

# Understanding People’s Privacy Attitudes Towards Video Analytics Technologies

Shikun Aerin Zhang  
*School of Computer Science*  
*Carnegie Mellon University*  
Pittsburgh, PA 15213 USA  
shikunz@cs.cmu.edu

Yuanyuan Feng  
*School of Computer Science*  
*Carnegie Mellon University*  
Pittsburgh, PA 15213 USA  
yuanyuanfeng@cmu.edu

Anupam Das  
*Department of Computer Science*  
*North Carolina State University*  
Raleigh, NC 27695 USA  
anupam.das@ncsu.edu

Lujo Bauer  
*School of Computer Science*  
*Carnegie Mellon University*  
Pittsburgh, PA 15213 USA  
lbauer@cmu.edu

Lorrie Cranor  
*School of Computer Science*  
*Carnegie Mellon University*  
Pittsburgh, PA 15213 USA  
lorrie@cmu.edu

Norman Sadeh  
*School of Computer Science*  
*Carnegie Mellon University*  
Pittsburgh, PA 15213 USA  
sadeh@cs.cmu.edu

**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 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. 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 people’s 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, or who it is shared with. Finally, we consider recent technical advances that would enable entities that deploy video analytics software to selectively apply it only to footage of people who provide consent (“opt in”). New regulations such as the General Data Protection Regulation actually mandate obtaining such consent “at or before the point of collection.” Because obtaining consent from users at or before each point of collection could result in significant user burden, we use our data to explore the development of predictive models that could one day help people manage such consent. Our results are rather encouraging.

**Index Terms**—usable privacy and security, facial recognition

## I. INTRODUCTION

In August of 2019, a high school in Sweden received the first fine under the European Union’s General Data Protection Regulation (GDPR) for its use of facial recognition to track students’ attendance [1]. This comes at a time when facial recognition has prompted increased scrutiny from both privacy advocates and regulators [2], [3]. Facial recognition is a type of video analytics technology that has become increasingly accurate with recent advances in deep learning and computer vision [4]. The increasing ubiquity of facial recognition is contributing to the collection and inference of vast amounts of personal information, including not only people’s whereabouts, their activities, who they are with but also their mood,

health and behavior. As the accuracy of algorithms improves and as data continues to be collected across an ever wider range of scenarios, inferences made from this data can be expected to reveal even more sensitive information about individuals. To make things worse, such data collection and usage often take place without people’s awareness or consent. While facial recognition could benefit different entities (e.g., law enforcement, businesses), its broad deployment raises important privacy questions [5]. In the US, the GAO and NIST have recommended more transparency when it comes to appropriate use of facial recognition [6], [7]. New regulations such as the GDPR and the California Consumer Privacy Act (CCPA) mandate specific disclosure and choice requirements, which extend to the deployment of facial recognition (e.g., opt-in or opt-out). While these regulations are important steps towards providing data subjects with more information about and more control over personal data privacy, they do not address significant issues such as how often people should be notified about the presence of facial recognition and how one could ultimately help them also take advantage of choice options required by these new regulations.

Our research aims to address these issues by developing a more comprehensive understanding of how people feel about the deployment of facial recognition 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. Our study is organized around two broad sets of questions.

The first set focuses on understanding people’s privacy expectations and preferences. This includes looking for possible social norms that might extend to larger population for particular deployment scenarios [8], or alternatively identifying variability in how different people respond to various deployments of facial recognition.

The second set of questions is motivated by recent tech-

nical advances introduced by Das et al. [9], namely (1) the development of real-time face denaturing functionality that enables video analytics software to only be applied to people who provide consent, and (2) the development of a privacy infrastructure for the Internet of Things (IoT) that enables entities deploying facial recognition software to publicize their data practices and allow data subjects to opt in or out of having their footage analyzed and/or shared. Such functionality effectively enables these entities to comply with regulations such as GDPR or CCPA, which require notifying data subjects and enabling them to opt in or out of some practices at or before the point of collection. Because expecting people to manually opt in or out of facial recognition each time they encounter such functionality entails a unrealistically high number of privacy decisions, we use our data to explore the feasibility of developing predictive models to assist users with their privacy decisions, and discuss different possible deployment strategies for such models.

The main contributions of this work are as follows:

- We conducted a first longitudinal in situ study of people’s privacy expectations and preferences across a wide range of video analytics deployment scenarios. We offer an in-depth analysis of the data collected as part of this study — 10-day study involving 123 participants who provided us with detailed insight into their degree of awareness and comfort as they related to a total of 2,328 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 are likely to receive if they are not selective in their notification preferences.
- We use the data collected as part of our study to explore the feasibility of developing predictive models to help people cope with the large number of allow/deny decisions they would otherwise have to make each time they encountered facial recognition deployments. We show that using clustering techniques, it is possible to accurately predict people’s privacy decisions across many deployment scenarios and discuss different possible configurations for using these models.

## II. RELATED WORK

### A. Privacy Challenges of Video Analytics

Video analytics, many with facial recognition functionality, is increasingly being integrated with the Internet of Things (IoT) systems, as one of the newest applications of ubiquitous computing [10]–[12]. Data privacy has been a central discussion in IoT [13] because IoT systems rely on the collection and use of contextual information in the environments (e.g.,

people, time, location, activity) that often contains identifiable personal data [14]–[16]. Researchers have explored technical solutions to safeguard user data in IoT [17]–[19]. However, transparency around IoT data privacy remains an unsolved issue [16], [20]. People often have no way to know the existence of video analytics applications in their daily environments, what personal data is being collected, how the footage is used for what purpose by whom, and how long the footage will be retained. Moreover, video analytics face unique data privacy challenges. First, it can collect people’s biometric data (e.g., facial features, body pose) [21] that is considered more sensitive than other digital identifiers like email addresses. Second, it can be applied later on video footage already collected by existing cameras for a myriad of purposes (e.g., security, operation optimization, targeted advertising).

These challenges indicate 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 [8] 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. Changes to these parameters are likely to cause a privacy norm violation and must be examined closely [22]. 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 surveying people’s opinions and attitudes towards the technology.

### B. Sampling and Modeling People’s 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 [23], [24]. However, vignette studies are often limited because participants have to imagine themselves in hypothetical scenarios that are not immediately relevant [25]. 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 [26]. Particularly, mobile-based ESM can prompt participants in light of context, enabling the collection of higher quality, more valid research data [27], [28]. 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.

Previous research on privacy preference modeling is also relevant. Researchers have used data-driven approach to identify patterns of people’s privacy expectations and preferences

within mobile app permission [29], [30] and some IoT contexts [24], [31], [32]. Previous research show privacy preferences vary greatly from person to person and from case to case [31], [33], [34]. This indicates there is no one-size-fits-all solution to accommodate people’s diverse privacy preferences to address data privacy issues in mobile and IoT contexts. Some research focuses on building machine learning models to predict people’s privacy preferences [35], [36], which is a promising approach to better provide privacy notices and help people configure their privacy settings [37], [38]. We extend this line of work by exploring predictive models of people’s privacy preferences towards different deployments of video analytics using data collected through ESM.

### C. Designing and Implementing Privacy Assistants

In web and mobile environments, many privacy settings are available to users, such as to stop websites from tracking personal browsing data or to prevent mobile apps from accessing location data on smartphones. In reality, users often struggle to configure privacy settings to match their privacy preferences, either due to unintelligible privacy notice [39] or the unreasonable amount of efforts required to manage their privacy [34], [40]. This renders most existing privacy settings unusable and thus ineffective for personal privacy management [33].

To address these usability issues, recent research advocates for “privacy assistants”, which are software tools capable of informing people about sensitive data privacy practices [41] and helping them configure a large number of privacy settings [42]. Privacy assistants can be enhanced by incorporating machine learning models for privacy preferences to further reduce user burden [38], [43]. For example, Liu et al. [37] have implemented an Android privacy assistant app which generates personalized privacy recommendations based on data-driven privacy profiles, operationalizing personalized privacy assistants in a real-world mobile context.

