

Shikun Zhang*, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman Sadeh*

“Did you know this camera tracks your mood?”: Understanding Privacy Expectations and Preferences in the Age of Video Analytics

Abstract: Cameras are everywhere, and are increasingly coupled with video analytics software that can identify our face, track our mood, recognize what we are doing, and more. We present the results of a 10-day in-situ study designed to understand how people feel about these capabilities, looking both at the extent to which they expect to encounter them as part of their everyday activities and at how comfortable they are with the presence of such technologies across a range of realistic scenarios. Results indicate that while some widespread deployments are expected by many (e.g., surveillance in public spaces), others are not, with some making people feel particularly uncomfortable. Our results further show that individuals’ privacy preferences and expectations are complicated and vary with a number of factors such as the purpose for which footage is captured and analyzed, the particular venue where it is captured, and whom it is shared with. Finally, we discuss the implications of people’s rich and diverse preferences on opt-in or opt-out rights for the collection and use (including sharing) of data associated with these video analytics scenarios as mandated by regulations. Because of the user burden associated with the large number of privacy decisions people could be faced with, we discuss how new types of privacy assistants could possibly be configured to help people manage these decisions.

Keywords: facial recognition, video analytics, privacy, experience sampling

DOI Editor to enter DOI
 Received ..; revised ..; accepted ...

***Corresponding Author: Shikun Zhang:** Carnegie Mellon University, E-mail: shikunz@cs.cmu.edu
Yuanyuan Feng: Carnegie Mellon University, E-mail: yuanyua2@cs.cmu.edu
Lujo Bauer: Carnegie Mellon University, E-mail: lbauer@cmu.edu
Lorrie Faith Cranor: Carnegie Mellon University, E-mail: lorrie@cmu.edu
Anupam Das: North Carolina State University, E-mail: anupam.das@ncsu.edu

1 Introduction

In August of 2019, a high school in Sweden was fined for unnecessarily relying on facial recognition to track students’ attendance, despite obtaining consent [3]. Over the past few years, the growing deployment of video analytics has prompted increased scrutiny from both privacy advocates and regulators [25, 29]. Yet, little is known about how people actually feel about the many different contexts where this technology is being deployed. While video analytics technologies such as facial recognition have become increasingly accurate thanks to recent advances in deep learning and computer vision [39], some deployments have also been shown to suffer from race and gender bias [40, 74]. The increasing ubiquity of video analytics is contributing to the collection and inference of vast amounts of personal information, including people’s whereabouts, their activities, whom they are with, and information about their mood, health, and behavior. As the accuracy of algorithms improves and as data continue to be collected across an ever wider range of contexts, inferences can be expected to reveal even more sensitive information about individuals. Unfortunately, such data collection and usage often take place without people’s awareness or consent. While video analytics technologies arguably have many potentially beneficial uses (e.g., law enforcement, authentication, mental health, advanced user interfaces), their broad deployment raises important privacy questions [90].

In the US, the GAO and NIST have recommended more transparency when it comes to appropriate use of facial recognition [4, 93]. New regulations such as the European Union’s GDPR and the California Consumer Privacy Act (CCPA) mandate specific disclosure and choice requirements that apply to the deployment of

***Corresponding Author: Norman Sadeh:** Carnegie Mellon University, E-mail: sadeh@cs.cmu.edu

video analytics technologies. While these regulations are important steps towards providing data subjects with more transparency and control over their data, they do not specify how people should be notified about the presence of video analytics, or how to effectively empower them to exercise their opt-in or opt-out rights. This includes addressing questions such as when to notify users, what to notify them about, how often to notify them, how to effectively capture their choices, and more. Our research aims to address these issues by developing a more comprehensive understanding of how people feel about video analytics deployments in different contexts, looking both at the extent to which they expect to encounter them at venues they visit as part of their everyday activities and at how comfortable they are with the presence of such technologies across a range of realistic scenarios.

The main contributions of this work are as follows:

- We offer an in-depth analysis of the data collected as part of 10-day in-situ study involving 123 participants who provided us with detailed insight into their degree of awareness and comfort across a total of 2,328 video analytics deployment scenarios.
- Our analysis reveals that many people have little awareness of many of the contexts where video analytics can be deployed and also show diverse levels of comfort with different types of deployment scenarios. Notification preferences are also shown to be diverse and complex, and seem to evolve over time, as people become more sophisticated in their expectations as well as in their realization of the number of notifications they may receive if they are not selective in their notification preferences.
- Finally, we review the implications of people’s rich and diverse privacy preferences when it comes to notifying them about different video analytics scenarios and to supporting opt-in or opt-out choices associated with the collection and use of their data under these scenarios. We focus in particular on the challenges resulting from the increasing deployment of these technologies and the corresponding burden on users. This includes a discussion of different possible configurations of privacy assistant functionality to selectively notify people about those scenarios they care to be notified about and to support different levels of delegation in managing opt-in/opt-out decisions.

2 Related Work

2.1 Privacy Challenges of Video Analytics

Video analytics, often equipped with facial recognition, is increasingly being integrated with the Internet of Things (IoT) systems [46, 47, 62]. Data privacy has been a central discussion in IoT [71] because IoT systems rely on the collection and use of contextual information (e.g., people, time, location, activity) in environments that often contains identifiable personal data [18, 72, 73]. Researchers have explored technical solutions to safeguard user data in IoT [30, 31, 83], including algorithms to avoid being tracked by video analytics [86, 87, 99]. However, transparency around IoT data privacy remains an unsolved issue [18, 78]. People often have no way to know the existence of video analytics deployments in their daily environments, what personal data is being collected, what purpose the footage is used for, and how long the footage will be retained. Moreover, video analytics has unique data privacy challenges. First, it can collect people’s biometric data (e.g., facial features, body pose) [75], which is considered more sensitive than digital identifiers like email addresses. Second, it can be applied later to video footage already collected by existing cameras for a myriad of purposes (e.g., security, operation optimization, targeted advertising).

These challenges indicate that the privacy implications of video analytics differ greatly in real-world scenarios, and should be evaluated case by case. Nissenbaum’s privacy as contextual integrity framework [67] is a theory best suited to evaluate the appropriateness of data practices of new technologies by considering important contextual factors. Under the framework, data practices can be evaluated against certain privacy norms in five information flow parameters — the sender, the recipient, the attribute, the subject, and the acceptable transmission principle. Changes to these parameters are likely to cause a privacy norm violation and must be examined closely [68]. However, privacy norms can vary across societies/cultures and may change over time, so existing privacy norms may not be suitable for new technologies like facial recognition in video analytics. Therefore, the first step to address data privacy challenges of video analytics is to establish a baseline of privacy norms by understanding people’s opinions and attitudes towards the technology.

2.2 Sampling and Modeling Privacy Preferences

Researchers have made initial progress in discovering privacy norms with IoT technologies in general by sampling people’s privacy expectations and preferences through vignette scenarios using large-scale online surveys [10, 65]. However, vignette studies are limited because participants have to imagine themselves in hypothetical scenarios that are not immediately relevant [6]. The experience sampling method (ESM), where both the context and content of individuals’ daily life are collected as research data, better examine links between external context and the contents of the mind [42]. Particularly, mobile-based ESM can prompt participants with the actual context they are in, enabling the collection of higher quality, more valid responses [13, 26]. This motivates us to use ESM to elicit people’s privacy expectations and preferences towards video analytics. As part of this study, we notify participants about realistic scenarios of video analytics deployment that could happen at the places they actually visit. Then, we ask about their privacy preferences towards these scenarios in situ, aiming to collect high quality responses to elucidate privacy norms regarding video analytics.

This study is also related to previous research on privacy preference modeling. Prior work has shown that individual privacy preferences vary greatly from one person to another and across different data collection and use scenarios [51, 57, 88]. One-size-fits-all models are often unable to capture individuals’ diverse privacy preferences when it comes to the collection and use of their data by mobile and IoT technologies. Research on mobile app permission preferences has shown that it is often possible to identify common patterns among the privacy preferences of different subgroups of users [56, 63]. Similar results have been reported in the context of IoT scenarios [51, 52, 65]. Some of this work has also demonstrated the use of machine learning models to predict individuals’ privacy preferences [59, 97] and help them manage their privacy decisions [58, 98].

2.3 Designing and Implementing Privacy Assistants

The past ten years have seen a proliferation of privacy settings, whether to enable users to block web trackers or to deny mobile apps access to their location. In practice however, users often struggle to configure privacy settings to match their privacy preferences, whether it

is because these settings are unintelligible [84], or because the number of available settings is unmanageable [7, 57, 59, 88], or both.

To overcome these usability challenges, recent research has advocated the introduction of “privacy assistants” to (1) notify people about sensitive data collection and use practices and motivate them to manage associated privacy settings [9], and to (2) also help them configure privacy settings [58, 79]. Privacy assistants can be enhanced by incorporating machine learning models of individuals’ privacy preferences to further reduce user burden [57–59, 91, 98]. For example, Liu et al. successfully demonstrated an Android privacy assistant app that relied on machine learning to generate personalized recommendations about which permission to grant or deny to different apps based on a small number of personalized questions answered by each user [58]. Users could review the recommendations and decide whether or not to accept them. The authors report on a pilot of this technology in the wild, with users indicating they saw value in the way in which this technology made it easier for them to manage a large number of privacy decisions without taking away control over their privacy decisions.

