15th IEEE Conference on Pervasive Computing and Communications, Kona, HI, 2017 (to appear)

# Privacy Preference Modeling and Prediction in a Simulated Campuswide IoT Environment

Hosub Lee and Alfred Kobsa
Donald Bren School of Information and Computer Sciences
University of California, Irvine
Irvine, USA
{hosubl, kobsa}@uci.edu

***Abstract*—With the advent of the Internet of Things (IoT), users are more likely to have privacy concerns since their personal information could be collected, analyzed, and utilized without notice by the networked IoT devices and services. Users may want to control all such activities by explicitly expressing their privacy preferences. However, it is becoming increasingly difficult for users to do so, not only because of the cognitive burden of continuously making privacy decisions for IoT services, but also because IoT devices have no, or only very restricted, user interfaces. Intelligent software helping users make better privacy decisions will be an important component of privacy-preserving IoT environments. In order to construct such a component, we aim to verify whether it will be possible to computationally model and predict users' privacy preferences in IoT. To that end, we survey 172 participants in a simulated campuswide IoT environment about their privacy preferences regarding hypothetical personal information tracking scenarios. Then, we cluster the scenarios based on the survey responses, arriving at four clusters with distinct associated privacy preferences. Based on the clustering results, we uncover contextual factors that induce privacy violations in IoT. Finally, we build machine learning models to predict users' privacy decisions, using both contextual information and the corresponding cluster membership as training data. The final trained model shows 77% accuracy in predicting users' decisions whether or not to allow the respective IoT scenario.**

***Keywords — IoT, privacy, preference modeling, experience sampling, Google Glass, K-modes clustering, conditional inference tree***

## I. INTRODUCTION

The Internet of Things (IoT) is a networked computing environment consisting of various types of physical objects (i.e., things) that are able to collect and exchange data over a network with minimum user intervention [1-4]. Sensors and devices in IoT can easily collect data about our personal characteristics and behavior. For individuals, there are many advantages of incorporating IoT into their lives. These advantages can come in various forms such as safety, financial benefits, social relationships, convenience, and health. For instance, IoT-based home automation systems can monitor users' behavior via motion sensors, Wi-Fi signals or facial recognition technology, to identify their presence in their homes and automatically control room temperature or lighting. IoT technologies can be embedded into virtually every situation that users encounter in their daily lives. IoT could improve users' overall quality of life if it works appropriately, but compromise their privacy if it does not. This is mainly because IoT devices can collectively gather massive amounts of personal information, without informing users let alone asking for their permission [5-8]. For these reasons, it is understood that safeguarding users' privacy is a big challenge to the widespread adoption of IoT products and services.

To protect users' privacy in ubiquitous computing environments, service providers increasingly ask them to make privacy decisions (e.g., grant or deny smartphone apps permission to access the user's location). However, users are increasingly unable to make these decisions due to limits in their available time, motivation, and their cognitive decision-making abilities [9, 10]. Therefore, many researchers proposed various mechanisms to predict users' privacy decisions via machine learning models trained on a subset of users' prior privacy behavior [11-18]. Software agents can then use these machine learning models to give users personalized privacy recommendations, thereby assisting them to better control their privacy. This kind of technology is going to become more important in IoT environments, not only because users need to make decisions much more frequently for pervasive IoT services, but also because of the lack of user interfaces for specifying privacy preferences to the services. It is therefore necessary to investigate whether it will be possible to model users' privacy preferences in such IoT environments as well as to predict their future privacy decisions.

In this vein, we conducted a series of studies consisting of privacy preference collection, privacy preference analysis, and privacy preference prediction. First, we collected people's decisions and opinions regarding their privacy in diverse privacy-invasive scenarios in simulated IoT environments, through the experience sampling method (ESM). We developed an app for Google Glass that can dynamically display a description of an IoT scenario related to the current location of the user. We then recruited participants and asked them to walk around a university campus while wearing Google Glass. They were instructed to answer survey questions whenever they received notifications from Google Glass describing an IoT scenario related to their current location. We utilized Google Glass to give participants an immersive virtual experience of being monitored by IoT devices, to ensure that our research would be as situated as is currently possible. As a result, we collected 33,090 valid survey responses from 172 participants over a period of three months.

Next, we clustered the collected responses using the K-modes clustering algorithm to quantitatively assess the impact of different contextual factors (e.g., what is monitored, by whom, etc.) on participants' desire for notification and control, and on their subjective evaluation of potential privacy risks. We found four distinct clusters in terms of their stated privacy preferences, and explored relationships between IoT contexts and user attitudes by comparing the survey responses in each cluster. Through this analysis, we can now understand how contextual factors influence people's behaviors and perceptions toward their privacy in IoT environments.

Finally, we tried to predict participants' privacy decisions by learning conditional inference trees, using the gathered survey responses as training data. We utilized the above-mentioned contextual factors as well as clustering results as variables (or features) for predicting how participants will make privacy decisions in the presented scenarios. The final trained model has a 77% 10-fold cross validation accuracy in predicting whether or not participants will allow personal information monitoring in a given IoT scenario.

## II. Related Work

In this section, we present a literature review of user privacy in various computing environments including IoT. To begin with, we survey previous studies aimed at understanding the causes and effects of users' privacy behavior in mobile/ ubiquitous computing environments. Next, we focus on techniques for predicting users' privacy decisions or preferences in such environments.

### A. *Privacy Preference Analysis*

In order to design privacy-preserving applications and services, we first need to understand the extent to which users' privacy preferences are shaped by the context in which the usage of these applications or services takes place. In this regard, many researchers have investigated several contextual factors that could influence users' privacy concerns in diverse application scenarios.

Lederer et al. [19] conducted a scenario-based online survey to evaluate the relative importance of two factors, requester and situation, in determining users' privacy preferences in ubiquitous computing environments. They presented a set of personal information disclosure scenarios to participants, and then collected participants' reactions to these scenarios. Specifically, users were asked to specify the preferred degree of disclosure of their personal information at three levels (i.e., full disclosure, vague disclosure, non-disclosure). By quantitatively analyzing the responses, the authors found that the identity of information requester (4 possible values: spouse, employer, stranger, merchant) is more significant than the current situation (2 possible values: working lunch and social evening) in making a privacy decision. However, there is no guarantee that this finding can also be applied to IoT contexts since the situation was too coarsely defined in this study.

Choe et al. [20] confirmed that users are less willing to share self-appearance, intimacy behavior, cooking or eating, media use, and oral expressions at home when various sensors are installed. They also conducted an anonymous online survey to collect personal behavior that people usually exhibit at home but would not want to be monitored. The authors concluded that designers and developers of in-home sensing systems should be careful not to monitor such private behaviors. Although this work gives useful insights into important contextual factors like location, the findings are restricted to the specific place investigated, namely people's homes.

