt, and Mihaly Csikszentmihalyi. *Experience sampling method: Measuring the quality of everyday life*. Sage, 2007.
- [52] Kashmir Hill. Wrongfully accused by an algorithm. <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>, 2020.
- [53] Mariko Hirose. Privacy in public spaces: The reasonable expectation of privacy against the dragnet use of facial recognition technology. *Connecticut Law Review*, (377), 2017.
- [54] Wilhelm Hofmann, Roy F. Baumeister, Georg Förster, and Kathleen D. Vohs. Everyday temptations: An experience sampling study of desire, conflict, and self-control. *Journal of Personality and Social Psychology*, 102(6):1318, 2012.
- [55] Gary B. Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. 2008.
- [56] Isabelle Hupont and Carles Fernández. DemogPairs: Quantifying the impact of demographic imbalance in deep face recognition. *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, pages 1–7, 2019.
- [57] Timothy Johnson. Shoplifters meet their match as retailers deploy facial recognition cameras. <https://www.mcclatchydc.com/news/nation-world/national/article211455924.html>, 2018.
- [58] Ira Kemelmacher-Shlizerman, Steven M. Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4873–4882, 2016.
- [59] Os Keyes. The misgendering machines: Trans/hci implications of automatic gender recognition. *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW), 2018.
- [60] Mehreen Khan. EU plans sweeping regulation of facial recognition. <https://www.ft.com/content/90ce2dce-c413-11e9-a8e9-296ca66511c9>, 2019.
- [61] Ingrid Kramer et al. A therapeutic application of the experience sampling method in the treatment of depression: a randomized controlled trial. *World Psychiatry*, 13(1):68–77, 2014.

- [62] Sarah Krouse. The new ways your boss is spying on you. <https://www.wsj.com/articles/the-new-ways-your-boss-is-spying-on-you-11563528604>, 2019.
- [63] Stephen Lepitak. Disney’s Dumbo and Accenture Interactive collaborate for the movie poster of the future. <https://www.thedrum.com/news/2019/03/10/disneys-dumbo-and-accenture-interactive-collaborate-the-movie-poster-the-future>, 2019.
- [64] Gil Levi and Tal Hassner. Emotion recognition in the wild via convolutional neural networks and mapped binary patterns. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, page 503–510, New York, NY, USA, 2015. ACM.
- [65] David Levine. What high-tech tools are available to fight depression? <https://health.usnews.com/health-care/patient-advice/articles/2017-10-06/what-high-tech-tools-are-available-to-fight-depression>, 2017.
- [66] David Levine. What your face may tell lenders about whether you’re creditworthy. <https://www.wsj.com/articles/what-your-face-may-tell-lenders-about-whether-youre-creditworthy-11560218700>, 2019.
- [67] Chengjun Liu and Harry Wechsler. Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition. *IEEE Transactions on Image Processing*, 11 4:467–76, 2002.
- [68] Brain Logan. Pay-per-laugh: the comedy club that charges punters having fun. <https://www.theguardian.com/stage/2014/oct/14/standup-comedy-pay-per-laugh-charge-barcelona>, 2014.
- [69] Martin Magdin, L’ubomír Benko, and Štefan Koprda. A case study of facial emotion classification using Affdex. *Sensors (Basel)*, 19(9):2140, 2019.
- [70] Tobias Matzner. Why privacy is not enough privacy in the context of “ubiquitous computing” and “big data”. *Journal of Information, Communication and Ethics in Society*, 12(2):93–106, 2014.
- [71] Darren Murph. SceneTap app analyzes pubs and clubs in real-time, probably won’t score you a Jersey Shore cameo. <https://www.engadget.com/2011/06/12/scenetap-app-analyzes-pubs-and-clubs-in-real-time-probably-won/>, 2011.
- [72] Sharon Nakar and Dov Greenbaum. Now you see me. Now you still do: Facial recognition technology and the growing lack of privacy. *Boston University Journal of Science & Technology Law*, (23):88–122, 2017.
- [73] NEC Corporation. New biometric identification tools used in theme parks. <https://www.nec.com/en/global/about/mitatv/03/3.html>, 2002.
- [74] Alfred Ng. With facial recognition, shoplifting may get you banned in places you’ve never been. <https://www.cnet.com/news/with-facial-recognition-shoplifting-may-get-you-banned-in-places-youve-never-been/>, 2019.
- [75] Helen Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- [76] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015.
- [77] PCMag Stuff. NEC unveils facial-recognition system to identify shoppers. <https://www.pcmag.com/archive/nec-unveils-facial-recognition-system-to-identify-shoppers-305015>, 2012.
- [78] Charith Perera, Rajiv Ranjan, Lizhe Wang, Samee U Khan, and Albert Y. Zomaya. Big data privacy in the Internet of Things era. *IT Professional*, 17(3):32–39, 2015.
- [79] Salil Prabhakar, Sharath Pankanti, and Anil K. Jain. Biometric recognition: Security and privacy concerns. *IEEE Security & Privacy*, 1(2):33–42, 2003.
- [80] Emilee Rader. Most Americans don’t realize what companies can predict from their data. <https://bigthink.com/technology-innovation/most-americans-dont-realize-what-companies-can-predict-from-their-data-2629911919>, 2019.
- [81] Inioluwa D. Raji et al. Saving face: Investigating the ethical concerns of facial recognition auditing. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’20, pages 145–151, New York, NY, USA, 2020. ACM.
- [82] Edith Ramirez, Julie Brill, Maureen K. Ohlhausen, Joshua D. Wright, and Terrell McSweeney. Data brokers: A call for transparency and accountability. Technical report, Federal Trade Commission, 2014.
- [83] Luis Felipe M. Ramos. Evaluating privacy during the covid-19 public health emergency: The case of facial recognition technologies. In *Proceedings of the 13th International Conference on Theory and Practice of Electronic Governance*, ICEGOV 2020, page 176–179, New York, NY, USA, 2020. ACM.
- [84] Robert W. Reeder et al. An experience sampling study of user reactions to browser warnings in the field. In

Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, pages 1–13, 2018.

- [85] Timothy Revell. Computer vision algorithms pick out petty crime in CCTV footage. <https://www.newscientist.com/article/2116970-computer-vision-algorithms-pick-out-petty-crime-in-cctv-footage/>, 2017.
- [86] David Rosen. Disney is spying on you! https://www.salon.com/test/2013/01/17/disney_is_spying_on_you/, 2013.
- [87] Andrew Ryan et al. Automated facial expression recognition system. In *43rd Annual 2009 International Carnahan Conference on Security Technology*, pages 172–177, 2009.
- [88] Norman Sadeh. Design of a privacy infrastructure for the internet of things. In *2020 USENIX Conference on Privacy Engineering Practice and Respect (PEPR 20)*. USENIX Association, 2020.
- [89] T. Scott Saponas, Jonathan Lester, Carl Hartung, Sameer Agarwal, and Tadayoshi Kohno. Devices that tell on you: Privacy trends in consumer ubiquitous computing. In *16th USENIX Security Symposium (USENIX Security '07)*, pages 55–70, 2007.
- [90] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015.
- [91] E. J. Schultz. Facial-recognition lets marketers gauge consumers’ real responses to ads. <https://adage.com/article/digital/facial-recognition-lets-marketers-gauge-real-responses/298635>, 2015.
- [92] Christie N. Scollon, Chu Kim-Prieto, and Ed Diener. Experience sampling: Promises and pitfalls, strengths and weaknesses. *Journal of Happiness Studies*, 4(1):5–34, 2003.
- [93] Ignacio Serna et al. SensitiveLoss: Improving accuracy and fairness of face representations with discrimination-aware deep learning. *arXiv*, abs/2004.11246, 2020.
- [94] Shawn Shan et al. Fawkes: Protecting privacy against unauthorized deep learning models. In *29th USENIX Security Symposium (USENIX Security '20)*, pages 1589–1604, 2020.
- [95] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*, pages 1528–1540, 2016.
- [96] Ed Silverstein. New Konami casino facial recognition technology could rival reward cards. <https://www.casino.org/news/new-konami-casino-facial-recognition-technology-could-rival-reward-cards/>, 2019.
- [97] Arron Smith. More than half of U.S. adults trust law enforcement to use facial recognition responsibly. Technical report, Pew Research Center, 2019.
- [98] Benjamin Snyder. This beer ad only works when women pass by. <https://fortune.com/2015/05/21/astra-beer-ad/>, 2015.
- [99] Daniel J. Solove. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3):477–564, 2006.
- [100] Luke Stark, Amanda Stanhaus, and Denise L. Anthony. “I don’t want someone to watch me while I’m working”: Gendered views of facial recognition technology in workplace surveillance. *Journal of the Association for Information Science and Technology*, 71(9):1074–1088, 2020.
- [101] Léa Steinacker, Miriam Meckel, Genia Kostka, and Damian Borth. Facial recognition: A cross-national survey on public acceptance, privacy, and discrimination. In *International Conference on Machine Learning - Law and ML Workshop*, 2020.
- [102] Steve Stemler. An overview of content analysis. *Practical assessment, research, and evaluation*, 7(1):17, 2000.
- [103] Francesca Street. How facial recognition is taking over airports. <https://www.cnn.com/travel/article/airports-facial-recognition/index.html>, 2019.
- [104] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, 2014.
- [105] U.S. Government Accountability Office. Facial recognition technology: Commercial uses, privacy issues, and applicable federal law. <https://www.gao.gov/products/GAO-15-621>, 2015.
- [106] Niels Van Berkel, Denzil Ferreira, and Vassilis Kostakos. The experience sampling method on mobile devices. *ACM Computing Surveys*, 50(6):1–40, 2017.