There is a growing body of research focusing on helping people manage their privacy in IoT contexts [27, 33]. This work ranges from the delivery of machine-readable privacy notices to users who are responsible for manually making all privacy decisions [44] to functionality that leverages models of individuals’ privacy preferences to help them manage their privacy. The latter includes the use of machine learning to generate privacy setting recommendations that users can review and accept (or reject) [58] as well as functionality that attempts to automate some privacy decisions on behalf of users [33]. Recent work generally indicates that people appreciate privacy assistant technology that helps them manage privacy decisions, while it also reveals that not everyone feels the same way about how much control they are willing to give up in return for a lighter user burden [22]. The work reported herein is intended to supplement this prior research by providing a more in-depth understanding of individuals’ privacy expectations and preferences in the context of a diverse set of video analytics scenarios. By understanding how rich and diverse people’s expectations and preferences actually are across these scenarios, we aim to build a better understanding of the complexity involved in notifying people about the presence of video analytics deployments and in enabling them to effectively manage associated privacy choices.

3 Study Design

3.1 Experience Sampling Method

Context has been shown to play an important role in influencing people’s privacy attitudes and decisions [68]. Studying people’s privacy attitudes through online surveys is often limited because participants answer questions about hypothetical scenarios and often lack context to provide meaningful answers. Accordingly, we conducted an experience sampling study to collect people’s responses to a variety of video analytics deployments (or “scenarios”) in the context of their regular everyday activities. The experience sampling method [42] has been repeatedly used in clinical trials [48, 95], psychological experiments [17, 43], and human-computer interaction (HCI) studies [36, 80], yielding “a more accurate representation of the participants’ natural behaviour” [94]. This enables us to engage and survey participants in a timely and ecologically valid manner as they go about their normal daily lives [70]. Participants are prompted to answer questions about plausible video analytics scenarios that could occur at the location in which they are actually situated.

3.2 Selecting Realistic Scenarios

Previous research mainly surveyed participants’ privacy attitudes in the context of generic IoT scenarios, including some facial recognition scenarios [52, 65]. By systematically exploring more concrete scenarios in actual settings associated with people’s day-to-day activities, we are able to elicit significantly richer reactions from participants and develop more nuanced models of their awareness, comfort level, and notification preferences pertaining to different deployment scenarios. The scenarios considered in our in-situ study were informed by an extensive survey of news articles about real-world deployments of video analytics in a variety of different contexts (e.g., surveillance [81], marketing [82], authentication [11], employee performance evaluation [28], and church attendance tracking [12]). These scenarios provided the basis for the identification of a set of relevant contextual attributes which were randomly manipulated and matched against the different types of venues our subjects visited.

Our baseline scenario described the use of generic surveillance cameras with no video analytics. All other scenarios in our study involved the use of some type

of video analytics. *Security-related* scenarios included automatic detection of petty crime [81], and identification of known shoplifters and criminals in public places [2, 24, 37, 45]. Scenarios for *commercial* purposes included helping businesses to optimize operations [64, 69, 82], displaying personalized advertisements based on the detection of demographic features [34, 37, 76, 92], collecting patrons’ facial reaction to merchandise [15, 16, 21, 85], and detecting users’ engagement at entertainment facilities [53, 60, 96]. Other significant use case scenarios revolve around *identification* and *authentication*. Here, we considered two broad categories of scenarios: (1) replacing ID cards with facial authentication in schools, gyms, libraries and places with loyalty programs [11, 32, 66, 89], and (2) attendance tracking in the workplace, at churches, and at gyms [11, 12, 38]. Lastly, we included a small number of plausible, yet hypothetical, scenarios inspired by emerging practices as discussed in news articles or as contemplated in research. This includes health insurance providers using facial recognition and emotion analysis to make health-related predictions [8, 55, 77]; employers using emotion analysis to evaluate employee performance [28, 35, 49, 54]; hospitals using emotion recognition to make health-related predictions [1, 35, 41].

In total, we identified 16 purposes, as shown in Table 1, representative of a diverse set of video analytics scenarios. A representative list of the scenarios as well as the corresponding text shown to participants to elicit their reactions can be found in the Appendix (Table 7). The scenario text was crafted through multiple iterations to sound plausible without deceiving participants.

3.3 Factorial Design

We employed a factorial study design and developed a taxonomy that captured a representative set of attributes one might expect to influence individuals’ privacy attitudes. These attributes are shown in Table 1. We specified a discrete set of possible values for each attribute, taking into account our desire to cover a broad spectrum of scenarios while also ensuring that we would be able to collect a sufficiently large number of data points for each scenario. Here, we differentiate between the retention time of raw footage and of video analytics results because raw video data, containing biometrics, can be very sensitive, and possibly be exploited for additional analyses subsequently.

3.4 Study Protocol and Procedures

The 10-day study comprised the following five stages.

Stage 1: Eligible participants completed the consent forms for this study and downloaded the study app from the Google Play Store. Upon installing the app, participants completed a pre-study survey about their perceived knowledge level, comfort level, and notification preference with regard to facial recognition.

Stage 2: Participants were instructed to go about their regular daily activities. The study app collected participants’ GPS locations via their smartphones. As they visited points of interest, namely places for which we had one or more plausible deployment 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, as illustrated in the app screenshots in Fig. 1(i)–(iv). The protocol limited the number of scenarios presented to each participant to six per day, though most of the time participants’ whereabouts would trigger a smaller number of scenarios—closer to three per day.

Stage 3: On the days participants received push notifications via the app, they 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 a few additional questions. See Fig. 1(v) for an example of the evening review.

Stage 4: After completing 10 days of evening reviews, participants concluded the study by filling out a post-study survey administered via Qualtrics. This survey contained free-response questions about their attitudes on facial recognition, the 10-item IUIPC scale on privacy concerns [61], as well as additional demographic questions like income, education level, and marital status.

Stage 5 (Optional): Participants who indicated they were willing to be interviewed in their post-study survey may be invited to an online semi-structured interview. The interview contained questions about study validity, perceptions of scenarios, and clarifications with regard to their earlier responses.

To maximize the contextual benefits provided by the experience sampling method [20], we designed a sophisticated payment scheme to incentivize prompt responses to in-situ notifications. Participants were compensated \$2 per day for each day of the study. They received an additional 25 cents per notification they responded to within 15 minutes, or 10 cents if they responded to the notification between 15 and 60 minutes. We also compensated them \$2 for the time spent on answering pre-study and post-study surveys. An additional \$15 was awarded when they finished the study. In total, participants could earn between \$37 and \$52 and were compensated with Amazon gift cards. Participants who completed the online interviews were awarded \$10.

Attribute Name	Values
Purpose	Generic Surveillance, Petty crime detection Known criminal detection (Anonymous) people counting (Individualized) jump the line offers (Anonymous) demographic ad targeting (Individualized) ad targeting (Anonymous) sentiment-based ad targeting (Individualized) sentiment-based ad targeting (Anonymous) sentiment-based customer service evaluation (Individualized) customer engagement detection Attendance tracking Using face as IDs Work productivity predictions Health predictions - eatery visits Health predictions - medical visits
Anonymity level	No video analytics Anonymous face detection Facial recognition
Retention of raw footage	ephemeral, 30 days, unspecified
Retention of analysis results	ephemeral, 30 days, unspecified
Sharing specified	Yes, No
Detection of who people are with	Yes, No
Type of places	store, eatery, workplace, education, hospital, service, alcohol, entertainment, fitness, gas, large public places, transportation, worship, library, mall, airport, finance

Table 1. Contextual attributes: Among all the possible combinations of these attributes, our study focused on a subset of 102 scenarios representative of common and emerging deployments of video analytics technology.

3.5 Ensuring Study Validity

Due to the complexity and the number of components of the study framework, we conducted several pilot rounds, with initial rounds involving members of our research team and later rounds involving a small number ($N=9$) of external participants. Each pilot round helped identify issues that needed to be addressed, whether in the form of small refinements of our protocol or adjustments to technical components of our system (e.g., study app, web survey app, study server). Below, we briefly discuss the two most important refinements that were made as a result of this process.

Because of the limitations of location tracking functionality, we determined that we could not automati-

cally pinpoint the location of our subjects and use that location to automatically identify a relevant video analytics scenario. Instead, we opted to use location tracking to automatically generate a drop-down list of venues near our subject. We then asked them to select the actual venue where they were. The drop-down list of venues always included three additional options: “I was somewhere else in the area,” “I was passing by,” and “I was not there.” This ensured that our protocols also accounted for missing venues, situations where our subjects were passing by a given location (e.g., being stuck in traffic), as well as situations where location tracking was potentially inaccurate. Participants still received payments for each scenario when they selected one of these three additional choices. In other words, they had no incentive to select a place that they did not visit.

During the first pilot, we found that some participants did not seem to pay close attention to some of the scenario attributes (Table 1). This was remedied by introducing two multiple-choice attention check questions (see Figure 1(ii)). These questions required participants to correctly identify two different and randomly selected contextual attributes assumed in the scenario (attributes in Table 1, excluding type of places). Participants were only allowed to proceed with the remaining in-situ questions once they had passed the two attention checks. These attention checks proved rather effective, as discussed in the Section 4.3.

3.6 Recruitment and Ethics

We recruited participants using four methods: posts on local online forums for the Pittsburgh area (e.g., Craigslist, Reddit), posts in a university-based research participant pool, promotional ads on Facebook, and physical flyers posted on local community bulletin boards and at bus stops. Potential participants were asked to take a short screening survey to determine eligibility (age 18 or older, able to speak English, using an Android smartphone with data plan). The screening survey also displayed the consent form for the study and collected basic demographic information such as age, gender, and occupation. Recruitment materials, the consent form, and the screening survey did not mention or refer to privacy. We tried to avoid convenience samples of undergraduate college students, and purposely looked for participants with a variety of occupations.

This research was approved by our university’s institutional review board (IRB) as well as the funding agency’s human research protection office. As location

data collected over a period of time can be particularly sensitive, we refrained from using off-the-shelf experience sampling software and developed our own system and an location-aware Android app.