Benisch et al. [21] performed a user study in order to identify contextual factors that influence users' willingness to share their location with others. Using a web-based online survey, they collected detailed preferences from 27 subjects for three weeks. Regarding the actual locations that each participant visited that day, the participant was asked afterwards to decide whether or not to share the locations with her/his acquaintances (e.g., friend and family). Participants also specified the preferred time spans for these location-sharing activities. By statistically analyzing the collected preferences, the authors discovered several contextual factors that significantly impact people's perception of location privacy, such as time of day, day of week, and exact location. They also found that privacy settings, comprised of these factors, make users have less privacy concerns compared to the conventional method, namely whitelists. This work also sheds light on important contextual factors like time and location, and suggests meaningful guidelines for designing mobile applications with a location-sharing functionality. However, it considered merely one of the many possible application scenarios realizable in IoT environments.

We find little research that comprehensively investigates diverse contextual factors influencing users' attitudes and preferences towards their privacy in an IoT environment. One of the aims of our work is to fill this gap.

### B. *Privacy Preference Prediction*

Substantial research efforts have been made to devise mechanisms that infer users' privacy decisions and proactively recommend privacy choices to users. Researchers claim that this kind of technology could help users alleviate the cognitive burden of privacy decision-making, thereby allowing them to make their preferred privacy choices more easily.

Sadeh et al. [12] proposed an automated mechanism in a mobile social networking application for making privacy decisions on behalf of users. To this end, the authors adopted a supervised machine learning algorithm named random forests. This algorithm was utilized to semi-automatically generate sharing policies for the current location of the users, based on their previous decisions. The authors showed that these machine-generated policies have better accuracy than the user-defined policies: 91% vs. 79% success rate in satisfying users' actual preferences, respectively. The reason is that users are generally not able to specify privacy policies consistent with their actual location-sharing behavior in the real world. However, the users' binary feedback (accept or reject) on the generated policies was relatively consistent with their actual behavior. Thus, the authors utilized user feedback as an additional input feature for training machine learning models, so as to achieve better prediction accuracy.

Fang et al. [13] presented a system that infers access control policies for personal information on online social networking services like Facebook. Similar to [12], they presented a supervised machine learning approach to learn users' privacy preferences by iteratively asking them questions regarding their sharing activities with friends. In doing so, the authors asked the users about privacy preferences that machine learning models are most uncertain about (i.e., active learning with uncertainty sampling). By using both the collected answers and personal profiles (e.g., gender, age, etc.) of the users, the authors continuously trained personalized machine learning models that can assign privileges to unlabeled friends of each user (e.g., friend A can see my photos). According to their study with 45 Facebook users, the system was effective in reducing user burden in configuring privacy settings on Facebook. In addition, the system showed 90% accuracy in predicting personal privacy policies with a small amount of labeled training data (25 out of 200 friends with privileges).

Bilogrevic et al. [15, 17] proposed a privacy-preserving information sharing platform named SPISM that semi-automatically decides whether or not to disclose different types of personal information and at what level of granularity. Like other previous research, the authors used a logistic classifier (a supervised machine learning method) to predict users' privacy decision-making. They employed contextual information (e.g., types of information requested, location, time, etc.) and past behavior as features for training a classifier, and verified that the trained classifier can make a prediction with 90% accuracy. Like [13], the authors adopted an active learning paradigm in the training procedure, thereby minimizing users' initial labeling efforts. The authors also deployed SPISM on the Android operating system so that users could be assisted in making decisions for a considerable amount of information sharing requests on mobile computing platforms. Similarly, Liu et al. [18] developed and evaluated a personalized privacy assistant (PPA) that proactively produces permission settings for Android applications on behalf of users. They first employed hierarchical clustering to categorize users into several groups based on their prior privacy attitudes (i.e., privacy profiles). Next, they built SVM classifiers to predict users' decisions for each permission request by using their privacy profiles and other available information related to such a request (e.g., app category, permission type, etc.) as input features. PPA was also designed to nudge users to make a correct privacy decision by giving them recommendations (classification results) at the operating system level. Through field experiments with 72 Android users, the authors confirmed that 78.7% of the recommendations made by PPA was accepted by the users.

All of the abovementioned works not only remind us of the importance of predicting privacy decisions in online or mobile computing environments, but also provide practical guidelines for adapting the prediction results to people's actual behaviors that can evolve over time. However, there is still a lack of research on this topic targeted at IoT environments. Here, we aim to study whether it is possible in IoT environments to model and predict users' privacy preferences through data mining and machine learning technologies.

## III. Phase I: Privacy Preference Collection

Our study is conducted in three main steps: (i) collect privacy preferences of users in simulated IoT environments, (ii) understand how the users make privacy decisions in such environments by analyzing the collected preferences, and (iii) build machine learning models to predict future privacy decisions of the users by using contextual information related to their privacy choices as training data. These steps will be described in this and the next two sections.

We adopted ESM to collect people's privacy preferences on various IoT service scenarios (mostly about monitoring of personal information). We used Google Glass, one of the representative wearable computers, for presenting the IoT scenarios to study participants because we intended to let them perceive the scenarios as realistically as possible. Specifically, we developed a Google Glass app called IoT Privacy to dynamically display scenarios based on participants' location. Participants were then asked to walk around our university campus wearing Google Glass. As participants move towards one of 130 selected locations on campus, the IoT Privacy app presents the scenario pertaining to this location. Participants then answer several questions on their preferred privacy protection in the given scenario. Our immersive spatial setup seems more suitable to collect accurate privacy preferences from participants than a traditional online survey system, since it situates them in scenarios and is therefore likely to better capture the situatedness of privacy decisions [22, 23]. In addition, location has been found to be a particularly critical component in understanding people's privacy decision-making [12, 15, 17, 21].

### A. Data Description

In order to formalize users' privacy preferences, we defined several parameters representing both contextual characteristics of IoT scenarios ("contextual parameters") and possible user reactions ("reaction parameters"). In earlier interview and online survey studies [24, 25], we had already identified five contextual parameters that have the most influence on the reaction parameters.

These five parameters define the place where the monitoring occurs (parameter "where"), the type of information being monitored ("what"), the entity that is monitoring ("who"), the reason for monitoring ("reason"), and the frequency of the monitoring ("persistence"). We also identified the most important reaction parameters that serve as proxies of people's privacy preferences, namely the desire to be notified about the monitoring (parameter "_notification") and the willingness to accept the monitoring ("_permission"). In addition, we also found it important to measure people's opinion on each monitoring activity in terms of comfort, risk, and appropriateness (parameters "_comfort", "_risk", "_appropriateness").

Tables I and II display the contextual and reaction parameters, respectively, together with their values which are all categorical or ordinal. Each scenario can be described by an expression that includes every contextual parameter together with its respective parameter value for this scenario.