With unique data privacy challenges, there is growing research exploring privacy assistants in IoT [44], [45]. Privacy assistants have different levels of automation: some only deliver machine-readable privacy notices and totally leave privacy decisions to users [46]; some provide personalized recommendations to help users make better privacy decisions [37]; some can make privacy decisions for users in an automated manner [44]. A recent user study reported that people have varying attitudes towards privacy assistants of different automation levels in IoT contexts, largely based on their personal weighing of desire for privacy control in IoT against the concern of cognitive overload [47].

Ongoing technology advances indicate the feasibility to implement privacy assistants for IoT systems with video analytics functionality [9]. For example, technology that facilitates privacy notice and choice in IoT [45] can serve as the infrastructure to make privacy settings of IoT systems accessible to users. Also, computer vision solutions like RTFace [48] provide people the capability of not having their facial data analyzed by obfuscating their faces in live video streams. This

further motivates our ESM study to understand people privacy expectations and preferences towards video analytics, which opens the door for privacy assistants that help safeguard users’ privacy in IoT systems with video analytics.

## III. DESIGNING AN EXPERIENCE SAMPLING STUDY

### A. Experience Sampling Method

Context has been shown to play an important role in influencing people’s privacy attitudes and decisions (e.g., contextual integrity [22]). 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 opted to design an experience sampling study, where we collected information about people’s responses to a variety of video analytics deployments (or “scenarios”) in the context of their regular everyday activities. The experience sampling method [26] has been repeatedly used in clinical trials [49], [50], psychological experiments [51], [52] and human-computer interaction (HCI) studies [53], [54], yielding “a more accurate representation of the participants’ natural behaviour” [55]. This enables us to engage and survey participants in a timely and ecologically valid manner as they go about their normal daily lives [56]. Participants are prompted to answer questions about plausible video analytics scenarios at places representative of their actual whereabouts.

### B. Selecting Realistic Scenarios

Previous research mainly surveyed participants’ privacy attitudes in the context of generic IoT scenarios, including some facial recognition scenarios [24], [32]. 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, marketing, authentication, employee performance evaluation, and even church attendance tracking). 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 to ensure that the scenarios presented to them were consistent with the scenarios identified in our survey.

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 facial recognition. *Security-related* scenarios of facial recognition included automatic detection of petty crime using computer vision [57], and identification of known shoplifters and criminals in public places [58]–[61]. Facial recognition scenarios for *commercial* purposes included helping businesses to optimize operations [62]–[64], displaying personalized advertisements based

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

**TABLE I:** 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 facial recognition technology.

on the detection of demographic features [61], [65]–[67], collecting patrons’ facial reaction to merchandise [68]–[71], detecting users’ engagement at entertainment facilities [72]–[74]. 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 [75]–[78], and (2) attendance tracking in the workplace, at churches, and at gyms [77], [79], [80]. 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 [81]–[83]; employers using emotion analysis to evaluate employee performance [84]–[87]; hospitals using emotion recognition to make health-related predictions [85], [88], [89].

In total, we identified 16 purposes, as shown in Table I, representative of a diverse set of video analytics scenarios. A list of the scenarios as well as the corresponding text showed to participants to elicit their reactions can be found in the Appendix (Table V). The text associated with each scenario was carefully crafted through multiple iterations to make scenario descriptions as plausible as possible without deceiving participants.

### C. Factorial Design

We employed a factorial study design and developed a taxonomy that captured a representative set of attributes one might expect to influence people’s privacy attitudes. These attributes are shown in Table I. 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 afterwards.

### D. Study Protocol and Procedures

The 10-day study comprised the following four stages.

**Stage 1:** Eligible participants who completed the consent forms for this study were able to download the study app from the Google Play Store. Upon first installing the app, participants were asked to complete a pre-study survey asking 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, designed and implemented by the research team, 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 6 per day, though most of the time participants’ whereabouts would trigger a smaller number of scenarios - closer to 3 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 app, implemented by the research team, 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 administrated via Qualtrics. This survey contained free-response questions about their attitudes on facial recognition technology, their responses to three scenarios, the 10-item IUIPC scale on privacy concerns [90], as well as additional

demographic questions like income, education level, marital status and whether they live with children under 18.

### E. Payment and Study Length

To maximize the contextual benefits provided by the experience sampling method [91], 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 rewarded when they finished the study. In total, participants can earn at least \$37 and up to \$52.

### F. Ensuring Study Validity

Due to the complexity and the number of different components of the study framework, we conducted several rounds of pilot studies first with members of the research team, and then with a small number (N=3) real participants each round. We further refined our study and ensured that our protocol and all technical components (study app, web survey app, study server) function as expected after each round of piloting.

1) *Option to say “I was not there”*: Due to varying accuracy of GPS on smartphones, instead of assuming participants’ locations solely based on GPS, in each notification we asked them to first select the place that they were at from a drop-down list of nearby places. Participants were provided 3 additional options: “I was somewhere else in the area”, “I was passing by”, and “I was not there” to account for missing places or cases when they were passing by (e.g., being stuck in traffic). Participants still received payments for each scenario when they have selected these three choices above so they have no incentive to select a place which they did not visit.

2) *Attention Check Questions*: After the first pilot, we found that some participants did not read the scenario text carefully. As a result, we implemented two multiple-choice attention check questions, each of which was randomly generated to query about one of the six attributes (attributes in Table I excluding type of places). They can only proceed to answer the four in situ questions once they passed the attention check. We recorded the number of attempts and the time spent to answer attention check questions for every scenario.

### G. Technical Optimizations

Conducting in situ studies often requires substantial efforts to deal with a number of uncertainties and challenges [92]. Here we give a few examples of ethical, pragmatic, and logistical concerns that we need to navigate as part of our in situ study. Firstly, we refrained from using off-the-shelf experience sampling software, and developed our own system and location-aware Android app because location data collected over a period of time can be particularly sensitive. Our app needed to be able to detect participants’ location at all times without the GPS draining their phone battery. To

optimize our app’s performance, we effectively incorporated the Android Activity Recognition API to not collect location while the phone remains still. Secondly, our study relies on accurately identifying the types of places that participants visit. This is crucial not only because places serve as an important contextual attribute, but also because we need to know the place types to show relevant scenarios. We queried both the Google Geo-coding API, which includes more locations without meaningful place categories, and the Foursquare API, which offers detailed place types but also contains noisy crowd-sourced entries. We cross-checked both sources to provide a list of nearby places for participants to choose from. Thirdly, since our study requires participants to visit places of interest to receive notifications, it is expected that people may not do so every day. We designed the study to allow participants to remain in the study for up to 5 days when they did not visit places that would trigger notifications.

### H. Recruitment

We recruited participants using four methods: posts on local online forums (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 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, occupation, and self-reported frequencies at which they typically visit different venues, like restaurants, shops, etc. All 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.

## IV. ANALYZING PRIVACY ATTITUDES

### A. 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, among which 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, as shown in Table II.

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 II:** Occupation of participants and respective %

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 different 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 while riding the bus, sitting in a car, biking or just walking. 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 commercial establishments, or small local stores.

For the 65% of the 3,589 notifications, once participants had reported their location, they were presented with a plausible scenario given their reported location, and were 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 in the form of an Amazon gift card.

