UNIVERSITY OF CALIFORNIA,
IRVINE

Modeling and Prediction of Privacy Decision-Making in IoT

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Information and Computer Sciences

by

Hosub Lee

Dissertation Committee:
Professor Alfred Kobsa, Chair
Professor Gloria Mark
Doctor Richard Chow

2019

DEDICATION

To my wife Jisu,
my sons Eunjun (Erik) and Minjun (Luke),
my parents and parents-in-law,
my brother Hwasub and sister-in-law Hyejin,
and my friends

in recognition of their unwavering support.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to express my gratitude to my academic advisor and doctoral committee chair, Professor Alfred Kobsa, who has supported me during my time at the University of California, Irvine. His invaluable advice on how to conduct research has helped me further advance my career as an industry research scientist. His *hands off approach with strategic interventions* has trained me to work independently and professionally. His keen insight into the field has not only broadened my viewpoint but also deepened my knowledge of machine learning, recommender systems, and usable privacy.

I would like to thank my doctoral committee members, Professor Gloria Mark and Doctor Richard Chow, who gave me thoughtful feedback on my dissertation topic proposal. Thanks to their help, I was able to successfully complete the final studies of my dissertation.

In addition, I would like to thank all my former collaborators and supervisers (there are too many people to list them all), who have shaped my views on practical implications of machine learning technologies, through our interactions while working together on research projects and/or writing papers about them.

Financial support was provided by the University of California, Irvine, NSF Grant SES-1423629/1640527, and gifts by Intel Corporation. The human subjects research described herein is covered under IRB protocol #2014-1600 at the University of California, Irvine.

CURRICULUM VITAE

Hosub Lee

EDUCATION

Doctor of Philosophy in Information and Computer Sciences **2019**
University of California, Irvine *Irvine, CA*

Master of Science in Computer Science and Engineering **2007**
Seoul National University *South Korea*

Bachelor of Science in Computer Science **2005**
Korea University *South Korea*

RESEARCH EXPERIENCE

Graduate Student Researcher **2014–2019**
University of California, Irvine *Irvine, CA*

Software Engineering Intern **2018–2019**
Samsung Research America *Irvine, CA*

Research Intern **2017–2017**
Intel Corporation *Santa Clara, CA*

Research Intern **2015–2015**
Samsung Research America *Mountain View, CA*

Research Staff Member **2013–2014**
Samsung Advanced Institute of Technology *South Korea*

Research Staff **2007–2013**
Samsung Advanced Institute of Technology *South Korea*

TEACHING EXPERIENCE

Reader **2015–2016**
University of California, Irvine *Irvine, CA*

Teaching Assistant **2005–2005**
Seoul National University *South Korea*

REFEREED JOURNAL PUBLICATIONS

Confident Privacy Decision-Making in IoT Environments (*under review*) **2019**
Transactions on Computer-Human Interaction

An Adaptive User Interface based on Spatiotemporal Structure Learning **2011**
IEEE Communications Magazine

REFEREED CONFERENCE PUBLICATIONS

Privacy Preference Modeling and Prediction in a Simulated Campuswide IoT Environment **Mar. 2017**
IEEE Int'l Conf. on Pervasive Computing and Communications

Understanding User Privacy in Internet of Things Environments **Dec. 2016**
IEEE World Forum on Internet of Things

HCI in Business: A Collaboration with Academia in IoT Privacy **July 2015**
International Conference on HCI in Business

Smart Pose: Mobile Posture-aware System for Lowering Physical Health Risk of Smartphone Users **Apr. 2013**
ACM Conference on Human Factors in Computing Systems

A New Posture Monitoring System for Preventing Physical Illness of Smartphone Users **Jan. 2013**
IEEE Consumer Communications and Networking Conference

Towards Unobtrusive Emotion Recognition for Affective Social Communication **Jan. 2012**
IEEE Consumer Communications and Networking Conference

An Adaptive User Interface based on Spatiotemporal Structure Learning **Jan. 2011**
IEEE Consumer Communications and Networking Conference

An Ontology-based Reasoning Approach for Energy-aware Smart Homes **Jan. 2011**
IEEE Consumer Communications and Networking Conference

REFEREED BOOK CHAPTER PUBLICATIONS

Towards Ubiquitous Privacy Decision Support: Machine Prediction of Privacy Decisions in IoT (*in press*) **2019**
Convergence of Artificial Intelligence and Internet of Things