4 Results

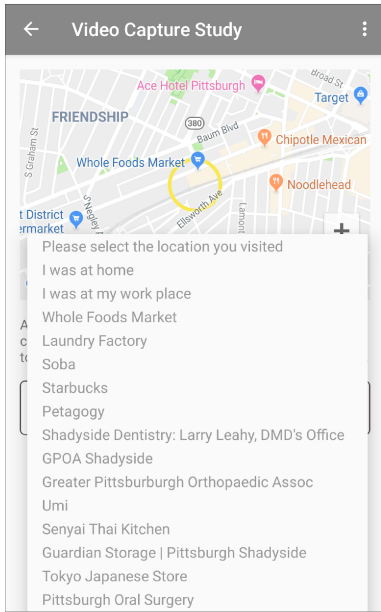
4.1 Participants and Responses

A total of 164 individuals (excluding 9 pilot participants) took part in the study and downloaded our study app from the Google Play Store between May and November 2019. Of these, 124 completed the 10-day study. One participant was removed due to poor response quality as that person selected “I was somewhere else” for all the notifications received. Among the remaining 123 participants, 10 (8%) were 18-24 years old, 67 (54.5%) were 25-34, 29 (23.6%) were 35-44, 10 (8%) were 45-54, 4 (3%) were 55-64, and 3 (2%) were between 65 and 74. In our sample, 58% identified as female, 41% as male, and 2% as other. Most participants were highly educated: 43 (35%) had bachelor’s degrees, and 46 (37%) had graduate degrees. Half of the participants were single and never married, and 42% were married or in a domestic partnership. The majority of our participants (82%) reported having no children under 18 living with them. Participants reported diverse occupations (see Table 5 in the Appendix). The average IUIPC factor scores of our participants are shown in Table 2. Comparing our results with those of a large MTurk sample from another study (N=1007) [65] using Mann-Whitney U tests, we found no difference in the collection and the awareness factors, and a significant difference in the control factor with a small effect size ($r = 0.1, p < 0.01$).

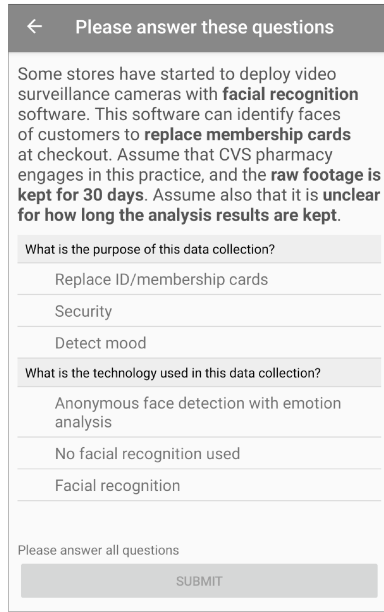
	Ours Mean [SD]	MTurk Mean [SD]	Reject H0
IUIPC-Collection	5.90 [1.04]	5.79 [1.11]	No
IUIPC-Control	6.21 [0.78]	5.95 [0.90]	Yes
IUIPC-Awareness	6.53 [0.66]	6.44 [0.82]	No

Table 2. Comparison of IUIPC scores of our participants (N=123) with an MTurk sample (N=1007). H0 stipulates that two samples come from the same population. Cannot reject H0 means that 2 groups are not significantly different.

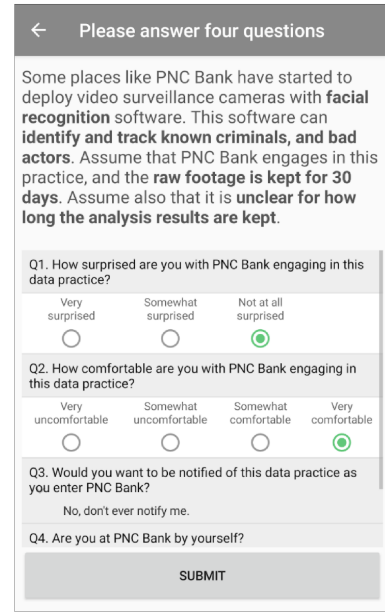
We recruited interviewees about halfway through the study. Participants were selected based on their de-



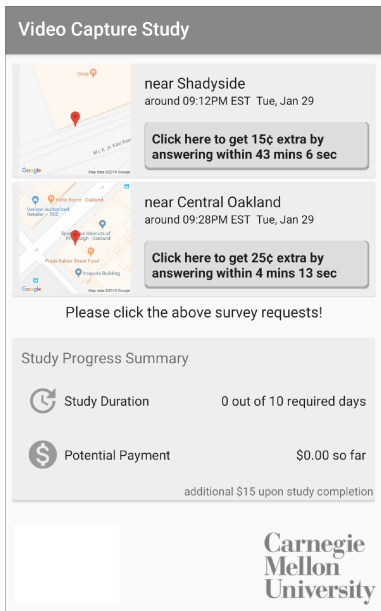
(i) Prompting users to clarify their location



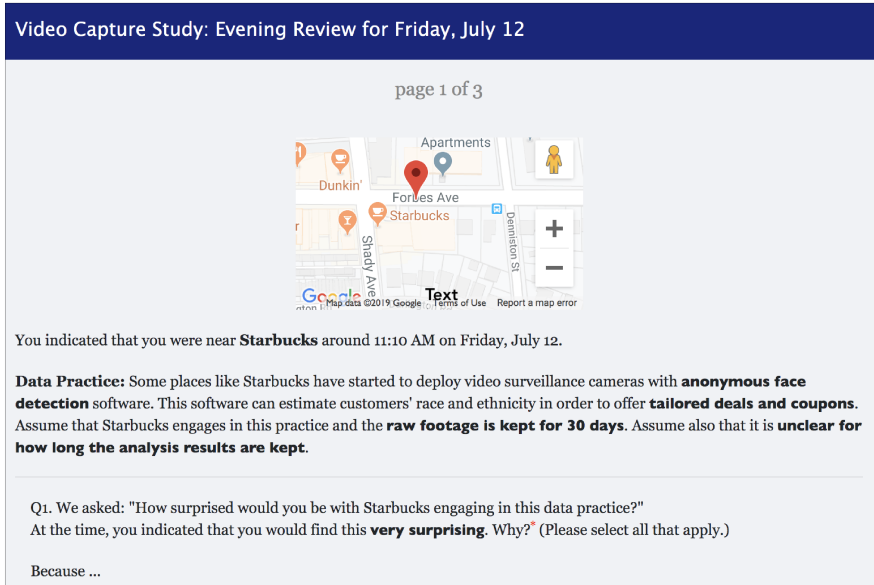
(ii) Two attention check questions designed to ensure participants read about relevant attributes



(iii) Four in-situ questions



(iv) Dashboard showing prompts to complete two in-situ surveys, including monetary incentives to respond as quickly as possible



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

Fig. 1. Screenshots of the study app and the web survey used for the evening review

mographics. We sent out 17 invitations and conducted online interviews with 10 participants who followed up. In total, participants were sent 3,589 notifications prompting them to identify their specific location (Fig. 1(i)). In the majority of cases (65%), our system was able to retrieve a scenario relevant to the location reported by the participant, such as the two scenarios shown in Fig. 1(ii) and (iii). For the remaining 35%, the system did not have a pre-identified scenario that

matched the response provided by the participant, in which case we were unable to elicit any additional information from the participant for that particular location. Based on answers provided by participants, common examples of such situations included the participant being at home or visiting a partner, friend, or relative. Other situations included the participant waiting for a bus or passing by a location. In some instances, participants reported that they did not see the location at which they

were in the drop-down menu shown to them (Fig. 1(i)). This seemed to most commonly occur when participants were in parks, parking lots, farmers’ markets, new establishments, or small local stores.

When the system was able to retrieve a plausible scenario relevant to the participant’s location, the participant was presented with the scenario and prompted to answer a few quick questions related to that scenario (e.g., see Fig. 1(ii) and (iii)). In addition to these in-situ responses, they were also requested to answer a more complete set of questions about the scenario in the evening. As a result, we were able to collect in-situ and evening responses for a total of 2,328 scenarios. Each participant on average provided in-situ and evening responses to 19 scenarios over a 10-day period, and received an average compensation of \$41.

4.2 Privacy Attitudes

When surveying participants’ responses to facial recognition scenarios, we focused on four related questions: how surprised they were by the scenario presented to them (**surprise level**), how comfortable they were with the collection and use of their data as assumed in that scenario (**comfort level**), to what extent they would want to be notified about the deployment scenario at the location they visited (**notification preference**), and whether, if given a choice they would have **allowed** or **denied** the data practices described in that scenario at that particular location at the time they visited that location (**allow/deny preference**). These questions are shown in Fig. 2.

How surprised are you with *Controller* engaging in this data practice?

Very surprised Somewhat surprised Not at all surprised

How comfortable are you with *Controller* engaging in this data practice?

Very uncomfortable Somewhat uncomfortable Somewhat comfortable Very comfortable

Would you want to be notified of this data practice as you enter *Controller*?

Yes, notify me every time it happens.
 Yes, but only once in a while to refresh my memory.
 Yes, but only the first time I enter this location.
 I don't care whether I am notified or not.
 No, don't ever notify me.

If you had the choice, would you allow or deny this data practice?