The median response time taken by participants to answer each scenario in situ was 42.0 seconds, of which 18.5 seconds were spent answering two attention check questions designed to ensure that they had familiarized themselves with relevant contextual attributes associated with the selected scenario - these attributes were highlighted using bold typeface, as illustrated in Fig. 1. The median response time to complete additional evening questions for a given scenario was 90.3 seconds. Fig. 1 displays box plots of the time spent completing each step associated with a given scenario. The data collected included a few outliers. For example, a few evening reviews

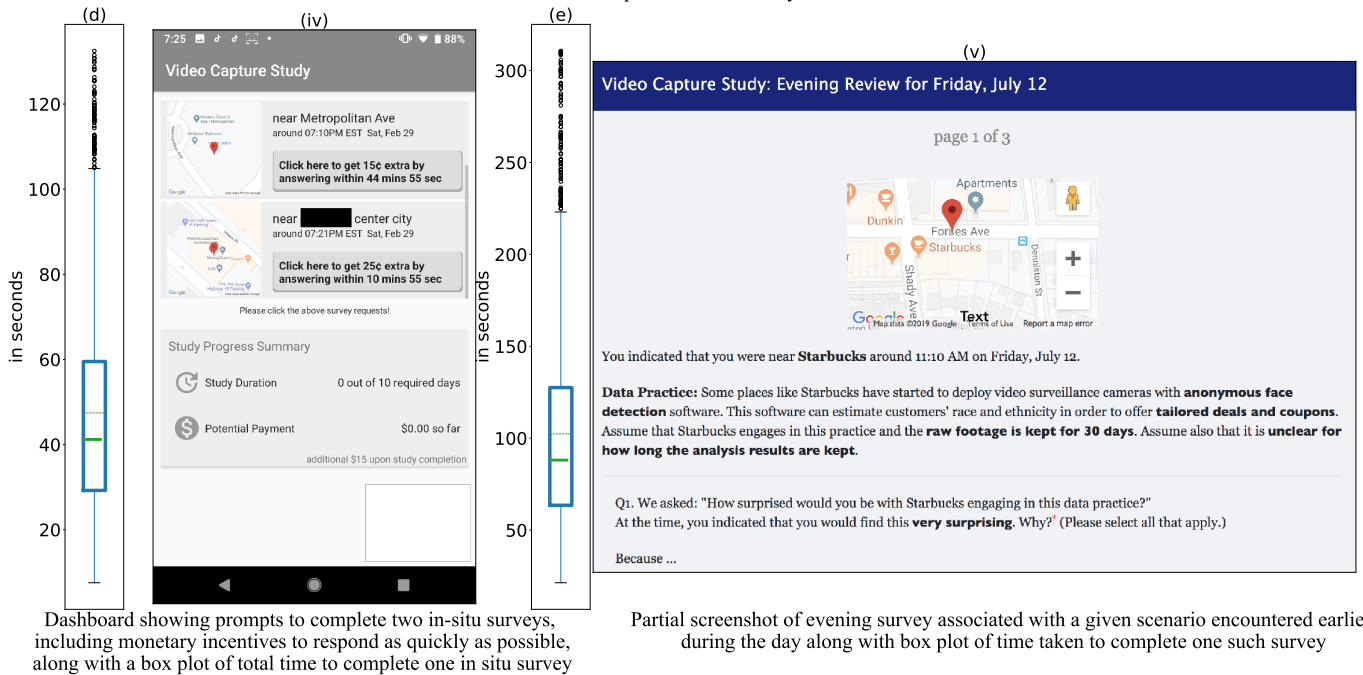
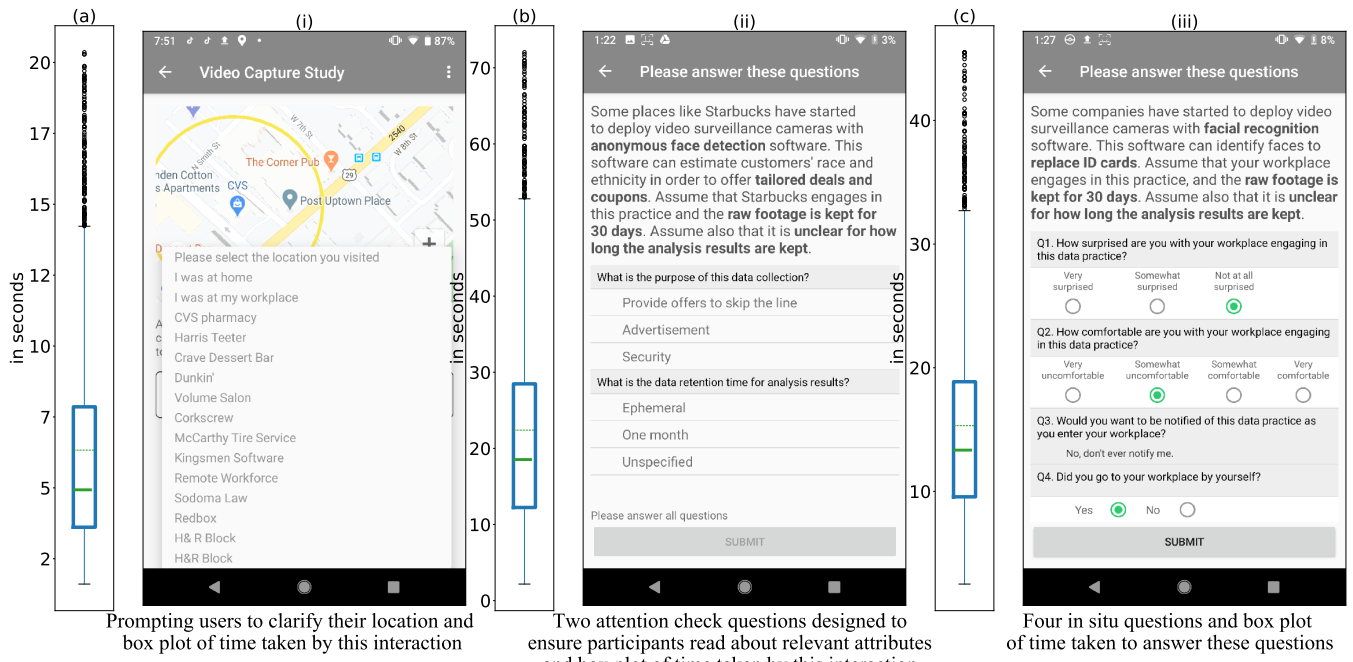
took participants more than 2500 seconds (i.e., more than 40 minutes) to answer, which most likely corresponds to participants temporarily interrupting their reviews and completing them later. These outliers were omitted in the plots shown in Fig. 1 to avoid distorting the scale.

Data on attention check questions and when participants responded to in situ notifications indicates relatively high validity of the responses received. 59% of the time, participants successfully completed both attention check questions associated with the scenarios assigned to them in their first attempt. 81% of the time, they did so within two attempts. See Fig. 2 for details. We interpret these results as an indication that the attention checks were necessary to make sure people paid attention to the contextual attributes associated with scenarios assigned to them. These results also suggest that the responses we collected most likely reflect privacy attitudes that take these contextual attributes into account, since people had no choice but to take the time to familiarize themselves with these details. Additionally, 68% of in situ questions were answered within 15 minutes and 87% within 1 hour. In other words, the actual context within which a participant had visited the location associated with the scenario was likely still fresh in their mind (e.g., what the participant was doing at that particular location, who they might have been with, etc.).


### B. Collecting People’s Privacy Attitudes

When surveying participants’ privacy responses to different facial recognition scenarios, we decided to focus on four related sets of questions, namely 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 (e.g., opt-in or opt-out), 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 were worded as follows - with *Controller* being a variable that would be instantiated with the name of the venue participants were visiting:

- 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 I am notified or not.



**Fig. 1:** Screenshots of the study app and the web survey app used for the evening review, including box plots of the response times required to complete each step associated with participants answering a given scenario in situ and in the evening.



**Fig. 2:** The cumulative percentage of scenarios answered plotted against the number of attempts

- No, don't ever notify me.
  - If you had the choice, would you allow or deny this data practice?
    - Allow, Deny
- Fig. 3 provides a summary of collected responses organized around the 16 different categories of scenarios (or “purposes”) introduced in Table I. As can be seen, people’s responses vary for each scenario. In other words, “one size fits all” would fail to capture people’s diverse preferences when presented with

these different 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 different 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.