Personalized Visual Recognition via Wearables: A First Step Toward Personal Perception Enhancement **Sep. 2017**
Personal Assistants: Emerging Computational Technologies

REFEREED WORKSHOP PUBLICATIONS

IoT Service Store: A Web-based System for Privacy-aware IoT Service Discovery and Interaction **Mar. 2018**
IEEE Int'l Conf. on Pervasive Computing and Communications

Personalized Object Recognition for Augmenting Human Memory **Sep. 2016**
ACM Int'l Joint Conf. on Pervasive and Ubiquitous Computing

Racial Violence Archive: Public Information System on Incidents of Violence during the Civil Rights Period **Mar. 2015**
iConference

Mobile Posture Monitoring System to Prevent Physical Health Risk of Smartphone Users **Sep. 2012**
ACM Int'l Joint Conf. on Pervasive and Ubiquitous Computing

Fit Your Hand: Personalized User Interface Considering Physical Attributes of Mobile Device Users **Oct. 2011**
ACM Symposium on User Interface Software and Technology

A New Privacy Protection Scheme for Lost Mobile Device **Sep. 2009**
MobileHCI Future Innovation

ABSTRACT OF THE DISSERTATION

Modeling and Prediction of Privacy Decision-Making in IoT

By

Hosub Lee

Doctor of Philosophy in Information and Computer Sciences

University of California, Irvine, 2019

Professor Alfred Kobsa, Chair

Smartphone apps and websites increasingly ask users to make privacy decisions, e.g., to grant or deny app permission to access their location. Previous research indicates that people are often unable to make these decisions in a reasonable manner due to limits of their available time, motivation, and cognitive decision-making abilities. This problem will continue to grow in ubiquitous computing environments like the Internet of Things (IoT), as an array of IoT devices around the user unobtrusively collects (or infers) his/her personal information. Even though this practice may enable IoT systems to realize highly personalized services for their users, it also raises privacy concerns that may lead users to stop using the service. Therefore, providing IoT services with minimized privacy risks is crucial for both protecting user privacy and keeping IoT ecosystems sustainable. One possible way to achieve this aim is to assist users with making better privacy decisions, by predicting decisions based on their and/or fellow users' historical decision-making behaviors and recommending the privacy settings accordingly (i.e., privacy decision support). To make it a reality, we investigated how to computationally model and predict people's privacy decision-making in IoT. Through both online and situated survey studies, we collected user-stated privacy attitudes and decisions toward a wide range of IoT service scenarios. We then extracted a set of context- and user-specific factors that could impact IoT users' privacy decision-making via statistical analysis on the collected dataset. Based on this dataset, we also conducted a series of

machine learning experiments so as to figure out the most suitable approach for constructing predictive models. These models are trained to infer the optimal privacy decisions toward IoT services that the user had never interacted with. Regarding this, we presented several practical insights we gained from building privacy decision prediction models. Lastly, we designed and proposed a novel privacy-aware system that aims not only to increase users' awareness about privacy implications of using IoT services but also to gather their privacy decision samples made with confidence, which can be utilized as high-quality training data for continuously updating machine learning models for the realization of privacy decision support in IoT.

Chapter 1

Motivation

Smartphone apps and websites increasingly ask users to make privacy decisions about personal information disclosure, e.g., to grant or deny app permission to access their location or their phone book. However, previous research indicates that people are often unable to make rational privacy decisions due to limits of their available time, motivation, and abilities to fully understand the tradeoff between utility benefits and privacy risks in the disclosure of personal information. This problem will continue to grow in ubiquitous computing environments like the Internet of Things (IoT), as an array of computing devices around the user unobtrusively collects his/her personal information or infers more sensitive information from various types of sensor data collected [8]. Clearly, rich personal information helps IoT systems better understand users, thus providing better-tailored services to them. At the same time, however, this leads to a considerable increase in privacy concerns that may lead users to stop using the IoT service [27, 90, 82]. Therefore, providing IoT services with minimized privacy risks is crucial for both protecting users' privacy and keeping IoT ecosystems sustainable. One possible way to achieve this objective is to assist users with making better privacy decisions, by predicting decisions based on their and/or fellow users' historical decision-making behaviors and recommending privacy settings accordingly (i.e., privacy

decision support).