TABLE I. CONTEXTUAL PARAMETERS (">" indicates the purpose)

| Parameter (id) | Values | |
|---|---|---|
| "where" ($C_1$) | 0. your place<br>1. someone else's place<br>2. semi-public space (e.g., restaurant)<br>3. public space (e.g., street) | |
| "what" ($C_2$) | 1. phoneID<br>2. phoneID>identity<br>3. location<br>4. location>presence<br>5. voice<br>6. voice>gender<br>7. voice>age<br>8. voice>identity<br>9. voice>presence<br>10. voice>mood<br>11. photo<br>12. photo>gender | 13. photo>age<br>14. photo>identity<br>15. photo>presence<br>16. photo>mood<br>17. video<br>18. video>gender<br>19. video>age<br>20. video>presence<br>21. video>mood<br>22. video>lookingAt<br>23. gaze<br>24. gaze>lookingAt |
| "who" ($C_3$) | 1. unknown<br>2. colleague/fellow<br>3. friend<br>4. own device | 5. business<br>6. employer/school<br>7. government |
| "reason" ($C_4$) | 1. safety<br>2. commercial<br>3. social | 4. convenience<br>5. health<br>6. none |
| "persistence" ($C_5$) | 0. once | 1. continuously |

TABLE II. REACTION PARAMETERS

| Parameter (id) | Values |
|---|---|
| "_notification" ($R_1$) | 1. notify me, always<br>2. notify me, just this time<br>3. don't notify me, just this time<br>4. don't notify me, always |
| "_permission" ($R_2$) | 1. allow, always<br>2. allow, just this time<br>3. reject, just this time<br>4. reject, always |
| "_comfort" ($R_3$) | 1. very uncomfortable/risky/inappropriate<br>2. uncomfortable/risky/inappropriate<br>3. somewhat uncomfortable/risky/inappropriate<br>4. neutral<br>5. somewhat comfortable/safe/appropriate<br>6. comfortable/safe/appropriate<br>7. very comfortable/safe/appropriate |
| "_risk" ($R_4$) | |
| "_appropriateness" ($R_5$) | |

### B. *Scenario Generation*

In our earlier online survey [25], we had created a broad range of 2,800 hypothetical IoT scenarios through random permutation of the values of the abovementioned five contextual parameters. This approach allowed us to diversify the range of scenarios without much time and effort. However, given that participants responded to the created scenarios at a time and location that bear no relationship to the scenarios described in the survey questions, there could have been a sense of decreased realism to the scenarios. This may have negatively influenced the quality and accuracy of their survey responses.

In the present study, we tackle this limitation by creating more realistic scenarios that are specifically related to known geographical locations, and by letting Google Glass prompt the scenarios based on the current location of the participant. To meet the former aim, our research team collaboratively created numerous scenario descriptions using Google My Maps, which lets multiple users create and update a custom Google Map. As shown in Fig. 1, we created landmarks with GPS coordinates and associated scenario descriptions containing all five contextual parameters. We aimed to make the scenarios as specific and realistic as possible by cross-validating the scenario texts with each other. Through this approach, we were able to improve the realism of the scenarios compared to our earlier work. We produced 130 IoT scenarios in total for our campus.

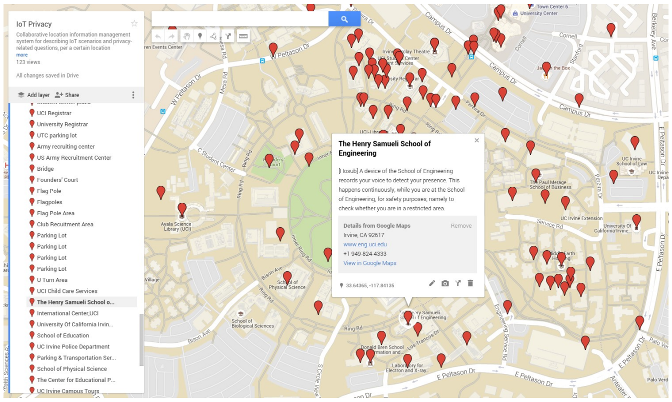

IoT Privacy

Collaborative location information management system for describing IoT scenarios and privacy-related questions, per a certain location

more

123 views

All changes saved in Drive

UCI Registrar
University Registrar
UTC parking lot
Army recruiting center
US Army Recruitment Center
Bridge
Founders' Court
Flag Pole
Flagpoles
Flag Pole Area
Club Recuitment Area
Parking Lot
Parking Lot
Parking Lot
Parking Lot
U Turn Area
UCI Child Care Services
The Henry Samueli School o...
International Center,UCI
University Of California Irvin...
School of Education
UC Irvine Police Department
Parking & Transportation Ser...
School of Physical Science
The Center for Educational P...
UC Irvine Campus Tours

The Henry Samueli School of Engineering

[Hosub] A device of the School of Engineering records your voice to detect your presence. This happens continuously, while you are at the School of Engineering, for safety purposes, namely to check whether you are in a restricted area.

Details from Google Maps

Irvine, CA 92617

www.eng.uci.edu

+1 949-824-4333

View in Google Maps

33.64365, -117.84135

Fig. 1. Collaborated Scenario Generation via Google My Maps

As Google My Maps provides functionality to export all entries into a machine-readable format such as XML, we extracted all created scenarios and relevant information as a single XML file and converted it to a more compact format in JSON (see Table III). Note that we also transformed each scenario description into a specific sequence of values of the contextual parameters (*context_param* in Table III) to make it analyzable by data mining and machine learning algorithms. For the parameters "where" ($C_1$) and "who" ($C_3$), we then replaced their written values with categorical values defined in Table I. For instance, the School of Engineering is mapped with $C_1$=3 because this place is considered as a public place.

TABLE III. SAMPLE JSON FILE

| Attribute | Value |
|---|---|
| location_name | School of Engineering |
| latitude | -117.841359 |
| longitude | 33.643657 |
| scenario | A device of the School of Engineering ($C_3$=6) records your voice to detect your presence ($C_2$=9). This happens continuously ($C_5$=1), while you are at the School of Engineering ($C_1$=3), for safety ($C_4$=1) purposes, namely to check whether you are in a restricted area. |
| scenario_id | 111 |
| context_param | {$C_1$=3, $C_2$=9, $C_3$=6, $C_4$=1, $C_5$=1} |

### C. *Location-based Scenario Display for Google Glass*

To operationalize our study, we designed and developed a novel Google Glass application named IoT Privacy that synchronizes the display of IoT scenario descriptions with the current location of survey respondents. Google Glass is a small computer that is worn like a pair of eyeglasses. Users can receive various information from its head-up display and built-in speaker, and also freely interact with their environments (i.e., hands-free user experience). Because Google Glass itself is not equipped with a GPS sensor, it needs to receive GPS information from a Bluetooth-paired smartphone.