### C. 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 [29]. We compiled pairwise correlations between the four types of responses collected from our participants across the 2,328 scenarios evaluated in our study (Table III). 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, 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 with 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.

To further confirm our finding of a moderate correlation between surprise and comfort, we also looked at correlation between these two variables when limiting our data set to the first instance when each participant encountered a given scenario - since seeing the exact same scenario multiple times is likely to result in less surprise over time. Computing correlation between these two variables when limiting ourselves to these first instances (total of 1,213 out of the 2,328 data points), yielded a similarly moderate effect size ( $r = 0.435, p < 0.001$ ).

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

TABLE III: Correlation matrix where \*\*\* indicates  $p < 0.001$

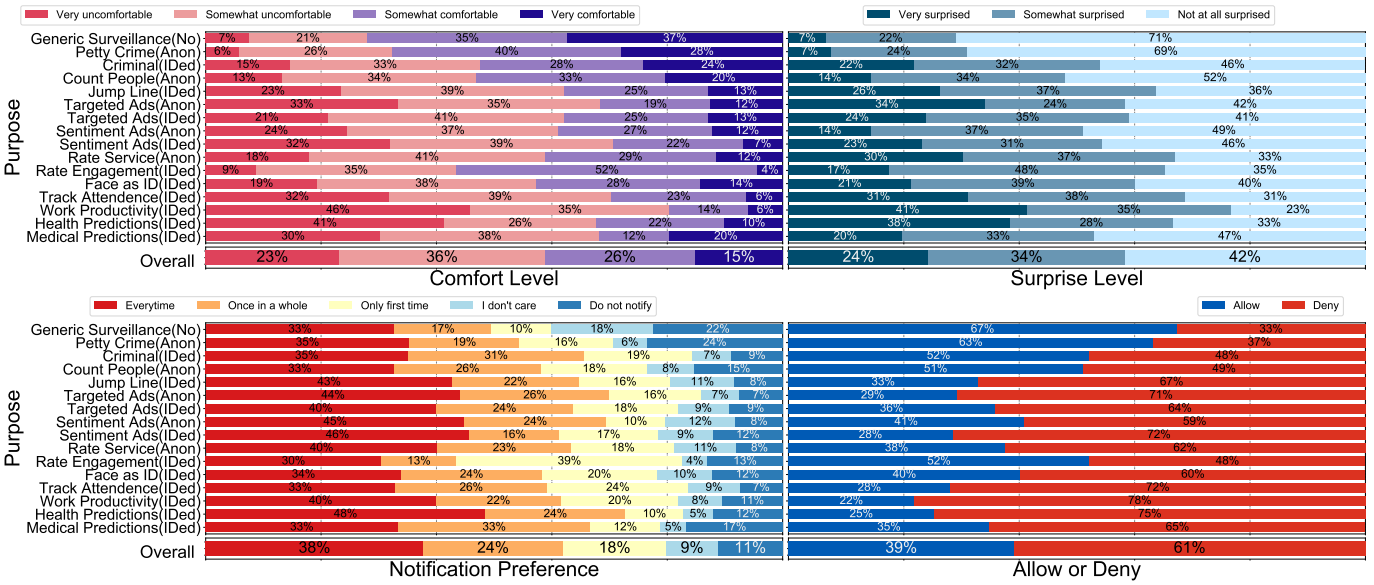
### D. Factors Impacting People’s Privacy Attitudes

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 I. 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 I, 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.

1) *Allow/Deny Decisions*: We first investigate whether people’s decisions to allow or deny data collection have a relationship with the contextual attributes in Table I. We constructed our model using generalized linear mixed model (GLMM) regression [93], 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 I, 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 practices (odds ratio =  $e^{1.70} = 5.5$ ). Three purposes turned out not to be significant: detecting petty crime, and 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 it 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, amusement parks, etc. Contrary to our expectations, we found that whether targeted ads relied on identifying individuals or whether they





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

treated them anonymously did not elicit substantially different responses from our participants. In fact, participants were 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.

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, and etc.), we find that participants were more likely to deny data practices at eateries (odds ratio =  $e^{1.09} = 3$ ), at libraries (odds ratio =  $e^{1.71} = 5.5$ ), at gas stations (odds ratio =  $e^{1.36} = 3.9$ ). Participants were less likely to respond with a “deny” to deployment scenarios at transportation locations (buses stops, train stations, metro stations) than at the baseline (odds ratio =  $e^{-1.87} = 0.23$ ).

The day of the study when participants were exposed to a scenario 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 IV).

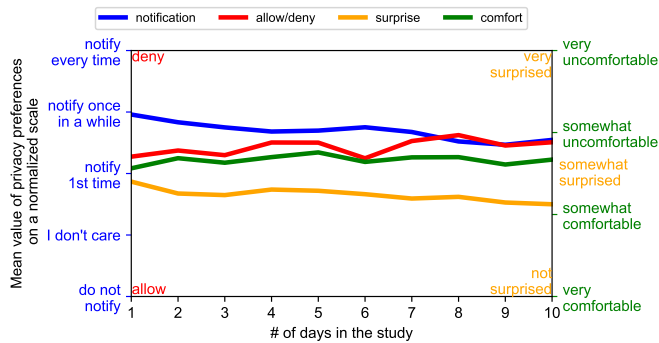
2) *Comfort Level, Surprise Level and Notification Preference:* Here we explore how the different contextual attributes considered in our study seem to influence participants’ comfort level, surprise level, and notification preferences. As those responses are not binary nor linear, GLMM is not suitable due to its inability to model ordinal dependent variables. Instead, we constructed three cumulative link mixed models (CLMM) fitted with the adaptive Gauss-Hermite quadrature approximation with 10 quadrature points using the R package `ordinal` [94] for each of the dependent variable, adopting

the same set of independent variables and random effect, as is the case with allow/deny decisions described in Section IV-D1.

Similarly to the case with allow/deny decisions, purpose remains the attribute exercising 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. We plotted the mean values of the 4 response variables in Fig. 4. We represented “allow” as 0 and “deny” as 1, and convert the other three response variables in Likert scales to integers. Values are normalized on a scale of [0, 1] for illustration. Fig. 4 shows the trends of these response variables as the study progressed. Participants’ surprise level seemed to go down, probably due to them coming across previously encountered scenarios. Participants tended to deny more, and their desire to be notified also appeared to become more selective.

### E. Attitude Change between Start and End of the Study

1) *Less Knowledge and More Concerned:* In our pre-study and post-study surveys, we asked participants the same questions about their knowledge of, comfort level with, notification preference for facial recognition and how likely or unlikely they would be to avoid places where this technology is deployed. Comparing answers from both surveys, we find that one third of the participants reported being less knowledgeable about facial recognition deployments at the end of the study



**Fig. 4:** Evolution of 4 response variables (normalized) during the course of the 10-day study. The respective scales for these variables are colored-coded along the vertical axes.

than they had at the start. 61% (N=75) of participants also felt more concerned than before. 60 out of those 75 (80%) participants attributed their heightened concern to increased awareness resulting from participation in the study. They did not know facial recognition could be used for so many different purposes, at such a diverse set of venues, and with this level of sophistication. One participant (P68) wrote, “Some of the scenarios and growth of the technology you mentioned, I had never considered. Freaked me out.” 17% emphasized the privacy issue: “It is pervasive without any notification. It is a violation of privacy” (P123). Some (15%) were concerned by the lack of notice or consent. For example, P40 elaborated on their thoughts: “It feels really easy for someone just to take data from you without your knowledge/consent. There’s a lot of directions the use of the info could go in and some of them (related to the govt) are scary.” The second part of the argument was also echoed by 12% who were worried about implications like how the data is shared, what could be inferred from it, and what they perceived as potential abuse.