Allow Deny

Fig. 2. *Controller* being a variable that would be instantiated with the name of the venue participants were visiting

Fig. 3 provides a summary of collected responses organized around the 16 categories of scenarios (or “purposes”) introduced in Table 1. As can be seen, people’s responses vary for each scenario. In other words, “one size fits all” would fail to capture individuals’ diverse preferences when presented with these scenarios. At the same time, some scenarios elicit more consistent responses from participants than others. For instance, generic surveillance scenarios appear to surprise participants the least and to elicit acceptance by the most (close to 70% would agree to such scenarios, if given a choice and fewer than 10% reported feeling “very uncomfortable” with such scenarios). Yet, even in the presence of such scenarios, 60% of participants reported they would want to be notified at least the first time they encounter these scenarios at a given venue and over 35% indicated they would want to be notified each time. At the other end of the spectrum, scenarios involving facial recognition for the purpose of evaluating employee productivity or tracking attendance at venues elicited the greatest level of surprise and lowest level of comfort among our participants, with barely 20% reporting that, if given a chance, they would consent to the use of these technologies for the purpose of evaluating employee productivity. Similarly, participants expressed significant levels of surprise and discomfort with scenarios involving the use of facial recognition to make health and medical predictions or to track the attendance of individuals.

4.3 Study Validity and Benefits of ESM

Below we report results on study validity, focusing on three aspects: whether participants carefully read the scenarios, whether they thought the scenarios could happen, and how the ESM helped anchor their responses to their everyday life experience.

Overall, 81% of the time participants successfully completed both attention check questions associated with the scenarios assigned to them within two attempts. Attention questions were found to be useful by 8 out of the 10 interviewees. For instance, one participant (P107) stated, “*I think you definitely had to read them [scenarios]. I think there was one or two that I saw the bold words, and thought that they were the same as older questions, so I picked the same answer, and it was a different one. So once I re-read it, I saw that it was a little different.*” Five interviewees reported attention questions helping them discern between retention for raw footage, and retention for analysis results, as




Fig. 3. Summary of collected responses organized around 16 different purposes. The bottom row shows the aggregated preferences across different purposes.

P55 said, “But the first couple of times, I mixed up the raw footage with the analysis results, but after that [the attention checks] I remembered to look for the distinction.” These comments suggest that the attention checks contributed to participants noticing the contextual attributes associated with each scenario and that the responses we collected most likely reflect privacy attitudes that take these contextual attributes into account.

As 68% of in-situ questions were answered within 15 minutes and 87% within 1 hour, the actual location visited by the participant and the context associated with the scenario were likely still fresh in their mind (e.g., what the participant was doing at that particular location, or who they might have been with). When asked about whether the scenarios matched actual video collection practices at the places participants were visiting in the exit interviews, most ($N = 7$) stated that they found the scenarios to be realistic, and “it is entirely possible that it is happening in those places” (P55). P107 explained, “I don’t know if they actually use any of the strategies right now, but they did seem to fit pretty well with the places like grocery stores offering coupons, or targeting some ads towards you.”

Furthermore, the experience sampling method provided context to participants’ responses, with participants reporting that context played an important role in influencing their attitudes towards different video analytics deployments. When the participants selected in situ that they felt somewhat or very uncomfortable about a scenario, in daily the evening reviews they can select multiple-choice options and provide additional

free responses to further explain their discomfort. Fig. 4 plots the reasons participants selected, many of which are directly related to the in-situ context. The figure also shows the percentages of participants who ever reported considering each reason: many reasons were taken into account by the majority of 123 study participants. Our




Fig. 4. Percent of participants/notifications reporting specific reasons for discomfort. Participants only selected reasons for notifications that they indicated discomfort ($N=1,369$). N is the used as the denominator to calculate the percent of notifications.

qualitative analyses of free responses in evening reviews also revealed that study participants had context in mind when they explained their in-situ comfort level. Their responses also reflected various aspects of data flows as by Nissenbaum’s framework of CI [67]. Example quotes listed by purpose are shown in Table 3.

Purpose	Example Quotes	Values
Generic Surveillance(No)	I'm fine with it to keep banks more safe. — P27	A,TP
Petty Crime(Anon)	When it comes to law enforcement and public safety I am more ok with giving up privacy. But there is an expectation that the data is protected. But any data collected for one reason is expected to stay within that original use. — P12	R,TP
Detect Criminal(IDed)	Because this is a bar, I feel like I would be more willing to acquiesce to a certain degree of surveillance for my own safety. — P59	A,TP
Count People(Anon)	It's anonymous and seems like a good use of the technology. — P68	TP
Jump Line(IDed)	The cafe is never super crowded when I go, and the space is small. I am surprised they would need something like that due to area and logistics. — P16	R
Targeted Ads(Anon)	I've heard that Target has the most advanced security, so it's kind of unsettling because I don't know exactly what they're doing. — P7	R,TP
Targeted Ads(IDed)	It's the facial recognition of it and keeping of derived data that bothers me. — P13	A,TP
Sentiment Ads(Anon)	It's anonymous so I don't care as much. Also I have pretty good brand loyalty to Target and trust them more than I probably should. — P40	TP,R
Sentiment Ads(IDed)	The errands I do there are acceptable for all audiences. — P9	A
Rate Service(Anon)	I would expect this practice from larger chains rather than a small, local store, so it weirded me out a little to think the surveillance technology was there. — P27	R
Rate Engagement(IDed)	It might help improve the experience. — P110	TP
Face as ID(IDed)	I trust this location with footage as it is my local gym, and it actually would be convenient in this case. — P106	R,TP
Track Attendance(IDed)	It's a military base with 100% ID check at the gate, so I know about it and basically trust them. — P25	R
Work Productivity(IDed)	Big Brother is watching. I did not consent. — P104	TP
Health Predictions(IDed)	I don't like sharing data with health insurance companies. — P13	TP
Medical Predictions(IDed)	Emotion analysis combined with facial recognition makes me more uneasy than other ways this tech is implemented, especially coming from a healthcare provider. — P58	TP,R

Table 3. Example quotes from participants' evening reviews explaining their in-situ answers. Their responses were coded by relevant parameter values of contextual integrity. A—Attribute: Any description of information type. R—Recipient: Any entity (person, company, etc.) that receives the information. TP—Transmission Principle: The conditions under which information may be used or collected [68].

4.4 Factors Impacting Privacy Attitudes

The responses collected as part of this in-situ study provide rich insight into people's awareness of the many different ways in which facial recognition is deployed, how comfortable they are with these deployments, and to what extent they would want to be notified about them. Our analysis is organized around the different contextual factors already identified in Table 1. On average each participant responded to a total of about 19 deployment scenarios. These 19 different scenarios covered an average of 9.9 different “purposes,” as defined in Table 1, and 5.9 different types of venues, thereby offering rich insight into how people feel about facial recognition deployments across a range of different situations.

4.4.1 Allow/Deny Decisions

We first investigate whether people's decisions to allow or deny data collection have a relationship with the con-

textual attributes in Table 1. We constructed our model using generalized linear mixed model (GLMM) regression [14], which is particularly useful for data analysis with repeated measures from each participant. Our GLMM model was fit by maximum likelihood (Laplace approximation) treating the user identifier as a random effect, using a **logistic link** function for the binary response (allow/deny).

Among all the attributes introduced in Table 1, we find that “purpose” exhibits the strongest correlation with the decision to allow or deny data practices associated with our scenarios. In particular, when compared against “generic surveillance” scenarios, 12 out of 15 other purposes came out as being significantly more likely to result in a “deny” decision. Participants were respectively 23.5 ($=e^{3.16}$) times and 29 ($=e^{3.37}$) times more likely to respond with a “deny” to deployment scenarios for predicting work productivity, and for predicting health, compared to generic surveillance scenarios with no facial recognition. The odds of participants denying purposes for targeted advertising were at least 6

($=e^{1.87}$) times and up to 16 ($=e^{3.16}$) times greater than the odds for generic surveillance. Even for the purpose of using faces for authentication and identification, participants were still more likely to deny data collection (odds ratio = $e^{1.70} = 5.5$). Three purposes turned out not to be significant: detecting petty crime, using anonymous facial detection to count the number of people in the facility, and using facial emotion detection to rate engagement. The last of the three purposes, despite being relatively intrusive in comparison with the previous two, did not seem to have an important impact. We suspect that this might be partially due to the low number of occurrences ($N = 23$) of this purpose as this scenario was only associated with visits to places like movie theaters, museums, and amusement parks.

Contrary to our expectations, we found that whether targeted ads relied on identifying individuals or whether they treated them anonymously did not elicit substantially different responses from our participants. In fact, participants reported being more likely to respond with a “deny” to facial recognition scenarios used in targeted ads based on demographic features like race or ethnicity than to scenarios which involved individually targeted ads. The interview data revealed that some participants (3 out of 10) were viewing advertising based on demographics (e.g., race and age) as a form of profiling. For example, P106 stated, *“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.”*

Some of the place type attributes were also found to have an influence on participants’ allow or deny decisions. When we compare different place types to the baseline of large public places (e.g., sports stadiums, parking garages, city hall buildings), we find that participants were more likely to deny data practices at eateries (odds ratio = $e^{1.09} = 3$), at libraries (odds ratio = $e^{1.71} = 5.5$), and at gas stations (odds ratio = $e^{1.36} = 3.9$). Participants were significantly less likely to respond with a “deny” to deployment scenarios at transportation locations (buses stops, train stations, metro stations) than at the more generic baseline (odds ratio = $e^{-1.87} = 0.23$). The number of days participants had been in the study also seemed to influence their allow/deny decisions. Participants proved more likely to respond with a “deny” as the study progressed. None of the other attributes were statistically significant ($p < 0.05$). We present the complete results from the regression in the Appendix (Table 6).