IoT Privacy operates in the following steps:

1. The app tracks participants' location every 40 seconds with GPS data received from a Bluetooth-paired smartphone,
2. The app continuously compares the current location with the GPS coordinates of all scenario descriptions stored in a JSON-formatted database mounted in Google Glass,
3. When the current distance to a stored scenario location is below a given threshold, the app displays the description and its unique scenario ID, as shown in Fig. 2, together with a sound notification. A scenario description is displayed only once, i.e., it does not appear any more if a participant returns to the same area later.

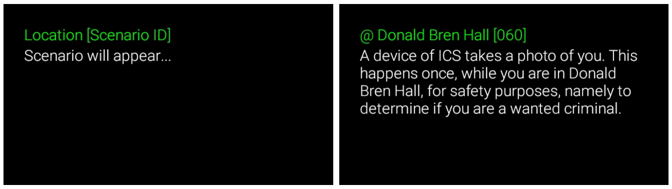

Location [Scenario ID]
Scenario will appear...

@ Donald Bren Hall [060]
A device of ICS takes a photo of you. This happens once, while you are in Donald Bren Hall, for safety purposes, namely to determine if you are a wanted criminal.

Fig. 2. IoT Privacy Screenshot

### D. *Study Procedure*

We recruited study participants on a university campus through e-mails and posted flyers. Participants needed to be at least 18 years old, be proficient in English, have a smartphone, and not have serious vision problems. They were briefed individually about the overall study procedure, basic usage of Google Glass (including Bluetooth pairing with their smartphone), and functional details of the IoT Privacy app. Participants were asked to walk around campus while wearing Google Glass. When a scenario description relating to a nearby location was displayed in Glass, participants were asked to read it, record the scenario ID, and answer the following questions:

1. Would you want to be notified about this monitoring? ($R_1$)
2. Would you want to allow this monitoring? ($R_2$)
3. How comfortable is the monitoring? ($R_3$)
4. How risky is the monitoring? ($R_4$)
5. How appropriate is the monitoring? ($R_5$)

Table II lists all available answer options. Subjects were asked to answer our questions on paper. While this seems technically unimpressive and made data collection cumbersome for the experimenters, our pilot tests showed this to be by far the best method for our participants, many of whom were first-time Google Glass users. Due to the small screen size of Glass, participants would otherwise have had to navigate through numerous pages to view each single question with all its answer options. The smartphone was also not a feasible entry device since the glare from near-permanent sunshine during the duration of the open-air experiment made the display hard to view.

Participants could carry out the experiment as long as they wished but were asked not to exceed three hours. After they returned Google Glass and the completed questionnaires, they took an exit survey and had a brief interview about their study experience. All participants received $10-60 in cash as compensation depending on how many questions they answered.

We recruited 172 participants in total over a period of three months: 106 males and 65 females (one person did not disclose her/his gender), with the majority (82%) being 18-25. Because we recruited the participants on campus, most of them have some university affiliation (107 undergraduate students, 63 graduate students, 1 postdoctoral fellow, 1 faculty member). Participants answered 39 scenario descriptions on average (std. dev: 14.72). After carefully checking our transcriptions and excluding a few invalid responses (e.g., answer number out of range), we wound up with a total of 33,090 privacy preferences for 6,618 IoT scenarios.

## IV. Phase II: Privacy Preference Analysis

We first analyze the collected privacy preferences to understand how the contextual factors impact people's reactions toward information monitoring activities in IoT environments. Specifically, we utilize the K-modes clustering algorithm to identify contexts (or situations) which might induce different privacy behavior of people.

### A. *K-modes Clustering*

K-means clustering is a well-known data mining technique to group data points into K clusters. Each data point is assigned to the cluster with the nearest mean, a representative value of the cluster. However, K-means can only process continuous numeric values as its input. As a variant of K-means, the K-modes clustering algorithm was designed to utilize the K-means paradigm in clustering categorical (or ordinal) values without data conversions. The K-modes algorithm modifies the original K-means by (1) replacing cluster means with cluster modes, (2) using the simple matching dissimilarity function instead of Euclidean distance to calculate the distance between categorical objects, and (3) updating modes with the most frequent categorical values in each iteration of the clustering [26, 27]. More specifically, K-modes clustering divides categorical objects into K groups such that the distance from objects to the assigned cluster modes is minimized. Default simple-matching distance is used to measure the dissimilarity between two categorical objects. It is computed by counting the number of mismatches in all variables. This distance is weighted by the frequencies of the clusters (modes) in the data. We used the *klaR* [28], an R implementation of K-modes, on the collected privacy preferences so as to discover cluster modes and assign each data point to the specific cluster according to its dissimilarity function in an iterative clustering process.

TABLE IV. MODES OF CLUSTERS

| Mode | Contextual Parameters | | | | | Reaction Parameters | | | | | Cluster | Label | Number of Instances | Color Code |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | $C_1$ | $C_2$ | $C_3$ | $C_4$ | $C_5$ | $R_1$ | $R_2$ | $R_3$ | $R_4$ | $R_5$ | | | | |
| $M_1$ | 3 | **2** | **6** | 4 | 0 | 1 | 1 | **6** | **6** | **6** | $CL_1$ | Acceptable | 2,608/6,618 | Green |
| $M_2$ | 2 | **16** | **5** | 2 | 0 | 1 | 2 | **4** | **4** | **4** | $CL_2$ | Neutral | 1,199/6,618 | Yellow |
| $M_3$ | 3 | **20** | **3** | 4 | 0 | 1 | 4 | **3** | **3** | **3** | $CL_3$ | Somewhat unacceptable | 1,492/6,618 | Red |
| $M_4$ | 3 | **17** | **7** | 3 | 1 | 1 | 4 | **1** | **1** | **1** | $CL_4$ | Very unacceptable | 1,319/6,618 | Black |

### B. *Determining Number of Clusters*

Determining the number of clusters (K) is the first step in the data clustering process. We need to find a balance between maximum data compression by assigning all data points into a single cluster (K=1) and maximum accuracy by assigning each data point into an individual cluster (K=n). Thus, we heuristically search for the optimal K by utilizing the well-known Elbow method [29]. First, we compute the sum of errors (SE) of the K-modes clustering with a maximum of 50 iterations, while increasing K from 2 to 10. The SE is defined as the sum of the distance between each instance of the cluster and the cluster's centroid (mode):

$$SE_K = \sum_{i=1}^{K} \sum_{x \in c_i} dist\left(x, c_i\right)$$

where $x$ is a data point belonging to the $i^{th}$ cluster and $c_i$ is the mode of the $i^{th}$ cluster. Next, we calculate the values for the difference between $SE_K$ and $SE_{K-1}$, and find that the largest decrease in errors occurs when we increased K from 3 to 4 (error difference: 1,080, see Fig. 3). Therefore, we chose 4 as a suitable number of clusters, and use it as a parameter (*modes*) for running the K-modes clustering algorithm on our data set. The algorithm then randomly chooses 4 categorical instances as the initial modes, and updates the modes through iterative clustering. Since we did not specify a maximum number of allowed iterations (*iter.max*), the algorithm continues until the clustering error is minimized.