2) *Reasons for not Feeling More Concerned:* At the end of the study, a total of 48 participants reported concern levels that were equal to or lower than those reported at the start. Out of those, 13 participants (27%) claimed that they were already familiar with facial recognition. For instance, participant P28 who falls in this group said “This study has taught me nothing I already didn’t know about the technology”. Others (23%) were not bothered by facial recognition and did not find the scenarios presented to them alarming: “its[sic] not the all seeing eye i once thought” (P93). A handful of participants (21%) expressed a sense of resignation, describing the technology as “ubiquitous and inevitable” (P55), and “out of my control” (P114). The fact that facial recognition happens in the physical world other than online also exacerbates this feeling “I’m limited as to what places I visit & don’t have much of choice” (P43). One participant summarized it: “I understand how we work in todays[sic] technology, privacy will now be the new luxury” (P78). Seven participants (15%) did not believe the data practices presented to them were real: “The study used hypothetical situations so it had no impact on my opinion” (P17). Others (13%) who learned the benefits of facial recognition seemed to become more accepting. “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.”, stated one participant (P106). Four people (8%) were already very concerned before and their concerns remained unchanged. One participant commented, “I was already very concerned and this study didn’t do anything to make me feel better” (P26).

3) *Dynamic 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, 55.3% of the participants ended up with different preferences by the time they completed the study regarding whether and how often they wanted to be notified about facial recognition deployments. Even though people felt more concerned in general, half of them wanted to be notified less often. 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 the previous section IV-D2, as well as the descending blue line in Fig. 4.

One possible explanation is that people became more sophisticated over time, developing a better appreciation for the broad deployment of these scenarios and fearing the privacy fatigue that would result from receiving a large number of 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.

## V. EXPLORING THE DEVELOPMENT OF PREDICTIVE MODELS

Under regulations such as GDPR data subjects are supposed to be notified and agree to having their footage captured by video analytics software at or before the point of collection. Recent technical advances introduced in prior work by Das et. al [9] open the door to scenarios where a user, with a “privacy assistant” app running on their smartphone, would be alerted to the presence of video analytics software and would be given the choice to opt in or out of such processing. Because of the increasingly widespread deployment of video analytics software, expecting people to manually opt in or out of video




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

analytics each time they encounter such functionality is likely to entail an unrealistically high number of decisions. In this section, we use our data to explore the feasibility of developing predictive models to assist users in managing these privacy decisions and discuss different possible deployment strategies for such models. Specifically, we focus on the development of models to predict people’s allow/deny decisions across the different types of scenarios using data collected as part of our in situ study.

#### A. Feature Selection and Clustering

As discussed in the Section IV-D1, purpose appears to be the most significant attribute when modeling people’s allow and deny decisions. Accordingly, we develop models that use purpose as feature — it is likely that more complex models could be developed with possibly even better results.

As prior work showed promising results of clustering like-minded users in the mobile app permission space [33], [37], we adopted a similar approach and applied agglomerative clustering with ward linkage on the feature vectors to cluster participants. After we obtained the resulting clusters of users, we calculated the privacy profiles of each cluster using two-thirds majority vote. If more than two thirds of participants in a given cluster allow (deny) a given data practice for a particular purpose, then the cluster profile recommends allowing (denying) that practice for that particular purpose. If there is no majority decision, or the number of data points in the cluster for the particular practice and purpose is too small, the cluster profile does not recommend allowing/denying the practice for the given purpose (i.e., no recommendation).

#### B. Predictive Power of Cluster Profiles

We want to evaluate how well the cluster profiles generated could help predict people’s allow/deny decisions for incoming users not present in the clusters. We first randomly select 90% of the participants to build clusters as described in the previous section, and use the remaining 10% of participants to evaluate the predictive power of the clusters by calculating the following two metrics *accuracy* and *efficiency*. *Accuracy* is defined as the percentage of time the prediction of a cluster profile (when such prediction is available) matches the actual allow/deny decisions made by users assigned to that profile. We define *efficiency* as the percentage of allow/deny decisions made by a user for which the assigned cluster of the user offers a prediction (or recommendation). In other words, if for every allow/deny decision a user needs to make, the cluster to

which the user is assigned offers a prediction, efficiency is 100 percent — theoretically the user does not need to manually make any decision, though the accuracy of the predictions could be less than 100 percent, as some predictions could be erroneous.





Fig. 6: Accuracy and efficiency of models plotted against the number of clusters used to build them.

We repeated 10 times the process of generating clusters from randomly drawing 90% of participants, and of evaluating the predictive power of these clusters using allow/deny decisions of the remaining 10% of participants. Average *accuracy* and *efficiency* results are shown in Fig. 6. As can be seen, there is a substantial increase in both accuracy and efficiency when we move from a global one-size-fits-all profile (single cluster) to models with two or more clusters. We can observe the trade-off between efficacy and accuracy as the number of clusters grows. Accuracy increases with the number of clusters, as these clusters become more targeted. Yet, efficiency decreases given that, as the number of clusters increases, the size (or population) of each cluster decreases, eventually making it more difficult to generate predictions as some entries have too few data points to obtain majority voting. The results for six clusters seem to provide the highest harmonic mean of accuracy and efficiency. It is worth noting that a model with 6 clusters achieves an efficiency of 93.9%, namely the clusters are able to predict 93.9% of the allow/deny decisions our participants had to make with an accuracy of 88.9%. It is likely that with additional data, more complex models, taking into account additional features beyond just purpose, could achieve even greater predictive power.

#### C. Example of Cluster Profiles

As shown in Fig. 6, one-size-fits-all models based on lumping all users in a single cluster fail to capture the rich



**Fig. 7:** Profiles associated with a 6-cluster model. Each cluster profile contains 3 columns: the left one displays the average mean value (deny=-1, allow=1), and the right column represents the cluster profile, where the blue color represents an allow decision, red means a deny, and white means no decision, either because not enough data points are available or for lack of a two-thirds majority. The middle column shows the variances, ranging from 0 to 1. The 3 numbers (D/A/T) in each entry in the the right column represent the distribution of deny (“D”) and allow (“A”) collected for members of the cluster for the corresponding purpose, with T=D+A representing the total number of decisions collected for the given purpose from members of the cluster.

and diverse responses towards facial recognition deployments captured in our study. However, models obtained by organizing participants in a small number of clusters seem to achieve much higher predictive power. Here we look at the profiles associated with a 6-cluster model, (see Fig. 7), namely the model that yielded the highest harmonic mean in the previous section, and discuss what these profiles tell us about how people report feeling towards different deployment scenarios.

As can readily be seen, participants in Cluster 1 and Cluster 5 represent polar extremes, with participants in Cluster 5 indicating they would largely respond with an “Allow” to all the deployment scenarios covered in our study, whereas participants in Cluster 1 would largely respond with a “Deny” to all these scenarios. It is worth also noting the low variances found in these two clusters for most deployment scenarios, indicating that people’s responses in these clusters tend to be particularly homogeneous. All other clusters also exhibit low variances for many scenarios, though each of these other 4 clusters has a few scenarios for which responses are less homogeneous, with each of these other 4 clusters having one or more deployment scenarios where the model is unable to make a prediction (e.g. “Rate Service (Anon)” in the case of Cluster 4). Comparing Cluster 3 with Cluster 5, we see that like in Cluster 5, participants in Cluster 3 tend to respond with an “Allow” to scenarios associated with a variety of different purposes, except when it comes to sensitive purposes like tracking attendance or evaluating work productivity. They tend to also be more reticent in the presence of facial recognition scenarios designed to support health predictions. Members of Cluster 2 exhibit significantly more conservative responses and are generally uncomfortable with a much larger set of deployment scenarios than members of Cluster 3, though they appear to be fine with the use of facial recognition to capture demographic information in support of anonymous targeted advertising scenarios (e.g. adjusting the ad shown in a store window based on demographic features of the person looking at the window [61], [65]–[67]). In comparison, members of Cluster 4 seem to exhibit somewhat different sensitivities.