- [107] Simone J. W. Verhagen, Laila Hasmi, Marjan Drukker, Jim van Os, and Philippe A. E. G. Delespaul. Use of the experience sampling method in the context of clinical trials. *Evidence-based Mental Health*, 19(3):86–89, 2016.
 - [108] Samuel D. Warren and Louis D. Brandeis. The right to privacy. *Harvard Law Review*, 4(5):193–220, 1890.
 - [109] Alan F. Westin. Privacy and freedom. *Washington & Lee Law Review*, 25:166, 1968.
 - [110] Jason Whitely. How facial recognition technology is being used, from police to a soccer museum. <https://www.wfaa.com/article/features/originals/how-facial-recognition-technology-is-being-used-from-police-to-a-soccer-museum/287-618278039>, 2018.
 - [111] Niels Wouters et al. Biometric mirror: Exploring ethical opinions towards facial analysis and automated decision-making. In *Proceedings of the 2019 on Designing Interactive Systems Conference*, DIS '19, page 447–461, New York, NY, USA, 2019. ACM.
 - [112] Elisa Wright. The future of facial recognition is not fully known: Developing privacy and security regulatory mechanisms for facial recognition in the retail sector. *Fordham Intellectual Property, Media & Entertainment Law Journal*, 29(2):611–685, 2019.
 - [113] John Wright, Allen Y. Yang, Arvind Ganesh, S. Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31:210–227, 2009.
 - [114] Huijuan Xu, Abir Das, and Kate Saenko. R-C3D: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
 - [115] Shikun Zhang et al. "Did you know this camera tracks your mood?": Understanding privacy expectations and preferences in the age of video analytics. *Proc. Priv. Enhancing Technol.*, 2021(2):282–304, 2021.
 - [116] Shikun Aerin Zhang et al. Understanding people's privacy attitudes towards video analytics technologies. Technical Report CMU-ISR-20-114, Carnegie Mellon University, School of Computer Science, 2020.
- We asked: How surprised would you be about [PLACE] engaging in this data practice? At the time, you indicated that you would find this _____. Why?
 - We asked: How comfortable would you feel about [PLACE] engaging in this data practice? At the time, you indicated that you would find this _____. Why?
 - We asked: How would you want to be notified as you enter [PLACE]? At the time, you indicated that you _____. Why?
 - If you had the choice, would you allow or deny this data practice?
 - Based on the data practice description above, do you believe the footage in which you appear could be made available to third parties for analysis with facial recognition?
 - Please indicate how much you agree or disagree with each of the following statements.
 - I feel that I benefit from this data practice
 - I feel that [PLACE] benefits from this data practice
 - I feel that the data practice enhances public safety
 - How would you feel about the raw footage being shared with the following entities?