4.4.2 Comfort Level, Surprise Level, and Notification Preference

Here we explore how the different contextual attributes considered in our study seem to influence participants’ comfort level, surprise level, and notification preferences. As those responses are not binary or linear, GLMM is not suitable due to its inability to model ordinal dependent variables. Instead, we opted for cumulative link mixed models (CLMM) fitted with the adaptive Gauss-Hermite quadrature approximation with 10 quadrature points using the R package `ordinal` [19]. We constructed one CLMM model for each dependent variable, adopting the same set of independent variables and random effect, as is the case with allow/deny decisions described in Section 4.4.1.

Similarly to the case with allow/deny decisions, purpose remains the attribute with the strongest influence on participants’ comfort level, surprise level, and notification preferences. Participants are more likely to feel uncomfortable, surprised, and are more likely to want to be notified when confronted with scenarios involving facial recognition than with our baseline “generic surveillance” scenario with no facial recognition. Data sharing with other entities seems to also contribute to a significant reduction in comfort among participants. As is the case with allow/deny decisions, we also found that the number of days in the study was significantly correlated with participants’ surprise level and notification preferences. Participants reported being less surprised over time, likely because they had already encountered similar scenarios earlier in the study. Over time, participants became slightly more inclined to deny scenarios, while their notification preferences became somewhat more selective. These results are further explored in Section 4.5.1 and 4.5.2.

4.5 Attitude Change Between Start and End of the Study

In our pre-study and post-study surveys, we asked participants the same questions about their understanding of, comfort level with, and notification preferences for facial recognition. In the post-study, we also asked them to provide open-ended responses to why their level of concern may have (not) changed. We analyzed these responses using inductive coding. Two researchers iteratively improved the codebook and independently coded all responses. Coding discrepancies were discussed and

reconciled. We reported results from comparing both surveys and qualitative coding.

4.5.1 Increased Awareness

By the end of the study, 60% of participants ($N = 74$) reported increased awareness resulting from participation in the study. They did not realize facial recognition could be used for so many different purposes, at such a diverse set of venues, and with this level of sophistication. For instance, P68 wrote, “Some of the scenarios and growth of the technology you mentioned, I had never considered. Freaked me out.” 11% of the above group reported learning the benefits of facial recognition. P106 explained, “In the beginning I was very uncomfortable with the fact that this tech could be abused or that law enforcement could use it. However, as the scenarios came up in the study, I realized it could be helpful in my life as long as there are safeguards in place to prevent abuse.” At the end of the study, when rating how much they thought they knew about facial recognition, one third of participants rated their knowledge of facial recognition lower than what they had reported at the start. This situation could be explained by the Dunning-Kruger effect, a cognitive bias wherein people tend to overestimate their knowledge in areas which they have little or no experience [50]. As participants grew more aware of possible video analytics deployments, they gained a more grounded estimate of their knowledge level. In interviews, 5 out of 10 interviewees indicated their awareness had increased. For instance, P50 mentioned “I didn’t know when I started there were so many different potential uses. I only thought that it could be used for tracking someone who committed a crime, so I was really surprised that there are so many different things being developed. And I definitely do think there are good uses and some that are more invasive.” Three interviewees described their deliberation on facial recognition usages as the study progressed. For example, P56 recounted “I feel like I might’ve started to get more negative about

the use of cameras... I could easily how see all of this information would go to very bad places... In some ways now that I am more aware of it, I’ve certainly put more thought into it and became more negative about it.”, and P107 gave an account of his thought process: “I think it’s just thinking about it more, being asked a couple of different times, and then you get asked once you just kind of answer it, but then twice and the third, I really think about it. It’s been in my mind already, so then the answer is probably more close to what I think... by the end, maybe I am not so sure about them having that information. But I think by the last 3 or 4 days, they were more consistent, consistently no for certain ones.” This could possibly explain why the number of days in the study was a significant predictor of participants’ allow and deny preferences and why they tended to deny more as the study progressed as reported at the end of Section 4.4.1.

4.5.2 Evolution of Notification Preferences

Before the study, 95.9% of all participants claimed that they wanted to be notified about facial recognition deployment scenarios, including 51.2% who indicated they wanted to be notified every time they came within range of facial recognition. As shown in Fig. 5, between the beginning and end of the study 55.3% of participants changed their preferences regarding whether and how often they wanted to be notified about facial recognition deployments. Among participants who originally wanted to be notified every time, 44% of them opted for less frequent notifications. This is also supported by the positive coefficient associated with the number of days predictor of the CLMM regression model for notification preferences, as stated in Section 4.4.2, as well as the descending line in Fig. 6, which plots the percentage of notifications where participants want to be notified every time or once in a while against the number of days in the study.




Fig. 5. A Sankey diagram shows the change of participants’ reported notification preferences before and after the study




Fig. 6. Participants’ desire to be notified decreases as the study progresses

One possible explanation is that people gradually developed a better appreciation for the broad deployment of these scenarios, and the possibility of receiving a large number of notifications, as P53 described, “*I think at first when I first started, I was saying once in a while and then I realized that would be really annoying to get multiple notifications.*” Some participants also expressed resignation. For instance, P89 said, “*The whole concept has become normal to me. I’ve definitely been reminded, through the app, that cameras with facial recognition are used in many, many places. I’ve become desensitized to the practice, and in fact, what I had considered in some ways[sic] to be negative because I want my privacy.*” It

is also worth noting that, as can be seen in Fig. 5, a simple “Ask on First Use” approach would not accommodate most users. If anything, changes identified in participants’ responses before and after the study indicate that people seem to become more sophisticated over time in their notification preferences with a substantially smaller fraction of participants requesting to be notified every time by the end of the study. The majority are looking for some type of selective notification solution.

On the other hand, we also noticed that a sizable minority of participants (shown in bottom of Fig. 7) stayed relatively consistent throughout the study with regards to their notification preferences, as they wanted to be notified every time facial recognition is in use. Results from interviews revealed that some participants would always want to be notified, like P56 noted “*The fact that I wanted everything to be always reminding me... I think it is worth letting people know upfront, and every time so you don’t get used to it and complacent.*” P52 also explained why he would always want to be notified at his workplace: “*At my work, if I didn’t think it was necessary or appropriate, then it wouldn’t register in my head that I was being watched. I would have to be reminded every time.*”

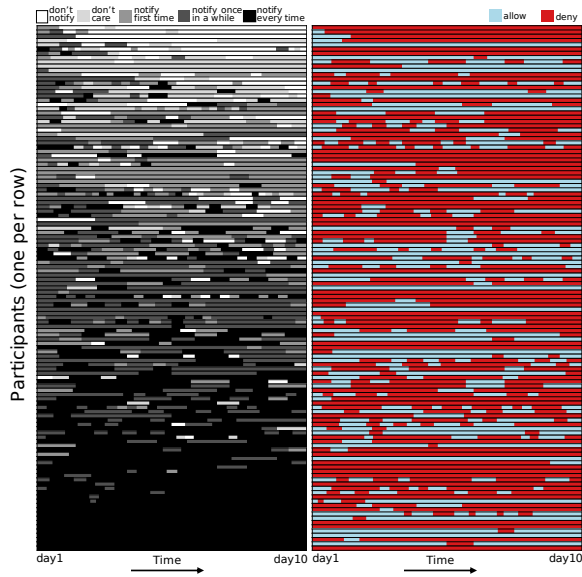


Fig. 7. This graph shows all participants’ notification and allow/deny preferences throughout the study in time order. Each row represents one participant. Participants are ordered increasingly by their desire to be notified on the left graph, and their corresponding allow/deny preferences were shown on the right. On the left graph, participants’ desire to get notified less over time is illustrated by the gradually brightening color from left to right, especially in the top right corner.

4.6 Correlation Between Privacy Expectations and Allow/Deny Preferences

Prior research has shown that comfort is often correlated with the degree of surprise people express towards different data collection and use practices [56]. We compiled pairwise correlations between the four types of responses collected from our participants across the 2,328 scenarios evaluated in our study (Table 4). Correlations were calculated using the Spearman rank correlation with Bonferroni-corrected p-values. Not too surprisingly, we find a significant correlation with a large effect size between people’s comfort level and whether they would allow or deny a given scenario. As reported in prior research [56], we also find a moderate correlation between surprise about some deployment scenarios and comfort with these scenarios. On the other hand, correlation between allow/deny decisions and desire to be notified seems nearly non-existent, suggesting people’s notification preferences do not simply correspond to their allow/deny preferences across different scenarios. An example of this case was mentioned in the previous section: only 30% of participants would deny data

practices for generic surveillance purposes, but 60% reported that they would like to be notified. Our qualitative results in Section 4.5.2 and Fig. 7 also seemed to suggest that individuals’ notification preferences are rather distinct from their allow/deny preferences, and serve different needs.

	comfort	surprise	notification	allow/deny
comfort	1			
surprise	0.442***	1		
notification	0.183***	0.214***	1	
allow/deny	0.604***	0.350***	0.046	1

Table 4. Correlation matrix where *** indicates $p < 0.001$

5 Discussion

5.1 Limitations

We do not claim that our results are representative of the general population. Our sample population skews young and more educated, which could have induced bias in our results. In addition, participants were recruited only from Pittsburgh, Pennsylvania, a mid-sized city in the United States.

Our study protocol determined the type and frequencies of scenarios participants saw, which in turn likely impacted their attitudes over time and in particular their notification preferences. We strived to keep the study realistic by presenting each participant with scenarios representative of the venues they visit in their everyday life. The actual frequency and types of video analytics participants would encounter could, however, be different from those in our study, and are likely to evolve over time. Our analyses were conducted using data provided by participants when presented with plausible deployment scenarios, rather than based on observations in the presence of actual deployments. While our use of an in-situ methodology was intended to mitigate this issue, it is possible that some of the data collected is not fully representative of participants’ actual behaviors.