While they too object to many deployment scenarios, they appear to be fine with the use of facial recognition to fight crime and to also anonymously count people.

#### D. Possible Application in the Context of Privacy Assistants

The above analysis sheds some light on how different groups of people share many privacy preferences when it comes to opting in or out of different video analytics scenarios and how these preferences vary across different groups. The analysis also suggests that it might be possible to predict many privacy decisions a user would otherwise have to manually make if given functionality to systematically opt in or out of video analytics software. While it is unlikely that people would want to fully delegate such decisions to software, as this would result in a significant loss of agency, it is easy to imagine configuring privacy assistant functionality where predictive models could be used to recommend some decisions and/or to automatically take care of otherwise tedious and repetitive decisions.

## VI. DISCUSSION

### A. Limitations

We would like to start by acknowledging some limitations of our study. Our sample population skews young and more educated, which could have induced bias in our results. Since our participants were recruited only from a mid-sized city in the United States, we do not claim that our results are fully representative of the general population. 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 always possible that some of the data collected is not fully representative of participants’ actual preferences, concerns and behaviors. We also acknowledge that more sophisticated predictive models could be built with even better performance, but believe that our results are sufficient to demonstrate the potential benefits of using such models.

### *B. 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 study participants and those for which, if given a chance, people 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 don’t, 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 comfort level and their desire to be notified for 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 need to give people the ability to choose whether or not to allow the collection and processing of their data, as people express diverse levels of comfort with these scenarios with many not feeling comfortable with at least some of them. These findings are also consistent with new requirements introduced by regulations such as GDPR or CCPA.

### *C. 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, how long the data is retained, but also where the user is at the time of collection and more. As such, our findings are another illustration of contextual integrity principles introduced by Nissenbaum [8]. The importance of purpose information identified in our study (i.e., for what purpose video analytics is being applied) is also consistent with results reported in earlier publications. This includes earlier work conducted by Lin et al. [33] and Smullen et al. [43] in their studies of people’s privacy preferences when it comes to configuring mobile app permission setting. This also includes prior work by Emami-Naeini et al. [24], looking at people’s privacy preferences across a number of IoT scenarios. In contrast to these earlier studies, our work did not rely on responses to online vignettes. Instead, in our work, people’s privacy attitudes were collected in situ in the context of their regular everyday activities. And obviously, our study takes a more systematic approach to exploring a range of video analytics scenarios, varying the type of analysis being carried out, the purpose for which the analysis is conducted, whether information is being shared with other entities, the venue

where video analytics is deployed; those factors all have an impact on people’s privacy attitudes.

### *D. Implications for the Design of Privacy Assistants*

Das et al. have introduced a privacy infrastructure for the Internet of Things, where users rely on “privacy assistant” mobile apps to discover nearby IoT resources such as cameras running video analytics software [45]. Using these privacy assistants, users can access opt-in or opt-out functionality made available by these IoT resources to indicate whether they agree or not to the collection and processing of their data. Das et al. have also reported customizing this infrastructure specifically for video analytics scenarios, including accommodating situations where some people provide consent and others do not, with denaturing software being applied in near real-time to obfuscate the faces of those people who did not provide consent [9]. Such functionality would effectively enable people to exercise those privacy rights granted to them under regulations such as GDPR or CCPA when it comes to video analytics scenarios. However, given the growing deployment of cameras, taking advantage of such functionality would still be hampered by the number of decisions a typical person would have to make each day when passing within range of cameras. Our work on developing models of people’s privacy preferences when it comes to granting permissions or not to the collection and processing of data by video analytics software under different scenarios, as presented in Section V, opens the door to the development of technology that could assist users with such decisions. Specifically, we demonstrated the feasibility of using simple clustering techniques to develop privacy profiles that could be used to accurately predict the majority of a user’s privacy decisions. These results are in line with prior research on helping users configure mobile app privacy permissions on their smartphones [33], [37], [43]. In the case of video analytics, such functionality could recommend to users privacy settings that could be repeatedly used to make allow/deny decisions on their behalf each time they encounter facial recognition with available opt-in/opt-out choices, saving them the effort to make these decisions manually each time. A recent study by Colnago et al. suggests that many people would see benefits to having this type of functionality, though not all of them would want to configure it the same way. In Colnago’s study some people express a desire to actually delegate many decisions while others indicate they would value the recommendations but would want to more closely control each decision [47]. Further work is needed to identify a simple set of configuration options to accommodate these different sensitivities.

### *E. Evolving Notification Preferences*

In our study, we observed that participants’ notification preferences evolved over time with people generally opting for somewhat 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



- [22] H. Nissenbaum, *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- [23] N. Apthorpe, Y. Shvartzshnaider, A. Mathur, D. Reisman, and N. Feamster, “Discovering smart home internet of things privacy norms using contextual integrity,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 2, no. 2, p. 59, 2018.
- [24] 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 *Thirteenth Symposium on Usable Privacy and Security (SOUPS) 2017*, 2017, pp. 399–412.
- [25] A. Acquisti and R. Gross, “Imagined communities: Awareness, information sharing, and privacy on the facebook,” in *International workshop on privacy enhancing technologies*. Springer, 2006, pp. 36–58.
- [26] J. M. Hektner, J. A. Schmidt, and M. Csikszentmihalyi, *Experience sampling method: Measuring the quality of everyday life*. Sage, 2007.
- [27] L. F. Barrett and D. J. Barrett, “An introduction to computerized experience sampling in psychology,” *Social Science Computer Review*, vol. 19, no. 2, pp. 175–185, 2001.
- [28] S. Consolvo and M. Walker, “Using the experience sampling method to evaluate ubicomp applications,” *IEEE pervasive computing*, vol. 2, no. 2, pp. 24–31, 2003.
- [29] 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 crowdsourcing,” in *Proceedings of the 2012 ACM conference on ubiquitous computing (Ubicomp ’12)*. ACM, 2012, pp. 501–510.
- [30] K. Martin and K. Shilton, “Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices,” *The Information Society*, vol. 32, no. 3, pp. 200–216, 2016.
- [31] 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)*, 2016, pp. 407–412.
- [32] 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)*. IEEE, 2017, pp. 276–285.
- [33] 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 *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. Menlo Park, CA: USENIX Association, 2014, pp. 199–212. [Online]. Available: <https://www.usenix.org/conference/soups2014/proceedings/presentation/lin>
- [34] 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*. ACM, 2015, pp. 807–816.
- [35] 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*, ser. WWW ’14. New York, NY, USA: ACM, 2014, pp. 201–212. [Online]. Available: <http://doi.acm.org/10.1145/2566486.2568035>
- [36] 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 (SP)*. IEEE, 2017, pp. 1077–1093.
- [37] 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 2016)*. Denver, CO: USENIX Association, 2016, pp. 27–41. [Online]. Available: <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu>
- [38] 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*, 2018, pp. 1–13.
- [39] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor, “A design space for effective privacy notices,” in *Eleventh Symposium On Usable Privacy and Security (SOUPS) 2015*, 2015, pp. 1–17.
- [40] A. Acquisti and J. Grossklags, “Privacy and rationality in individual decision making,” *IEEE security & privacy*, vol. 3, no. 1, pp. 26–33, 2005.
- [41] H. Almuhammedi, 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*, 2015, pp. 787–796.
- [42] 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)*, 2015, pp. 296–304.
- [43] D. Smullen, Y. Feng, S. Zhang, and N. Sadeh, “The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences,” *Proceedings on Privacy Enhancing Technologies*, vol. 2020, pp. 195–215, 01 2020.
- [44] 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*. IEEE, 2013, pp. 848–855.
- [45] 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*, vol. 17, no. 3, pp. 35–46, 2018.
- [46] 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*, 2004, pp. 177–189.
- [47] 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 Systems*, ser. CHI ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1–13. [Online]. Available: <https://doi.org/10.1145/3313831.3376389>
- [48] J. Wang, B. Amos, A. Das, P. Pillai, N. Sadeh, and M. Satyanarayanan, “Enabling live video analytics with a scalable and privacy-aware framework,” *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 14, no. 3s, pp. 1–24, 2018.
- [49] 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*, vol. 19, no. 3, pp. 86–89, 2016.
- [50] I. Kramer, C. J. Simons, J. A. Hartmann, C. Menne-Lothmann, W. Viechtbauer, F. Peeters, K. Schruers, A. L. van Bommel, I. Myin-Germeys, P. Delespaul *et al.*, “A therapeutic application of the experience sampling method in the treatment of depression: a randomized controlled trial,” *World Psychiatry*, vol. 13, no. 1, pp. 68–77, 2014.
- [51] 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*, vol. 102, no. 6, p. 1318, 2012.
- [52] L. L. Carstensen, B. Turan, S. Scheibe, N. Ram, H. Ersner-Hershfield, G. R. Samanez-Larkin, K. P. Brooks, and J. R. Nesselroade, “Emotional experience improves with age: evidence based on over 10 years of experience sampling,” *Psychology and aging*, vol. 26, no. 1, p. 21, 2011.
- [53] 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*, 2018, pp. 1–13.
- [54] 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*, 2014, pp. 91–100.
- [55] N. Van Berkel, D. Ferreira, and V. Kostakos, “The experience sampling method on mobile devices,” *ACM Computing Surveys (CSUR)*, vol. 50, no. 6, pp. 1–40, 2017.
- [56] V. Pejovic, N. Lathia, C. Mascolo, and M. Musolesi, “Mobile-based experience sampling for behaviour research,” in *Emotions and personality in personalized services*. Springer, 2016, pp. 141–161.
- [57] T. Revell. (2017, January) Computer vision algorithms pick out petty crime in cctv footage. [Online]. Available: <https://www.newscientist.com/article/2116970-computer-vision-algorithms-pick-out-petty-crime-in-cctv-footage/>
- [58] T. Johnson. (2018, may) Shoplifters meet their match as retailers deploy facial recognition cameras. [Online]. Available: <https://www.mcclatchydc.com/news/nation-world/national/article211455924.html>
- [59] B. Conarck. (2019, January) Florida court: Prosecutors had no obligation to turn over facial recognition evidence. [Online].