7.2 Post Study Survey

- What is the first thing that comes to your mind when you think about facial recognition technology?
- In what context(s) do you find the use of facial recognition technology to be particularly beneficial? (Enter up to 5 types of contexts)
- In what context(s) do you find the use of facial recognition technology to be particularly concerning? (Enter up to 5 types of contexts)
- Do you feel that you have a general understanding of where this type of technology is likely to be used and why?
- Please rate your comfort level when visiting stores and other locations that use facial recognition technology.
- How likely would you be to intentionally avoid stores that use facial recognition technology?
- Has your level of concern about facial recognition technology changed over the course of the study?
- 10 IUIPC Questions
- Show scenarios: Petty Crime/Sentiment Ads(IDed)/Health Predictions
 - Within what timeframe, do you believe this data practice will be commonplace?
 - Would you like to be notified about this data practice?
 - What sensitive information do you think could be inferred from this data collection practice?
 - How concerned would you be about this sensitive information being inferred? Why?
 - How likely would you be to avoid visiting those places following the introduction of this data practice?
 - What do you think is a reasonable timeframe for those places to retain the footage they capture of you?
 - In what manner would you like to receive notification about those places' use of this data practice?

7 Appendix

7.1 Evening Review

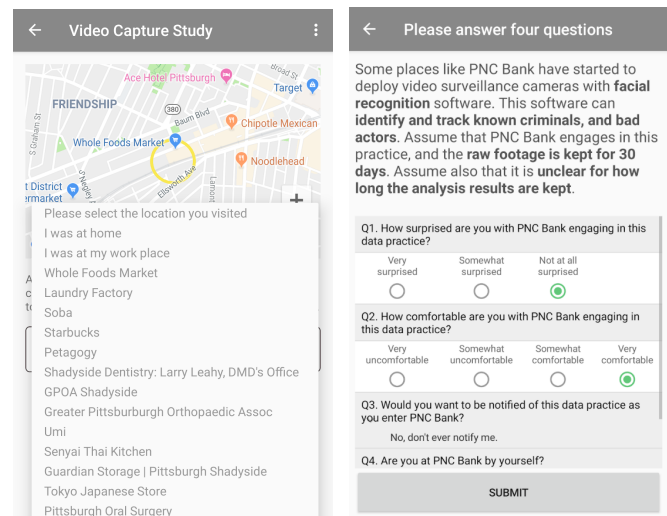
[Show a map, timestamp and scenario for each notification]

7.3 Interview Scripts

- Introduction: Thank you for agreeing to this interview. My name is _____. I will be audio-recording our session. How are you doing today? Just to fresh your memory. You started the study around [DATE], and finished the study around [DATE]. For this interview, we will be asking you some additional questions and clarifications about your experience during this study.
- Where did you find about our study?
- When did you download the app?
- Did you find participating in this study to be demanding?
- On average, how much time would you say you spent answering our questions each day?
- Were there days when you didn't receive any prompts?
- On the whole, do you feel that we covered most of the interesting places you went to during the course of the study, or would you say we missed some interesting places? If so, which interesting places did we miss? Would you expect cameras to be present at these places and what do you think these cameras could be doing?
- While going through the evening reviews, did you ever feel that you wished you could modify some of the answers you provided during the day? If so, can you specifically remember some of the scenarios and in which way you would have liked to modify your answers (e.g., less surprised or more surprised, less comfortable or more comfortable?)
- [CHECK DATA] For scenarios where we only collected your answers in the evening, because you didn't have time to answer them when the scenario occurred. Do you believe that you might have given different answers if you had responded at the time we first prompted you? If so, how different would your answer have been and why?
- [SHOW INSTANCES] Do you remember when you did not answer those scenarios on site / in-situ, why you could not answer them, and what you were you doing at the time?
- If you remember, each scenario came with two questions designed to check whether you had carefully read the description of the scenario. Do you remember those?
- Did you find that answering these questions could easily be defeated, or did you actually have to carefully read the scenarios to answer the questions? Feel free to tell us that the questions were easy to guess without reading the scenarios. We are trying to understand to what extent these questions help, or to what extent they are just not terribly useful.
- How often did you think that the scenarios we described matched actual video collection practices at the places you were visiting?
- Did you actually look for cameras, or start paying more attention to cameras?
- Have you discussed the study or scenarios with others?
- On the whole, do you feel that you have grown more concerned or less concerned about the types of video analytics scenarios used in our study? Or would you have you remain equally concerned or unconcerned?