While describing study scenarios, we strived to maintain a balanced narrative without overly emphasizing benefits or potential risks associated with different deployments, but rather leaving it to participants to decide how they felt about them. This being said, we acknowledge that the phrasing of these types of scenar-

ios is an art and that on occasions our phrasing might have primed participants in one direction or the other.

5.2 Lack of Awareness and Desire for Greater Transparency

Our results clearly indicate that many people were taken by surprise when encountering a variety of video analytics scenarios considered in our study. While many expect surveillance cameras to be widely deployed, few are aware of other types of deployments such as deployments for targeted advertising, attendance, productivity, and more. These less expected scenarios are also those that generally seem to generate the greatest discomfort among participants and those for which, if given a chance, they would often opt out (or not opt in). These results make a strong case for the adoption of more effective notification mechanisms than today’s typical “this area under camera surveillance” signs. Not only are people likely to miss these signs, but even if they do not, these signs fail to disclose whether video analytics is being used, for what purpose, who has access to the footage and results, and more. Our study shows that many of these attributes have a significant impact on people’s desire to be notified about deployments of video analytics. And obviously, these signs do not provide people with the ability to opt in or out of these practices.

Our findings support new disclosure requirements under regulations like GDPR, which mandates the disclosure of this information at or before the point of collection. Our findings also demonstrate the urgent need to provide people with choices to decide whether or not to allow the collection and processing of their data, as our participants expressed diverse levels of comfort with these scenarios with many not feeling comfortable with at least some of them. Regulatory disclosure requirements help improve transparency of video analytics deployments. While some study participants grew more concerned about facial recognition, we observed others becoming more accepting of it as they learned about potential benefits of some deployments. These findings suggest that increased transparency and awareness would help data subjects make informed decisions.

5.3 Privacy Preferences Are Complex and Context-Dependent

Our findings show that people’s privacy preferences are both diverse and complex. They depend on a number of contextual attributes such as the purpose for using video analytics, who has access to the results, where the user is at the time of collection, and other factors. As such, our findings are another illustration of contextual integrity principles introduced by Nissenbaum [67]. The importance of purpose information identified in our study (i.e., for what purpose video analytics is being applied) is largely consistent with results reported in earlier publications. This includes earlier work conducted by Lin et al. [57] and Smullen et al. [91] in their studies of privacy preferences when it comes to configuring mobile app permission settings. This also includes prior work by Emami-Naeini et al. [65] looking at privacy preferences across generic IoT scenarios. In contrast to these earlier studies, our work took a more systematic approach to exploring the nuances in video analytics scenarios, including the type of analysis, the purpose for which the analysis is conducted, whether information is being shared with other entities, and the venue where video analytics is deployed; those factors all have an impact on individuals’ privacy attitudes.

5.4 Implications for the Design of Privacy Assistants

Our findings can also inform the design of privacy assistants that help users manage privacy decisions related to the deployment of video analytics and other Internet of Things (IoT) technologies. Das et al. have introduced a privacy infrastructure for IoT, where users rely on “privacy assistant” mobile apps to be notified about the presence of nearby IoT resources such as cameras running video analytics software [27]. Using these privacy assistants, users can access opt-in or opt-out functionality made available by IoT resources to indicate whether they agree or not to the collection and processing of their data. However, given the growing deployment of cameras, taking advantage of such functionality would still be hampered by the number of notifications and decisions a typical person would be confronted to each day when passing within range of cameras.

A more practical approach would involve allowing users to configure privacy assistants to only notify them about those deployments they care to be notified about, and to possibly also configure any available opt-in/opt-

out settings in accordance with their individual preferences. Based on our findings, it is easy to see that different users would likely select different configurations of their settings, namely different notification settings and different combinations of opt-ins/opt-outs. To keep user burden manageable, one would likely include settings that allow users to automatically opt in or out of scenarios for which they have pretty definite preferences (e.g., “I want to opt out of any video analytics deployment that shares my data with insurance companies”). For other scenarios, they would be notified and prompted to make manual opt-in or -out decisions. Given how rich and diverse people’s privacy preferences are, enabling users to accurately specify their notification and opt-in/opt-out preferences would require a large number of privacy settings (e.g., differentiating between a variety of different video analytics deployments, different notification preferences). Recent work on privacy assistants has shown that it is possible to use machine learning to reduce user burden when it comes to configuring such complex privacy settings. For instance, Liu et al. have demonstrated the use of machine learning techniques to help users configure mobile app privacy settings [58]. Similar results have been observed by the authors using data collected as part of the present video analytics study, where models of privacy preferences were built to predict participants’ allow/deny decisions [100]. The idea is that these models are used to recommend settings to users, who can review the recommendations and decide whether or not to accept them.

Our results showing that individual’s preferences for notification of video analytics deployments are quite diverse suggest that different people would select different setting configurations, with some people preferring to be systematically informed about each deployment and being prompted to manually decide whether to opt in or out, and other people preferring more selective notification settings and greater delegation of opt-in opt-out decisions. This is also consistent with results from a recent study by Colnago et al. [22] It goes without saying that effective implementation of notification functionality and opt-in/opt-out settings such as those we just discussed, settings that our findings seem to call for, would substantially benefit from the development of standardized APIs. Ideally such APIs would enable privacy assistant functionality to (1) discover video analytics deployments in the vicinity of their users, (2) selectively notify their users, and (3) transmit opt-in or opt-out requests on their behalf (whether these requests are made manually or derived from settings selected by users).

5.5 Evolving Notification Preferences

In our study, we observed that participants’ notification preferences evolved over time with many people opting for less frequent notifications as time passes. This change in preferences is attributed to some level of fatigue as people got a better appreciation for the number of times they were likely to be notified about the same or similar scenarios, and as their level of surprise in the face of some of these scenarios also diminished over time. Even taking into account this general trend in receiving less frequent notifications over time, it is clear that people’s notification preferences are not adequately met if one relies on a simple “Ask on First Use” approach—as is typically the case today when dealing with mobile app permissions, for instance. Individuals’ notification preferences are more complex and also more diverse, ultimately requiring a more sophisticated set of configurations that users could choose from and also modify over time, as their preferences evolve. Here again we see opportunities for the use of AI-based privacy assistant functionality [23, 58] that would adapt to their user’s preferences over time, possibly through a combination of nudges designed to motivate users to think about options available to them [5, 9] and dialogues designed to capture people’s evolving preferences. Our study also uncovers how individuals’ allow/deny preferences are distinct from their notification preferences. However, how to properly notify people without overwhelming them remains an understudied direction as the majority of work on modeling privacy preferences focused on allow/deny “choice” rather than “notice.”

6 Conclusions

We reported on a 10-day experience sampling study designed to help understand individuals’ privacy attitudes related to increasingly diverse video analytics scenarios. Our study collected in-situ responses for a total of 2,328 deployment scenarios from 123 participants as they went about their regular daily activities, presenting them with video analytics scenarios that could realistically be deployed at the venues they visited. The study was informed by a systematic review of recent articles describing existing use of video analytics in support of a range of different purposes. The data collected through this study provides rich insight into people’s awareness of, comfort with, and notification preferences associated with these deployments.

As the deployment of video analytics continues to proliferate and as regulations require that users be notified about these deployments and be given opt-in/opt-out choices, it will become increasingly important for people to have access to settings that help them filter notifications and manage their privacy decisions. The complexity and diversity of individuals’ notification and opt-in/opt-out preferences observed in our study across a representative selection of video analytics scenarios suggests that the privacy settings required to capture these preferences would need to be fairly complex themselves (e.g., differentiating between a number of purposes, a number of different entities with which analysis results might be shared and more). Our study indicates that different people would likely configure these settings differently with some users preferring to see more notifications and make more decisions manually, and others opting for more selective notifications and preferring to delegate many opt-in/opt-out decisions to more powerful privacy settings. We see an opportunity for the use of privacy assistants that can help users configure such settings to accommodate their particular preferences while mitigating user burden. We also see a need for standardized APIs that help notify users and help communicate users’ opt-in/opt-out decisions.

Acknowledgments

We thank our reviewers and our shepherd, David Evans, for their time and feedback. We also thank Dr. Martin Degeling for his input and help with server development.