- Available: <https://www.jacksonville.com/news/20190123/florida-court-prosecutors-had-no-obligation-to-turn-over-facial-recognition-evidence>
- [60] (2018, April) Chinese man caught by facial recognition at pop concert. BBC News. [Online]. Available: <https://www.bbc.com/news/world-asia-china-43751276>
- [61] C. Frey. (2016, March) Revealed: how facial recognition has invaded shops – and your privacy. [Online]. Available: <https://www.theguardian.com/cities/2016/mar/03/revealed-facial-recognition-software-infiltrating-cities-saks-toronto>
- [62] (2012, November) Nec unveils facial-recognition system to identify shoppers. PCMag. [Online]. Available: <https://www.pcmag.com/archive/nec-unveils-facial-recognition-system-to-identify-shoppers-305015>
- [63] D. Rosen. (2013, January) Disney is spying on you! [Online]. Available: [https://www.salon.com/test/2013/01/17/disney\\_is\\_spying\\_on\\_you/](https://www.salon.com/test/2013/01/17/disney_is_spying_on_you/)
- [64] D. Murph. (2011, June) Scenetap app analyzes pubs and clubs in real-time, probably won't score you a jersey shore cameo. [Online]. Available: <https://www.engadget.com/2011/06/12/scenetap-app-analyzes-pubs-and-clubs-in-real-time-probably-won/>
- [65] D. Etherington. (2016, December) Baidu and kfc's new smart restaurant suggests what to order based on your face. [Online]. Available: <https://techcrunch.com/2016/12/23/baidu-and-kfcs-new-smart-restaurant-suggests-what-to-order-based-on-your-face/>
- [66] (2013, November) Tesco's plan to tailor adverts via facial recognition stokes privacy fears. Tesco Press Association. [Online]. Available: <https://www.theguardian.com/business/2013/nov/03/privacy-tesco-scan-customers-faces>
- [67] B. Snyder. (2015, May) This beer ad only works when women pass by. [Online]. Available: <https://fortune.com/2015/05/21/astra-beer-ad/>
- [68] (2012, November) Mannequins collect data on shoppers via facial-recognition software. Bloomberg News. [Online]. Available: [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](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)
- [69] L. Clark. (2012, November) Mannequins are spying on shoppers for market analysis. [Online]. Available: <https://www.wired.co.uk/article/mannequin-spies-on-customers>
- [70] E. J. SCHULTZ. (2015, May) Facial-recognition lets marketers gauge consumers' real responses to ads. [Online]. Available: <https://adage.com/article/digital/facial-recognition-lets-marketers-gauge-real-responses/298635>
- [71] D. Burrows. (2017, May) Facial expressions show mars the adverts that will drive sales. [Online]. Available: <https://www.foodnavigator.com/Article/2017/03/23/Facial-expressions-show-Mars-the-adverts-that-will-drive-sales>
- [72] B. Logan. (2014, October) Pay-per-laugh: the comedy club that charges punters having fun. [Online]. Available: <https://www.theguardian.com/stage/2014/oct/14/standup-comedy-pay-per-laugh-charge-barcelona>
- [73] S. Lepitak. (2019, March) Disney's dumbo and accenture interactive collaborate for the movie poster of the future. [Online]. Available: <https://www.thedrum.com/news/2019/03/10/disneys-dumbo-and-accenture-interactive-collaborate-the-movie-poster-the-future>
- [74] J. Whitley. (2018, November) How facial recognition technology is being used, from police to a soccer museum. [Online]. Available: <https://www.wfaa.com/article/features/originals/how-facial-recognition-technology-is-being-used-from-police-to-a-soccer-museum/287-618278039>
- [75] M. Ehrenkranz. (2017, December) Burger joint teams up with surveillance giant to scan your face for loyalty points. [Online]. Available: <https://gizmodo.com/burger-joint-teams-up-with-surveillance-giant-to-scan-y-1821498988>
- [76] New biometric identification tools used in theme parks. NEC. [Online]. Available: <https://www.nec.com/en/global/about/mitatv/03/3.html>
- [77] R. Bachman. (2017, June) Your gym's tech wants to know you better. [Online]. Available: <https://www.wsj.com/articles/your-gyms-tech-wants-to-know-you-better-1497281915>
- [78] E. Silverstein. (2019, October) New konami casino facial recognition technology could rival reward cards. [Online]. Available: <https://www.casino.org/news/new-konami-casino-facial-recognition-technology-could-rival-reward-cards/>
- [79] S. F. Gale. (2013, March) Employers turn to biometric technology to track attendance. [Online]. Available: <https://www.workforce.com/news/employers-turn-to-biometric-technology-to-track-attendance>
- [80] S. P. Bailey. (2015, July) Skipping church? facial recognition software could be tracking you. [Online]. Available: [washingtonpost.com/news/acts-of-faith/wp/2015/07/24/skipping-church-facial-recognition-software-could-be-tracking-you/](http://washingtonpost.com/news/acts-of-faith/wp/2015/07/24/skipping-church-facial-recognition-software-could-be-tracking-you/)
- [81] D. Levine. (2019, June) What your face may tell lenders about whether you're creditworthy. [Online]. Available: <https://www.wsj.com/articles/what-your-face-may-tell-lenders-about-whether-youre-creditworthy-11560218700?mod=searchresults&page=1&pos=5&mod=djelmAIPro>
- [82] M. Allen. (2018, July) Health insurers are vacuuming up details about you — and it could raise your rates. [Online]. Available: <https://www.propublica.org/article/health-insurers-are-vacuuming-up-details-about-you-and-it-could-raise-your-rates>
- [83] E. Rader. (2019, February) Most americans don't realize what companies can predict from their data. [Online]. Available: <https://bigthink.com/technology-innovation/most-americans-dont-realize-what-companies-can-predict-from-their-data-2629911919>
- [84] B. J. Davidson. (2019, November) How your business can benefit from facial recognition technology. [Online]. Available: <https://percentotech.com/how-your-business-can-benefit-from-facial-recognition-technology/>
- [85] T. X. Ingrid Fadelli. (2019, December) Analyzing spoken language and 3-d facial expressions to measure depression severity. [Online]. Available: <https://techxplore.com/news/2018-11-spoken-language-d-facial-depression.html>
- [86] S. Krouse. (2019, July) The new ways your boss is spying on you. [Online]. Available: <https://www.wsj.com/articles/the-new-ways-your-boss-is-spying-on-you-11563528604>
- [87] D. Levine. (2017, October) What high-tech tools are available to fight depression? [Online]. Available: <https://health.usnews.com/health-care/patient-advice/articles/2017-10-06/what-high-tech-tools-are-available-to-fight-depression>
- [88] (2018, May) Facial recognition: School id checks lead to gdpr fine. Augmented Mental Health. [Online]. Available: <https://www.augmentedmentalhealth.com/blog/augmented-mental-health-revolutionary-mental-health-care-using-emotion-recognition>
- [89] 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*, vol. 25, no. 1, pp. 60–64, 2019. [Online]. Available: <https://doi.org/10.1038/s41591-018-0279-0>
- [90] N. K. Malhotra, S. S. Kim, and J. Agarwal, "Internet users' information privacy concerns (iupc): The construct, the scale, and a causal model," *Information systems research*, vol. 15, no. 4, pp. 336–355, 2004.
- [91] 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*, vol. 4, no. 1, pp. 53–78, 2003.
- [92] Y. Rogers and P. Marshall, "Research in the wild," *Synthesis Lectures on Human-Centered Informatics*, vol. 10, no. 3, pp. i–97, 2017.
- [93] D. Bates, M. Mächler, B. Bolker, and S. Walker, "Fitting linear mixed-effects models using lme4," *Journal of Statistical Software*, vol. 67, no. 1, pp. 1–48, 2015.
- [94] R. H. B. Christensen, "ordinal—regression models for ordinal data," 2019, r package version 2019.12-10. <https://CRAN.R-project.org/package=ordinal>.
- [95] 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*, ser. UbiComp '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 1542–1547. [Online]. Available: <https://doi.org/10.1145/2968219.2968546>
- [96] 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)*, vol. 50, no. 3, pp. 1–41, 2017.