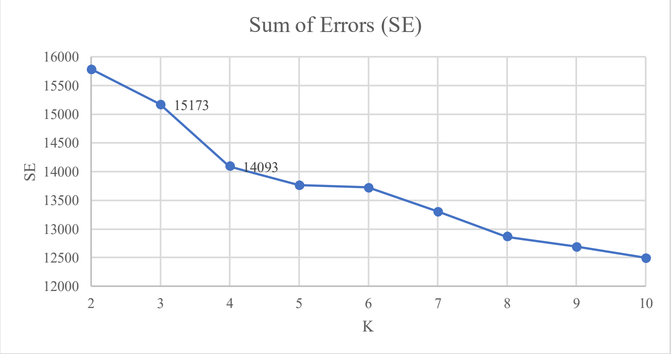

Fig. 3. Sum of Errors

### C. *Interpretation of Clusters*

Table IV presents the resulting cluster modes, which are composed of both contextual and reaction parameter values. The clusters are quite distinct from each other, primarily in the contextual parameters "what" ($C_2$) and "who" ($C_3$). Each mode has a unique categorical value for these parameters, which indicates that $C_2$ and $C_3$ characterize clusters relatively more influentially than the other contextual parameters. Additionally, each mode has identical and unique values for the reaction parameters "_comfort" ($R_3$), "_risk" ($R_4$, reverse-coded), and "_appropriateness" ($R_5$). These parameters represent people's privacy attitudes about IoT scenarios on a scale of 1 to 7. For example, $R_3$=1, $R_4$=1, and $R_5$=1 indicate that a specific scenario is perceived by a participant as "very uncomfortable", "very risky", and "very inappropriate", respectively (see Table II). On the other hand, the remaining reaction parameters such as "_notification ($R_1$)" and "_permission ($R_2$)" do not show unique values per cluster.

Since the reaction parameters $R_3$, $R_4$, and $R_5$ have unique values for each mode, we can mark the clusters using these parameters. We labeled scenarios belonging to the cluster $CL_1$ as "acceptable" to the study participants as its mode has the second highest value for $R_3$, $R_4$, and $R_5$ (namely 6 on a 7-item scale). Likewise, we labelled scenarios for $CL_2$ as "neutral", scenarios for $CL_3$ as "somewhat unacceptable" (since the value of its reaction parameters (3) falls slightly below the scale average), and scenarios for $CL_4$ as "very unacceptable." As a result, 39.4% of the scenario descriptions were grouped into the "acceptable" while 19.9% were grouped into the "very unacceptable" cluster.

### D. *Verification of Clustering Results*

To validate the distinctiveness of the resulting clusters, we performed three Welch's t-tests on the $R_3$ parameter between the following pairs of clusters: ($CL_1$, $CL_2$), ($CL_2$, $CL_3$), and ($CL_3$, $CL_4$). The reason for using Welch's t-test is that all clusters have different variances in the $R_3$ parameter. The tests confirm that the difference in the means of the $R_3$ parameter between each pair of the clusters is statistically significant ($p < 0.016$, Bonferroni-corrected for three comparisons). Next, we also conducted Welch's t-tests on the $R_4$ and $R_5$ parameters, and drew the same conclusion. Therefore, we find the clusters are sufficiently distinct from each other in terms of participants' reactions to the scenarios pertaining to each cluster.

### E. *Analysis of Results*

In this section, we compare the clusters with regard to the five contextual parameters to comprehend how people's reactions to and perceptions of the given IoT scenarios vary depending on the contextual parameters.

*1) where*

Regarding the "where" parameter (see Fig. 4), participants consider monitoring activities as very unacceptable if they occur at their own private places like home ("where"=0, see

$CL_4$; $p < .0001$, Cohen's $d = 0.6069$[1]). This is mainly because people do not exercise self-control in such places, and thus do not want to be monitored. We confirm that these findings are consistent with existing research results such as [20]. In contrast, participants consider monitoring that occurs at public spaces as acceptable ("where"=3, see $CL_1$; $p = 0.000113$, Cohen's $d = 0.2016$). As for semi-public spaces ("where"=2) like a restaurant, participants' attitude is somewhat neutral (see $CL_2$) since it can be perceived as both a personal and a public place, depending on other contextual factors like "what" and "who."

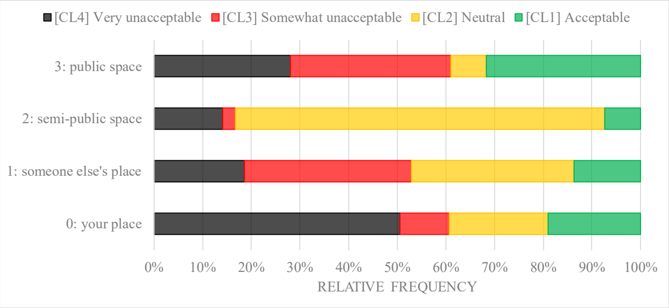

Fig. 4. Relative Distribution of "where" Parameter per Cluster

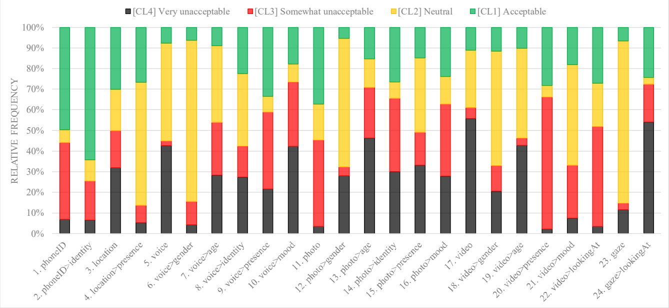

Fig. 5. Relative Distribution of "what" Parameter per Cluster

### *2) what*

In regard to the "what" parameter (see Fig. 5), participants do not allow situations in which someone is videotaping them without a clear purpose ("what"=17, see $CL_4$; $p < .0001$, Cohen's $d = 0.804$) or monitoring their eye movements to figure out what they are looking at ("what"=24, see $CL_4$; $p < .0001$, Cohen's $d = 0.6539$). In this context, participants also consider video monitoring as somewhat unacceptable even if it has some purpose ("what"=20, 22, see $CL_3$; $p < .0001$, Cohen's $d = 0.7449$). Photo-taking ("what"=11, see $CL_1$) is relatively more acceptable to the participants than video monitoring ("what"=17, see $CL_1$); however, they still worry about this activity if it aims to detect their personal information like age ("what"=13, see $CL_4$). Therefore, we can conclude that photo-taking and/or video monitoring of individuals could present significant privacy threats in IoT environments. On the other hand, participants are very open to provide information about their personal devices such as a unique phone identifier ("what"=1, 2, see $CL_1$; $p < .0001$, Cohen's $d = 0.9571$), presumably because they perceive this information not to directly represent their personal behavior.