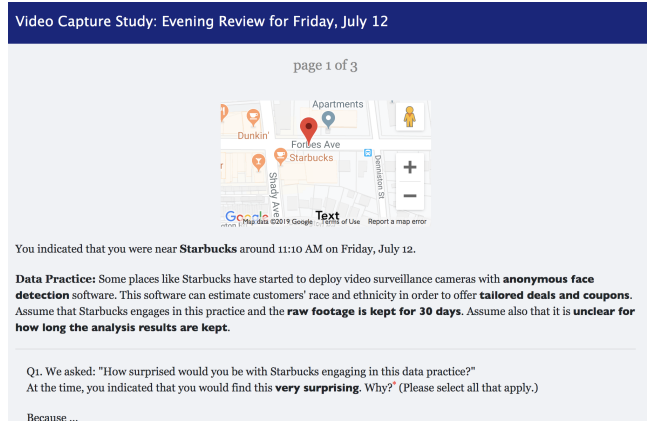
- If you remembered, there are a lot of scenarios you encountered as part of this study, were there scenarios that you found particularly surprising? Were there some scenarios that you found particular concerning? Or would you say that all these scenarios are to be expected and do not feel particularly concerning?
- Do you feel that, if you were to retake the study and be presented with the same scenarios, most of your answers would be the same? If some of your answers are likely to be different, could you identify some of the scenarios for which you would likely have different answers?
- Questions/Clarifications related to the interviewee's post-study and evening answers (different case by case)

7.4 Screenshots



(a) Users clarify their location

(b) In-situ scenarios



(c) Partial screenshot of the evening survey associated with a given scenario encountered earlier during the day

Figure 1: Screenshots of study instruments

7.5 Scenario Texts

Purpose	Scenario Text
Generic Surveillance	Some places like %s have started to deploy video surveillance cameras to deter crime.
Petty Crime	Some places like %s have started to deploy video surveillance cameras to deter crime. These cameras are equipped with software that can automatically detect and record petty crime (e.g. pickpocketing, car break-ins, breaking store windows).
Known Criminal Detection	Some places like %s have started to deploy video surveillance cameras with facial recognition software. This software can identify and track known shoplifters, criminals, and bad actors.
Count people	Some places like %s have started to deploy video surveillance cameras with anonymous face detection software. This software can estimate the number of customers in the facility in order to optimize operation, such as personnel allocation.
Jump Line	Some places like %s have started to deploy video surveillance cameras with facial recognition software. This software can identify patrons in line and push individualized offers to skip the wait-line for a fee.
Targeted Ads(Anon)	Some places like %s have started to deploy video surveillance cameras with anonymous face detection software. This software can estimate customers' race and ethnicity in order to offer tailored deals and coupons.
Targeted Ads(IDed)	Some places like %s have started to deploy video surveillance cameras with facial recognition software. This software can match detected faces against individual customer profiles in order to offer tailored deals and coupons.
Sentiment Ads(Anon)	Some places like %s have started to deploy video surveillance cameras with anonymous face detection and emotion analysis software. This software can estimate customers' age, gender and ethnicity, and analyze their reactions to items displayed. This software is used to generate tailored deals and coupons for different demographic groups.
Sentiment Ads(IDed)	Some places like %s have started to deploy video surveillance cameras with facial recognition and emotion analysis software. This software can recognize people, and analyze their reactions to items displayed. Then the software matches detected faces against individual customer profiles to send tailored deals and coupons to their phones.
Rate Service	Some places like %s have started to deploy video surveillance cameras with anonymous emotion analysis software. This software can gauge customer satisfaction with the service provided by its employees. They can use the results for employee evaluation and training purposes.
Rate Engagement	Some places like %s have started to deploy video surveillance cameras with facial recognition and emotion analysis software. This software can identify each patron, and measure their engagement at the facility.
Face as ID	Some places have started to deploy video surveillance cameras with facial recognition software. This software can identify faces to replace ID cards.
Track Attendance	Some companies have started to deploy video surveillance cameras with facial recognition software. This software can track the work time attendance of its employees.
Word Productivity	Some companies have started to deploy video surveillance cameras with emotion analysis and facial recognition software. This software can detect the mood of its employees, and predict their productivity. This software can record your presence and who you hang out with.
Health Predictions	Some eatery chains like %s have started to deploy video surveillance cameras with emotion analysis and facial recognition software. This software can detect your mood, and record data about your orders. This information can be shared with health insurance providers. The health insurance providers could use such data to estimate your likelihood of developing depression, diabetes, and obesity, which in turn can impact your health insurance premium.
Medical Predictions	Some medical facilities have started to deploy video surveillance cameras with emotion analysis and facial recognition software. This software can automatically detect some physical and mental health problems. This information can be shared with health insurance providers, which could impact your health insurance premium.

Table 4: Scenarios text shown to participants