APPENDIX A  
GLMM TABLE FOR ALLOW/DENY

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
eatory	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 IV:** 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 <b>deter crime</b> . (This footage can be shared with <b>law enforcement</b> .) Assume that you are captured by such a camera, and the <b>raw footage is kept for 30 days</b> .
Petty Crime	Some places like %s have started to deploy video surveillance cameras to <b>deter crime</b> . These cameras are equipped with software that can automatically <b>detect and record petty crime</b> (e.g. pickpocketing, car break-ins, breaking store windows). When a suspicious scene is believed to have been detected, it is <b>recorded for further analysis (possibly including facial recognition) and kept for 30 days</b> . <b>Otherwise the data is immediately discarded</b> . Assume that you are captured by such a camera.
Known Criminal Detection	Some places like %s have started to deploy video surveillance cameras with <b>facial recognition</b> software. This software can <b>identify and track known shoplifters, criminals, and bad actors</b> . Assume that %s engages in this practice, and the <b>raw footage is discarded immediately, with the analysis results being kept for 30 days</b> .
Count people	Some places like %s have started to deploy video surveillance cameras with <b>anonymous face detection</b> software. This software can estimate the number of customers in the facility in order to <b>optimize operation</b> , such as personnel allocation. Assume that %s engages in this practice and it is <b>unclear for how long all the data (raw footage and analysis results) is kept</b> .
Jump Line	Some places like %s have started to deploy video surveillance cameras with <b>facial recognition</b> software. This software can identify patrons in line and push individualized offers to <b>skip the wait-line for a fee</b> . This software can also record your presence and <b>who you are with</b> . Assume that %s engages in this practice and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Targeted Ads(Anon)	Some places like %s have started to deploy video surveillance cameras with <b>anonymous face detection</b> software. This software can estimate customers' race and ethnicity in order to offer <b>tailored deals and coupons</b> . Assume that %s engages in this practice and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Targeted Ads(IDed)	Some places like %s have started to deploy video surveillance cameras with <b>facial recognition</b> software. This software can match detected faces against individual customer profiles in order to offer <b>tailored deals and coupons</b> . Assume that %s engages in this practice and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Sentiment Ads(Anon)	Some places like %s have started to deploy video surveillance cameras with <b>anonymous face detection</b> and <b>emotion analysis</b> software. This software can estimate customers' age, gender and ethnicity, and analyze their reactions to items displayed. This software is used to generate <b>tailored deals and coupons</b> for different demographic groups. Assume that %s engages in this practice and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Sentiment Ads(IDed)	Some places like %s have started to deploy video surveillance cameras with <b>facial recognition</b> and <b>emotion analysis</b> software. This software recognizes people, and analyzes their reactions to items displayed. Then the software matches detected faces against individual customer profiles to send <b>tailored deals and coupons</b> to their phones. Assume that %s engages in this practice and the <b>raw footage is kept for 30 days</b> , and it is <b>unclear for how long the analysis results are kept</b> .
Rate Service	Some places like %s have started to deploy video surveillance cameras with <b>anonymous emotion analysis</b> software. This software can <b>gauge customer satisfaction</b> with the service provided by its employees. They can use the results for <b>employee evaluation and training purposes</b> . Assume that %s engages in this practice and it is <b>unclear for how long all the data (raw footage and analysis results) is kept</b> .
Rate Engagement	Some places like %s have started to deploy video surveillance cameras with <b>facial recognition</b> and <b>emotion analysis</b> software. This software can identify each patron, and <b>measure their engagement</b> at the facility. This software can be used to record your presence and also identify <b>who you are with</b> . Assume that %s engages in this practice and the <b>raw footage is kept for 30 days</b> , and it is <b>unclear for how long the analysis results are kept</b> .
Face as ID	Some stores have started to deploy video surveillance cameras with <b>facial recognition</b> software. This software can identify faces of customers to <b>replace membership cards</b> at checkout. Assume that %s engages in this practice, and the <b>raw footage is discarded immediately</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Track Attendance	Some companies have started to deploy video surveillance cameras with <b>facial recognition</b> software. This software can track the work time <b>attendance of its employees</b> . This software can also identify how long you participate in different activities and <b>who you hang out with</b> . Assume that your workplace engages in this practice, and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Word Productivity	Some companies have started to deploy video surveillance cameras with <b>emotion analysis</b> and <b>facial recognition</b> software. This software can detect the mood of its employees, and <b>predict their productivity</b> . Assume that your workplace engages in this practice, and it is <b>unclear for how long all the data (raw footage and analysis results) is kept</b> .
Health Predictions	Some eatery chains like %s have started to deploy video surveillance cameras with <b>emotion analysis</b> and <b>facial recognition</b> 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 <b>likelihood of developing depression, diabetes, and obesity</b> , which can impact your <b>health insurance premium</b> . Assume that %s engages in this practice, and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .
Medical Predictions	Some medical facilities have started to deploy video surveillance cameras with <b>emotion analysis</b> and <b>facial recognition</b> software. This software can automatically detect some physical and mental health problems. This information can be shared with health insurance providers, and impact your <b>health insurance premium</b> . Assume that %s engages in this practice, and the <b>raw footage is kept for 30 days</b> . Assume also that it is <b>unclear for how long the analysis results are kept</b> .

TABLE V: Scenarios text shown to participants