### *3) who*

In previous studies, we found that the identity of the information requester is an important determinant of people's privacy decisions on various information monitoring activities [19, 24]. Through the present cluster analysis (see Fig. 6), we further confirm that participants' responses to the given scenarios are very privacy-conservative if the entity of the monitoring is unknown to them ("who"=1, see $CL_4$; $p < .0001$, Cohen's $d = 1.1071$), or if it is the government ("who"=7, see $CL_4$; $p < .0001$, Cohen's $d = 1.0858$). Participants also have some privacy concerns if their school/employer ("who"=6, see $CL_3$; $p < .0001$, Cohen's $d = 0.6562$) tracks their personal information and behavior. Interestingly, a fair number of participants feel safe if the monitoring was performed by their school/employer ("who"=6, see $CL_1$; $p < .0001$, Cohen's $d = 0.9128$). These responses are probably because most participants were students who typically trust what their school does.

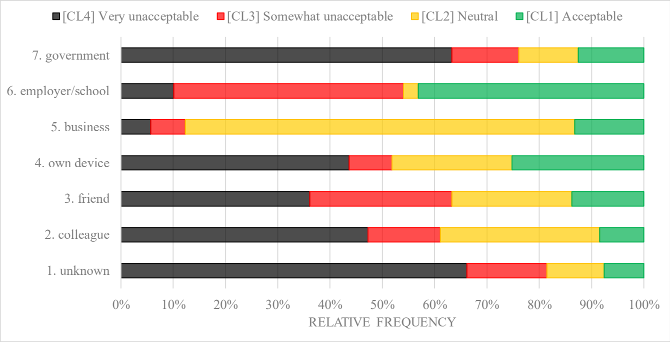

Fig. 6. Relative Distribution of "who" Parameter per Cluster

### *4) reason*

The monitoring purpose can take on one of six values: safety, commercial, social, convenience, health, and "not specified" (see Fig. 7). Participants consider monitoring as very unacceptable when it is performed for social ("reason"=3, see $CL_4$; $p < .0001$, Cohen's $d = 0.9691$) or safety-related purposes ("reason"=1, see $CL_4$; $p < .0001$, Cohen's $d = 0.6245$). This means that these purposes are not convincing enough for participants to allow the respective monitoring activities. For instance, some participants commented that they could not understand why an IoT service would try to recommend new friends to them. Also, participants tend to consider a university campus as safe, thus having difficulties envisioning safety-related IoT service scenarios (e.g., finding wanted criminals through face recognition). Conversely, health is the most significant purpose for participants to accept a given scenario ("reason"=5, see $CL_1$; $p < .0001$, Cohen's $d = 0.6089$). In addition, convenience is also a reasonable justification ("reason"=4, see $CL_1$; $p < .0001$, Cohen's $d = 0.9004$).

[1] In the Social Sciences, effect sizes less than 0.3 are commonly regarded as small, effect sizes between 0.3 and 0.6 as medium, and effect sizes larger than 0.6 as large [30].

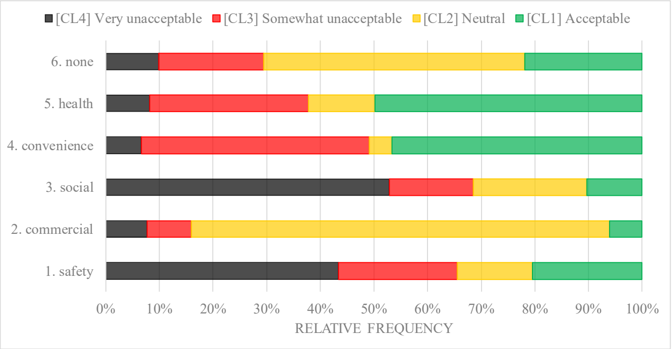

Fig. 7. Relative Distribution of "reason" Parameter per Cluster

### 5) *persistence*

Considering the frequency of monitoring, participants are usually concerned about the risk of privacy violations if IoT devices monitor them continuously, rather than just once (see Fig. 8). Participants are clearly unwilling to accept scenarios with continuous monitoring of personal information ("persistence"=1, see $CL_4$; $p < .0001$, Cohen's $d = 0.7252$). In contrast, one-time monitoring is generally acceptable to them ("persistence"=0, see $CL_1$; $p < .0001$, Cohen's $d = 0.3842$).

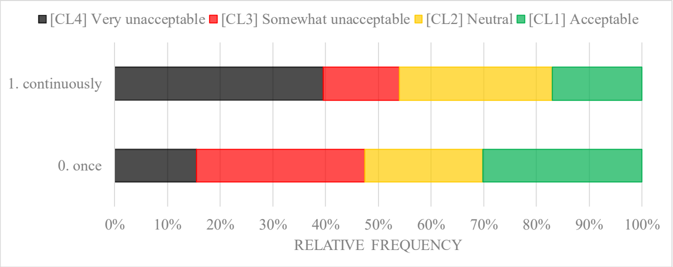

Fig. 8. Relative Distribution of "persistence" Parameter per Cluster

## V. Phase III: Privacy Preference Prediction

Finally, we build machine learning models to predict people's privacy preferences in the analyzed scenarios. Specifically, we aim to predict participants' response to the question

> If this situation [= scenario] happens, would you want to *allow* it?

using contextual parameter values and cluster membership of the scenarios as input features. We focus on the parameter "_permission" because it may reflect people's substantive privacy decisions in IoT environments. We utilize a conditional inference tree for building machine learning models.

### A. Conditional Inference Tree

Conditional inference tree (CTree) is a statistics-based decision tree learning algorithm that uses non-parametric tests as splitting criteria [31]. Unlike other learning algorithms such as recursive partitioning and regression trees (rpart), CTree uses a significance test procedure to select variables to split, instead of information measures like the Gini coefficient. In other words, CTree chooses predictor variables that have a statistically significant relationship ($p < 0.05$) with the response variable as internal nodes of the tree. Because the algorithm performs multiple test procedures (i.e., permutation tests) to determine whether there exist statistical associations between any of the covariates and the response variable, it can not only avoid potential overfitting but guarantee unbiased predictor selection. We used *party* [32], an R implementation of the CTree algorithm, for training CTree decision tree models on our data set.