This research was supported in part by grants from DARPA and AFRL under the Brandeis project (FA8750-15-2-0277) and in part by grants from the National Science Foundation (NSF) Secure and Trustworthy Computing program (CNS-1513957, CNS-1801316, CNS-1914486). 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>, May 2018. Accessed: 2020-12-15.
- [2] Chinese man caught by facial recognition at pop concert. <https://www.bbc.com/news/world-asia-china-43751276>, April 2018. Accessed: 2020-12-15.
- [3] Facial recognition: School ID checks lead to GDPR fine. <https://www.bbc.com/news/technology-49489154>, August 2019. Accessed: 2020-12-15.
- [4] Facial recognition technology: Ensuring transparency in government use. <https://www.nist.gov/speech-testimony/facial-recognition-technology-ensuring-transparency-government-use>, June 2019. Accessed: 2020-12-15.
- [5] A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. F. Cranor, S. Komanduri, P. G. Leon, N. Sadeh, F. Schaub, M. Sleeper, et al. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3):1–41, 2017.
- [6] A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the facebook. In *International Workshop on Privacy Enhancing Technologies*, pages 36–58. Springer, 2006.
- [7] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005.
- [8] M. 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>, July 2018. Accessed: 2020-12-15.
- [9] H. Almuhiemedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal. Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, pages 787–796, 2015.
- [10] N. Aphorpe, Y. Shvartzshnaider, A. Mathur, D. Reisman, and N. Feamster. Discovering smart home internet of things privacy norms using contextual integrity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(2), July 2018.
- [11] R. Bachman. Your gym's tech wants to know you better. <https://www.wsj.com/articles/your-gyms-tech-wants-to-know-you-better-1497281915>, June 2017. Accessed: 2020-12-15.
- [12] S. P. 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/>, July 2015. Accessed: 2020-12-15.
- [13] L. F. Barrett and D. J. Barrett. An introduction to computerized experience sampling in psychology. *Social Science Computer Review*, 19(2):175–185, 2001.
- [14] D. Bates, M. Mächler, B. Bolker, and S. Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015.
- [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, November 2012. Accessed: 2020-12-15.
- [16] D. 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>, May 2017. Accessed: 2020-12-15.
- [17] L. L. Carstensen, B. Turan, S. Scheibe, N. Ram, H. Ersner-Hershfield, G. R. Samanez-Larkin, K. P. Brooks, and J. R. Nesselrode. Emotional experience improves with age: Evidence based on over 10 years of experience sampling. *Psychology and Aging*, 26(1):21, 2011.
- [18] R. Chow. The last mile for IoT privacy. *IEEE Security & Privacy*, 15(6):73–76, 2017.
- [19] R. H. B. Christensen. ordinal—regression models for ordinal data, 2019. R package version 2019.12-10. <https://CRAN.R-project.org/package=ordinal>.
- [20] T. C. Christensen, L. F. Barrett, E. Bliss-Moreau, K. Lebo, and C. Kaschub. A practical guide to experience-sampling procedures. *Journal of Happiness Studies*, 4(1):53–78, 2003.
- [21] L. Clark. Mannequins are spying on shoppers for market analysis. <https://www.wired.co.uk/article/mannequin-spies-on-customers>, November 2012. Accessed: 2020-12-15.
- [22] J. Colnago, Y. Feng, T. Palanivel, S. Pearman, M. Ung, A. Acquisti, L. F. Cranor, and N. Sadeh. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing System (CHI '20)*, pages 1–13, 2020.
- [23] J. Colnago and H. Guardia. How to inform privacy agents on preferred level of user control? In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*, pages 1542–1547, 2016.
- [24] B. 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>, January 2019. Accessed: 2020-12-15.
- [25] K. Conger, R. Fausset, and S. F. Kovaleski. San Francisco bans facial recognition technology. <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html>, May 2019. Accessed: 2020-12-15.
- [26] S. Consolvo and M. Walker. Using the experience sampling method to evaluate ubicomp applications. *IEEE Pervasive Computing*, 2(2):24–31, 2003.
- [27] A. Das, M. Degeling, D. Smullen, and N. Sadeh. Personalized privacy assistants for the Internet of Things: Providing users with notice and choice. *IEEE Pervasive Computing*, 17(3):35–46, 2018.
- [28] B. J. Davidson. How your business can benefit from facial recognition technology. <https://percentotech.com/how-your-business-can-benefit-from-facial-recognition-technology/>, November 2019. Accessed: 2020-12-15.

- [29] D. 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>, October 2019. Accessed: 2020-12-15.
- [30] B. Djellali, K. Belarbi, A. Chouarfia, and P. Lorenz. User authentication scheme preserving anonymity for ubiquitous devices. *Security and Communication Networks*, 8(17):3131–3141, 2015.
- [31] Y. Duan and J. Canny. Protecting user data in ubiquitous computing: Towards trustworthy environments. In *International Workshop on Privacy Enhancing Technologies*, pages 167–185. Springer, 2004.
- [32] M. 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>, December 2017. Accessed: 2020-12-15.
- [33] M. Elkhodr, S. Shahrestani, and H. Cheung. A contextual-adaptive location disclosure agent for general devices in the internet of things. In *38th Annual IEEE Conference on Local Computer Networks - Workshops*, pages 848–855. IEEE, 2013.
- [34] D. 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/>, December 2016. Accessed: 2020-12-15.
- [35] I. 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>, December 2019. Accessed: 2020-12-15.
- [36] D. Ferreira, J. Goncalves, V. Kostakos, L. Barkhuus, and A. 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.
- [37] C. 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>, March 2016. Accessed: 2020-12-15.
- [38] S. F. Gale. Employers turn to biometric technology to track attendance. <https://www.workforce.com/news/employers-turn-to-biometric-technology-to-track-attendance>, March 2013. Accessed: 2020-12-15.
- [39] P. Grother, M. Ngan, and K. Hanaoka. Ongoing face recognition vendor test (FRVT) part 2: Identification. <https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8238.pdf>, November 2018. Accessed: 2020-12-15.
- [40] P. Grother, M. Ngan, and K. Hanaoka. Face recognition vendor test (FRVT) part 3: Demographic effects. <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf>, December 2019. Accessed: 2020-12-15.
- [41] Y. Gurovich, Y. Hanani, O. Bar, G. Nadav, N. Fleischer, D. Gelbman, L. Basel-Salmon, P. M. Krawitz, S. B. Kamphausen, M. Zenker, L. M. Bird, and K. W. Gripp. Identifying facial phenotypes of genetic disorders using deep learning. *Nature Medicine*, 25(1):60–64, 2019.
- [42] J. M. Hektner, J. A. Schmidt, and M. Csikszentmihalyi. *Experience sampling method: Measuring the quality of everyday life*. Sage, 2007.
- [43] W. Hofmann, R. F. Baumeister, G. Förster, and K. D. Vohs. Everyday temptations: An experience sampling study of desire, conflict, and self-control. *Journal of Personality and Social Psychology*, 102(6):1318, 2012.
- [44] J. I. Hong and J. A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services (MobiSys '04)*, pages 177–189, 2004.
- [45] T. Johnson. Shoplifters meet their match as retailers deploy facial recognition cameras. <https://www.mcclatchydc.com/news/nation-world/national/article211455924.html>, May 2018. Accessed: 2020-12-15.
- [46] E. Kanjo, L. Al-Husain, and A. Chamberlain. Emotions in context: examining pervasive affective sensing systems, applications, and analyses. *Personal and Ubiquitous Computing*, 19(7):1197–1212, 2015.
- [47] D. Korgut and D. F. Pigatto. An internet of things-based house monitoring system. In *2018 IEEE Symposium on Computers and Communications (ISCC)*, pages 01149–01152, June 2018.
- [48] I. 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.
- [49] S. 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>, July 2019. Accessed: 2020-12-15.
- [50] J. Kruger and D. Dunning. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6):1121–1134, 1999.
- [51] H. Lee and A. Kobsa. Understanding user privacy in internet of things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, pages 407–412, 2016.
- [52] H. Lee and A. Kobsa. Privacy preference modeling and prediction in a simulated campuswide IoT environment. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 276–285, 2017.
- [53] S. 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>, March 2019. Accessed: 2020-12-15.
- [54] D. 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>, October 2017. Accessed: 2020-12-15.
- [55] D. 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>, June 2019. Accessed: 2020-12-15.
- [56] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and purpose: Understanding users' mental models of mobile app privacy through crowdsourc-

- ing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*, pages 501–510. ACM, 2012.
- [57] J. Lin, B. Liu, N. Sadeh, and J. I. Hong. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Proceedings of the Tenth Symposium on Usable Privacy and Security (SOUPS '14)*, pages 199–212, 2014.
- [58] B. Liu, M. S. Andersen, F. Schaub, H. Almuhammedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS '16)*, pages 27–41, 2016.
- [59] B. Liu, J. Lin, and N. Sadeh. Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In *Proceedings of the 23rd International Conference on World Wide Web (WWW '14)*, pages 201–212, New York, NY, USA, 2014.
- [60] B. 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>, October 2014. Accessed: 2020-12-15.
- [61] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004.
- [62] L. Y. Mano et al. Exploiting IoT technologies for enhancing health smart homes through patient identification and emotion recognition. *Computer Communications*, 89-90:178–190, 2016.
- [63] K. Martin and K. Shilton. Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices. *The Information Society*, 32(3):200–216, 2016.
- [64] D. 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/>, June 2011. Accessed: 2020-12-15.
- [65] P. E. Naeini, S. Bhagavatula, H. Habib, M. Degeling, L. Bauer, L. F. Cranor, and N. Sadeh. Privacy expectations and preferences in an iot world. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS '17)*, pages 399–412, 2017.
- [66] NEC Corporation. New biometric identification tools used in theme parks. <https://www.nec.com/en/global/about/mitatv/03/3.html>, 2002. Accessed: 2020-12-15.
- [67] H. Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1):119, 2004.
- [68] H. Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- [69] PCMag Stuff. NEC unveils facial-recognition system to identify shoppers. <https://www.pcmag.com/archive/nec-unveils-facial-recognition-system-to-identify-shoppers-305015>, November 2012. Accessed: 2020-12-15.
- [70] V. Pejovic, N. Lathia, C. Mascolo, and M. Musolesi. *Mobile-Based Experience Sampling for Behaviour Research*, pages 141–161. Springer International Publishing, 2016.
- [71] C. Perera, R. Ranjan, L. Wang, S. U. Khan, and A. Y. Zomaya. Big data privacy in the internet of things era. *IT Professional*, 17(3):32–39, 2015.
- [72] C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos. Context aware computing for the internet of things: A survey. *IEEE Communications Surveys & Tutorials*, 16(1):414–454, 2013.
- [73] P. Porambage, M. Ylianttila, C. Schmitt, P. Kumar, A. Gurtov, and A. V. Vasilakos. The quest for privacy in the internet of things. *IEEE Cloud Computing*, 3(2):36–45, 2016.
- [74] J. Porter. Federal study of top facial recognition algorithms finds 'empirical evidence' of bias. <https://www.theverge.com/2019/12/20/21031255/facial-recognition-algorithm-bias-gender-race-age-federal-nest-investigation-analysis-amazon>, December 2019. Accessed: 2020-12-15.
- [75] S. Prabhakar, S. Pankanti, and A. K. Jain. Biometric recognition: Security and privacy concerns. *IEEE Security & Privacy*, 1(2):33–42, March 2003.
- [76] Press Association. 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>, November 2013. Accessed: 2020-12-15.
- [77] E. 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>, February 2019. Accessed: 2020-12-15.
- [78] E. Ramirez, J. Brill, M. K. Ohlhausen, J. D. Wright, and T. McSweeney. Data brokers: A call for transparency and accountability. Technical report, Federal Trade Commission, May 2014.
- [79] B. Rashidi, C. Fung, and T. Vu. Dude, ask the experts!: Android resource access permission recommendation with RecDroid. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pages 296–304, 2015.
- [80] R. W. Reeder, A. P. Felt, S. Consolvo, N. Malkin, C. Thompson, and S. Egelman. 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 (CHI '18)*, pages 1–13, 2018.
- [81] T. 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/>, January 2017. Accessed: 2020-12-15.
- [82] D. Rosen. Disney is spying on you! https://www.salon.com/test/2013/01/17/disney_is_spying_on_you/, January 2013. Accessed: 2020-12-15.
- [83] T. S. Saponas, J. Lester, C. Hartung, S. Agarwal, and T. Kohno. Devices that tell on you: Privacy trends in consumer ubiquitous computing. In *16th USENIX Security Symposium (USENIX Security '07)*, pages 55–70, 2007.
- [84] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A design space for effective privacy notices. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS '15)*, pages 1–17, 2015.
- [85] 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>, May 2015. Accessed: 2020-12-15.

- [86] S. Shan, E. Wenger, J. Zhang, H. Li, H. Zheng, and B. Y. Zhao. Fawkes: Protecting privacy against unauthorized deep learning models. In *29th USENIX Security Symposium (USENIX Security '20)*, pages 1589–1604, August 2020.
- [87] M. Sharif, S. Bhagavatula, L. Bauer, and M. 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.
- [88] F. Shih, I. Liccardi, and D. Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, pages 807–816, 2015.
- [89] E. 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/>, October 2019. Accessed: 2020-12-15.
- [90] A. Smith. More than half of U.S. adults trust law enforcement to use facial recognition responsibly. Technical report, Pew Research Center, September 2019.
- [91] D. Smullen, Y. Feng, S. Zhang, and N. M. Sadeh. The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences. *Proc. Priv. Enhancing Technol.*, 2020(1):195–215, 2020.
- [92] B. Snyder. This beer ad only works when women pass by. <https://fortune.com/2015/05/21/astra-beer-ad/>, May 2015. Accessed: 2020-12-15.
- [93] U.S. Government Accountability Office. Face recognition technology: FBI should better ensure privacy and accuracy. <https://www.gao.gov/assets/680/677098.pdf>, May 2016. Accessed: 2019-11-22.
- [94] N. Van Berkel, D. Ferreira, and V. Kostakos. The experience sampling method on mobile devices. *ACM Computing Surveys (CSUR)*, 50(6):1–40, 2017.
- [95] S. J. Verhagen, L. Hasmi, M. Drukker, J. van Os, and P. A. Delespaul. Use of the experience sampling method in the context of clinical trials. *Evidence-based Mental Health*, 19(3):86–89, 2016.
- [96] J. 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>, November 2018. Accessed: 2020-12-15.
- [97] P. Wijesekera, A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov. The feasibility of dynamically granted permissions: Aligning mobile privacy with user preferences. In *2017 IEEE Symposium on Security and Privacy*, pages 1077–1093, 2017.
- [98] P. Wijesekera, J. Reardon, I. Reyes, L. Tsai, J.-W. Chen, N. Good, D. Wagner, K. Beznosov, and S. Egelman. Contextualizing privacy decisions for better prediction (and protection). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, pages 1–13, 2018.
- [99] Z. Wu, S.-N. Lim, L. S. Davis, and T. Goldstein. Making an invisibility cloak: Real world adversarial attacks on object detectors. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *Computer Vision – ECCV 2020*, pages 1–17. Springer International Publishing, 2020.
- [100] S. A. Zhang, Y. Feng, A. Das, L. Bauer, L. Cranor, and N. Sadeh. Understanding people’s privacy attitudes towards video analytics technologies. Technical Report CMU-ISR-20-114, Carnegie Mellon University, School of Computer Science, December 2020.

Appendix

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	Homemaker	2.4
Education	8.1	Skilled labor	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 5. Occupation of participants and respective %

Factors	Est.	Std. Err	Z	p
Intercept	-1.79965	0.60789	-2.96	0.003072**
purpose:baseline = Generic Surveillance				
Petty Crime(Anon)	0.57922	0.52134	1.111	0.266563
Criminal Detection(IDed)	1.08567	0.43613	2.489	0.012799*
Count People(Anon)	0.54011	0.56511	0.956	0.339187
Jump Line(IDed)	2.12133	0.53749	3.947	7.92E-05***
Targeted Ads(Anon)	2.77327	0.56614	4.899	9.66E-07***
Targeted Ads(IDed)	1.87295	0.5265	3.557	0.000375***
Sentiment Ads(Anon)	2.03323	0.70039	2.903	0.003696**
Sentiment Ads(IDed)	2.7837	0.59923	4.645	3.39E-06***
Rate Service(Anon)	1.92574	0.55494	3.47	0.00052***
Rate Engagement(IDed)	0.9621	0.92536	1.04	0.298478
Face as ID(IDed)	1.70491	0.51797	3.292	0.000997***
Track Attendance(IDed)	2.56281	0.60284	4.251	2.13E-05***
Work Productivity(IDed)	3.15627	0.63879	4.941	7.77E-07***
Health Predictions(IDed)	3.37146	0.58706	5.743	9.30E-09***
Medical Predictions(IDed)	1.92103	0.7824	2.455	0.014077*
Raw retention:baseline=30 days				
Ephemeral	0.10859	0.3799	0.286	0.775005
Unspecified	0.23487	0.4079	0.576	0.564742
Analytics retention:baseline=unspecified				
Ephemeral	-0.02068	0.81819	-0.025	0.979836
30 days	-0.22812	0.30495	-0.748	0.454423
Association: baseline=No				
associationID	0.27251	0.18042	1.51	0.130937
Shared: baseline=No				
sharedID	-0.09074	0.26258	-0.346	0.729666
dayIndex	0.79628	0.27167	2.931	0.003378**
placeType:baseline=large public places				
store	0.73456	0.42748	1.718	0.085732
eatery	1.09194	0.41956	2.603	0.009252**
work	0.46835	0.50123	0.934	0.350094
education	-0.48813	0.50161	-0.973	0.330493
hospital	1.11144	0.65184	1.705	0.088178
service	0.67614	0.52179	1.296	0.195037
alcohol	0.81001	0.4635	1.748	0.08053
entertainment	0.80385	0.61804	1.301	0.193377
fitness	1.06873	0.66162	1.615	0.10624
gas	1.36253	0.58379	2.334	0.019598*
transportation	-1.48697	0.5998	-2.479	0.013171*
worship	-0.27275	0.81689	-0.334	0.738463
library	1.71228	0.71968	2.379	0.01735*
mall	1.19774	0.89793	1.334	0.182241
airport	0.08364	0.96362	0.087	0.930832
finance	-1.13355	1.16506	-0.973	0.33058

Table 6. Generalized Linear Mixed Model Regression with Logit Link. A positive coefficient(estimate) shows likeliness of participants' to deny a data collection

Purpose	Scenario Text
Generic Surveillance	Some places like %s have started to deploy video surveillance cameras to deter crime. (This footage can be shared with law enforcement.) Assume that you are captured by such a camera, and the raw footage is kept for 30 days.
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). When a suspicious scene is believed to have been detected, it is recorded for further analysis (possibly including facial recognition) and kept for 30 days. Otherwise the data is immediately discarded. Assume that you are captured by such a camera.
Known Criminal	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. Assume that %s engages in this practice, and the raw footage is discarded immediately, with the analysis results being kept for 30 days.
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. Assume that %s engages in this practice and it is unclear for how long all the data (raw footage and analysis results) is kept.
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. This software can also record your presence and who you are with. Assume that %s engages in this practice and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.
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. Assume that %s engages in this practice and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.
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. Assume that %s engages in this practice and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.
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. Assume that %s engages in this practice and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.
Sentiment Ads(IDed)	Some places like %s have started to deploy video surveillance cameras with facial recognition and emotion analysis software. This software recognizes people, and analyzes their reactions to items displayed. Then the software matches detected faces against individual customer profiles to send tailored deals and coupons to their phones. Assume that %s engages in this practice and the raw footage is kept for 30 days, and it is unclear for how long the analysis results are kept.
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. Assume that %s engages in this practice and it is unclear for how long all the data (raw footage and analysis results) is kept.
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. This software can be used to record your presence and also identify who you are with. Assume that %s engages in this practice and the raw footage is kept for 30 days, and it is unclear for how long the analysis results are kept.
Face as ID	Some stores have started to deploy video surveillance cameras with facial recognition software. This software can identify faces of customers to replace membership cards at checkout. Assume that %s engages in this practice, and the raw footage is discarded immediately. Assume also that it is unclear for how long the analysis results are kept.
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. This software can also identify how long you participate in different activities and who you hang out with. Assume that your workplace engages in this practice, and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.
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. Assume that your workplace engages in this practice, and it is unclear for how long all the data (raw footage and analysis results) is kept.
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 can impact your health insurance premium. Assume that %s engages in this practice, and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.
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, and impact your health insurance premium. Assume that %s engages in this practice, and the raw footage is kept for 30 days. Assume also that it is unclear for how long the analysis results are kept.

Table 7. Scenarios text shown to participants