### B. Experimental Setup

To investigate whether it is possible to predict people's future privacy choices, we learn CTree models (classifiers) to predict values of the parameter "_permission" ($R_2$) for the presented IoT scenarios. Among the attributes of our data set, we chose the five contextual parameters, "where" ($C_1$), "what" ($C_2$), "who" ($C_3$), "reason" ($C_4$), and "persistence" ($C_5$), for specifying a basic feature vector for the classifiers. We saw in Section IV that all these parameters influence people's privacy decision-making. We then added cluster membership ($CL_K$), assigned by the K-modes clustering algorithm, as an additional input feature, to analyze its impact on the predictive power of the decision tree models. Since there are 4 possible values in the parameter $R_2$ (1: allow always, 2: allow just this time, 3: reject just this time, 4: reject always), a prediction for this parameter can be formalized as a multi-level classification problem. We also noticed that many researchers have tried to predict people's binary privacy decisions, namely whether to allow or reject (recommended) privacy settings for personal information disclosure [15, 17, 18]. Therefore, we also build and evaluate CTree models as binary classifiers by converting $R_2$=1, 2 into "allow" and $R_2$=3, 4 into "reject."

### C. Experimental Results

We used 10-fold cross-validation accuracy for estimating prediction performance of the CTree models. In addition, we also computed Cohen's Kappa coefficient for gauging inter-rater agreement in predicting the response variable. In general, Kappa coefficients ranging from 0.4 to 0.6 denote a moderate agreement between two classifiers [33]. For the binary classification, we also measured the F1 score to consider both precision and recall for the classification results.

Table V summarizes the prediction accuracy of the learned CTree models. For multi-level privacy decisions (4 class), the model can predict future decisions with the maximum accuracy of 62%. When we narrowed the possible range of decisions to binary (allow or reject), the accuracy increased to 77%. As can be seen, adding cluster membership as an additional feature improves the performance of both the multi-level and the binary classifiers; it led to an accuracy increase of 21% and 11%, respectively. Performance measures of previous classifications of binary privacy decisions [15, 17, 18] are not directly comparable because each work uses different data sets and definitions of "privacy decision." However, when considering both the F1 score (0.701) and the Kappa coefficient (0.511), our binary classifier shows a prediction accuracy at least above the average of other works. We expect the performance could be further enhanced with the collection of extra training data. This is because a more sufficient amount of data would reduce the uncertainty for the classifier. For these reasons, we believe that it is practically feasible to predict privacy preferences of users in IoT environments if we can extract and model privacy-related contexts from the IoT environment.

TABLE V. CLASSIFICATION PERFORMANCE

| Response Variable | Predictor Variables | Acc. | F1 | Kappa |
|---|---|---|---|---|
| $R_2$ (4 class) | $C_1 + C_2 + C_3 + C_4 + C_5$ | 0.41 | - | 0.116 |
| $R_2$ (4 class) | $C_1 + C_2 + C_3 + C_4 + C_5 + CL_K$ | 0.62 | - | 0.461 |
| $R_2$ (binary) | $C_1 + C_2 + C_3 + C_4 + C_5$ | 0.66 | 0.358 | 0.148 |
| $R_2$ (binary) | $C_1 + C_2 + C_3 + C_4 + C_5 + CL_K$ | 0.77 | 0.701 | 0.511 |

## VI. DISCUSSION AND FUTURE WORK

We showed that it is feasible to group privacy scenarios into clusters with distinct user reactions, and to predict privacy preferences using data mining and machine learning techniques. Yet, our work still has some issues that need to be considered and addressed.

### A. Skewed Participants

Our study participants were skewed toward students aged 18-25 (82%) since we recruited them on campus. This may result in a sampling bias that makes our results less general. For instance, we had earlier conducted a cluster analysis on data collected from Amazon MTurk workers whose age was predominantly between 25 and 40 (57.5%) [25]. We used the same algorithm, but the outcomes were slightly different. Regarding the "who" parameter, for instance, MTurkers trust their own personal device ("who"=4) the most, while participants recruited for the present study trust their school or employer ("who"=6) the most. Thus, we might also need to consider demographic information when building machine learning models for the prediction of privacy preferences. In this regard, we plan to validate our approach and arguments with more representative samples, thereby establishing a future direction of this research (e.g., into personalized privacy preference prediction).

### B. Usability of Google Glass

Many participants mentioned that they became interested in our study because of Google Glass. They wanted to get hands-on experience with Glass as it is considered the most famous smart glass, and has been currently discontinued by its manufacturer. However, participants also complained about its usability. The major issue was the visibility of text shown in the Google Glass display. Google Glass users need to glance slightly upwards to view the screen rather than look straight. For this reason, a few participants felt dizzy shortly after using Google Glass and one even withdrew their participation early. Moreover, some participants had difficulty reading the scenario descriptions displayed in small letters on the screen. As discussed before, we also had to ask participants to record their responses on paper questionnaires because many screens would need to be navigated to see questions and all answer options in Google Glass. For these reasons, we need to devise a new way of letting participants interact with Google Glass. A voice user interface to our IoT Privacy app could be a possible approach for achieving better interaction: text-to-speech for presenting scenarios and questions, and speech recognition for collecting user responses. Its feasibility in practice will still need to be verified though.

### C. Privacy Paradox

We analyzed stated privacy preferences collected in a simulated IoT environment, and not actual behavior in a working IoT environment. Although we tried to make participants believe they were in a real situation, we still do not know how they would actually behave in real world situations. Previous research [34-36] has confirmed that people's stated privacy preferences are often inconsistent with their actual behaviors. However, operational IoT environments are not available to us yet, and hence our setup represents the closest possible approach to privacy decision behavior in the wild.

## VII. CONCLUSION

In this paper, we investigated how people's privacy decision-making in IoT environments can be modeled and predicted. We aimed to simulate user experience in a real IoT environment as realistically as possible, by letting users walk around campus wearing Google Glass, and occasionally asking them about their preferences regarding hypothetical privacy-invasive information tracking at nearby locations. We then performed a cluster analysis on the collected preferences in order to understand users' privacy concerns towards IoT applications and services. The results of the analysis show that IoT scenarios can be grouped into four distinct clusters in terms of their perceived privacy risks. By comparing the resulting clusters, we also extracted a number of contextual factors causing privacy threats in IoT. Lastly, we built decision tree models to predict users' future privacy decisions by utilizing both contextual information and its cluster membership as training data. The trained model showed 77% accuracy in predicting a binary privacy decision whether to accept or reject a specific IoT scenario.

## ACKNOWLEDGMENT

This research has been supported by NSF grant SES-1423629. The human subjects research described herein is covered under IRB protocol #2014-1600 at the University of California, Irvine. The authors would like to thank the pilot test users for their help in designing and developing our IoT Privacy app, and Gloria Mark and the anonymous conference reviewers for their valuable comments on earlier versions of this paper.

## REFERENCES

[1] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," Computer Networks, vol. 54, no. 15, pp. 2787–2805, Oct. 2010.

[2] D. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac, "Internet of Things: vision, applications and research challenges," Ad Hoc Networks, vol. 10, no. 7, pp. 1497–1516, Sep. 2012.

[3] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of Things (IoT): A vision, architectural elements, and future directions," Future Generation Computer Systems, vol. 29, no. 7, pp. 1645–1660, Sep. 2013.

[4] J. A. Stankovic, “Research directions for the Internet of Things,” IEEE Internet of Things Journal, vol. 1, no. 1, pp. 3–9, Feb. 2014.

[5] C. M. Medaglia and A. Serbanati, “An overview of privacy and security issues in the Internet of Things,” The Internet of Things, D. Giusto, A. Iera, G. Morabito, and L. Atzori, Eds. Springer New York, 2010, pp. 389–395.

[6] D. Christin, A. Reinhardt, S. S. Kanhere, and M. Hollick, “A survey on privacy in mobile participatory sensing applications,” Journal of Systems and Software, vol. 84, no. 11, pp. 1928–1946, Nov. 2011.

[7] J. Virkki and L. Chen, “Personal perspectives: Individual privacy in the IoT,” Advances in Internet of Things, vol. 3, no. 2, pp. 21–26, Apr. 2013.

[8] D. J. Liebling and S. Preibusch, “Privacy considerations for a pervasive eye tracking world,” in Proc. 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication, New York, NY, USA, 2014, pp. 1169–1177.

[9] A. Acquisti, L. Brandimarte, and G. Loewenstein, “Privacy and human behavior in the age of information,” Science, vol. 347, no. 6221, pp. 509–514, Jan. 2015.

[10] D. J. Solove, “Privacy self-management and the consent dilemma,” Harvard Law Review, vol. 126, no. 7, pp. 1880–1903, May 2013.

[11] G. Danezis, “Inferring privacy policies for social networking services,” in Proc. 2nd ACM Workshop on Security and Artificial Intelligence, Chicago, IL, USA, 2009, pp. 5–10.

[12] N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao, “Understanding and capturing people’s privacy policies in a mobile social networking application,” Personal and Ubiquitous Computing, vol. 13, no. 6, pp. 401–412, Aug. 2009.

[13] L. Fang and K. LeFevre, “Privacy wizards for social networking sites,” in Proc. 19th International Conference on World Wide Web, Raleigh, NC, USA, 2010, pp. 351–360.

[14] G. Bigwood, F. B. Abdesslem, and T. Henderson, “Predicting location-sharing privacy preferences in social network applications,” in Proc. AwareCast, Newcastle, UK, 2012.

[15] I. Bilogrevic, K. Huguenin, B. Agir, M. Jadliwala, and J.-P. Hubaux, “Adaptive information-sharing for privacy-aware mobile social networks,” in Proc. 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, New York, NY, USA, 2013, pp. 657–666.

[16] Y. Zhao, J. Ye, and T. Henderson, “Privacy-aware location privacy preference recommendations,” in Proc. 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, London, UK, 2014, pp. 120–129.

[17] I. Bilogrevic, K. Huguenin, B. Agir, M. Jadliwala, M. Gazaki, and J.-P. Hubaux, “A machine-learning based approach to privacy-aware information-sharing in mobile social networks,” Pervasive and Mobile Computing, vol. 25, pp. 125–142, Jan. 2016.

[18] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. (Aerin) Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, “Follow my recommendations: A personalized privacy assistant for mobile app permissions,” in Proc. 12th Symposium on Usable Privacy and Security (SOUPS 2016), Denver, CO, USA, 2016, pp. 27–41.

[19] S. Lederer, J. Mankoff, and A. K. Dey, “Who wants to know what when? privacy preference determinants in ubiquitous computing,” in Proc. CHI ’03 Extended Abstracts on Human Factors in Computing Systems, Lauderdale, FL, USA, 2003, pp. 724–725.

[20] E. K. Choe, S. Consolvo, J. Jung, B. Harrison, and J. A. Kientz, “Living in a glass house: a survey of private moments in the home,” in Proc. 13th International Conference on Ubiquitous Computing, Beijing, China, 2011, pp. 41–44.

[21] M. Benisch, P. G. Kelley, N. Sadeh, and L. F. Cranor, “Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs,” Personal Ubiquitous Computing, vol. 15, no. 7, pp. 679–694, Oct. 2011.

[22] C. Hine and J. Eve, “Privacy in the marketplace,” The Information Society, vol. 14, no. 4, pp. 253–262, Nov. 1998.

[23] P. A. Norberg, D. R. Horne, and D. A. Horne, “The privacy paradox: personal information disclosure intentions versus behaviors,” Journal of Consumer Affairs, vol. 41, no. 1, pp. 100–126, Mar. 2007.

[24] R. Chow, S. Egelman, R. Kannavara, H. Lee, S. Misra, and E. Wang, “HCI in Business: A collaboration with academia in IoT privacy,” HCI in Business, F. F.-H. Nah and C.-H. Tan, Eds. Springer International Publishing, 2015, pp. 679–687.

[25] H. Lee and A. Kobsa, “Understanding user privacy in Internet of Things environments,” Internet of Things (WF-IoT), 2016 IEEE 3rd World Forum on, in press.

[26] Z. Huang, “A fast clustering algorithm to cluster very large categorical data sets in data mining,” in Research Issues on Data Mining and Knowledge Discovery, Tucson, AZ, USA, 1997, pp. 1–8.

[27] Z. Huang, “Extensions to the k-Means algorithm for clustering large data sets with categorical values,” Data Mining and Knowledge Discovery, vol. 2, no. 3, pp. 283–304, Sep. 1998.

[28] C. Neumann, kmodes {klaR}: K-Modes Clustering. CRAN repository.

[29] T. S. Madhulatha, “An overview on clustering methods,” arXiv preprint arXiv:1205.1117, May 2012.

[30] J. Cohen, Statistical power analysis for the behavioral sciences (revised ed.). New York: Academic Press, 1977.

[31] T. Hothorn, K. Hornik, and A. Zeileis, “Unbiased recursive partitioning: A conditional inference framework,” Journal of Computational and Graphical Statistics, vol. 15, no. 3, pp. 651–674, Sep. 2006.

[32] T. Hothorn, CTree {party}: Conditional Inference Trees. CRAN repository.

[33] D. G. Altman, “Inter-rater agreement,” in Practical Statistics for Medical Research, CRC Press, 1990, pp. 403–408.

[34] A. Acquisti and J. Grossklags, “Privacy attitudes and privacy behavior,” Economics of Information Security, vol. 12, L. Camp and S. Lewis, Eds. Springer US, 2004, pp. 165–178.

[35] C. Jensen, C. Potts, and C. Jensen, “Privacy practices of Internet users: self-reports versus observed behavior,” International Journal of Human-Computer Studies, vol. 63, no. 1–2, pp. 203–227, July 2005.

[36] K. Connelly, A. Khalil, and Y. Liu, “Do I do what I say?: observed versus stated privacy preferences,” Human-Computer Interaction – INTERACT 2007, vol. 4662, C. Baranauskas, P. Palanque, J. Abascal, and S. Barbosa, Eds. Springer Berlin / Heidelberg, 2007, pp. 620–623.