Most previous research on privacy decision support has focused on personal information disclosure in online or mobile social network services. Researchers employed supervised learning approaches that build predictive models by utilizing users' past privacy behaviors as training data, and then utilized the trained models to recommend the most appropriate privacy decisions in the given situation [85, 36, 87, 88, 91, 16]. They also verified that privacy decision support systems alleviate users' cognitive burden, thereby allowing them to make their preferred decisions more easily. This kind of technology will become more crucial in IoT environments, not only because users will need to make privacy decisions more frequently, but also because there are no or only limited user interfaces available for users to state their preferences or decisions to the services (e.g., there are no standard keyboards and displays).

To realize privacy decision support in IoT environments, (1) we first need to understand the extent to which users' privacy decisions are shaped by the context in which the interactions with the IoT services take place [74] or by the individual characteristics such as privacy self-efficacy [58]. In other words, it is necessary to extract meaningful context- and user-specific factors that could influence users' privacy decision-making in IoT. We can then utilize the extracted factors as potential input features for constructing machine learning (ML) models that can learn and predict each user's privacy decisions toward unseen IoT services. In addition, (2) we need to figure out what the best approach is to training the ML model with a reasonable predictive performance, thereby enabling privacy decision support in IoT. To be specific, we need to find the most suitable ML algorithm, feature engineering methodology, and model training strategy, while considering privacy perception and behavior of the user who interacts with IoT. Lastly, (3) we need to investigate how to systematically enhance people's privacy risk awareness regarding IoT service, thereby allowing them to make more informed decisions. This is important for privacy decision support since informed privacy decision samples (training data) would contain consistent behavioral patterns that might be

effectively learned by ML models.

With these aims in mind, we collected, analyzed, and predicted the privacy decisions made by users who are situated in diverse IoT contexts in order to derive practical implications for realizing ubiquitous privacy decision support. To maximize the efficacy of privacy decision support in IoT, we also designed and prototyped a privacy-aware IoT system that aims not only to increase users' awareness of potential privacy risks (i.e., privacy awareness) but also to collect their privacy decisions made with confidence.

This dissertation is organized as follows. In Chapter 2, we summarize the literature concerning the modeling and prediction of privacy decision-making in diverse computing environments. In Chapter 3, we report both online and campus-wide situated survey studies to understand which contextual factors impact users' privacy decisions in IoT and how. We verified the effects of each contextual factor through unsupervised cluster analysis and statistical significance tests. In Chapter 4, we summarize the results of our machine learning experiments to present the optimal ML approach for predicting people's future privacy decisions toward IoT services. We trained several privacy decision prediction models on the dataset collected in a situated survey study, varying the ML algorithm, input features, and training strategy. We then compared the predictive performance of the trained ML models to choose the best. In Chapter 5, we report another online survey study to uncover the latent relationships between IoT users' privacy awareness and their confidence in privacy decision-making. We verified the role and importance of privacy awareness in making confident decisions via both mixed-effect logistic regression and Random Forests ML algorithm. In Chapter 6, we also report a new campus-wide online survey study to validate our findings with regard to privacy awareness. We analyzed qualitative responses and performed ML experiments to further support our previous claims. In Chapter 7, we present a novel web-based privacy-aware system that aims to maximize user understanding of the privacy implications of using IoT technologies, thereby allowing us to secure more reliable privacy

decision samples which can be helpful for training highly accurate privacy decision prediction models. In Chapter 8, we summarize the dissertation and propose future research directions.

Chapter 2

Literature Survey

In this chapter, we summarize the literature regarding the modeling and prediction of people's privacy decision-making in various computing environments.

2.1 Modeling of Privacy Decision-Making

We first presented a literature review of previous studies aimed at understanding the causes and effects of users' privacy behavior in mobile/networked computing environments. In this regard, researchers have investigated several contextual factors that could influence users' privacy concerns in diverse application scenarios.

Lederer et al. conducted a scenario-based online survey to evaluate the relative importance of two factors, requester and situation, in determining users' privacy preferences in networked computing environments [59]. 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 analyz-

ing 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. 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 [23]. 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. performed a user study in order to identify contextual factors that influence users' willingness to share their location with others [13]. Using a web-based online survey, they collected detailed preferences from human subjects ($N = 27$) 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. Hu et al. developed a context-aware location sharing system [42]. Their assumption was that users are more likely to share location information with IoT devices in *emergency situations* than under normal circumstances. However, there exist no user studies or experiments to support this claim.

2.2 Prediction of Privacy Decisions

Here, we presented a literature review of privacy decision prediction based on machine learning techniques. A considerable research effort has been made toward developing intelligent agents which infer and recommend users' privacy decisions on diverse personal information disclosure scenarios. Researchers have also claimed that this kind of technology could help users make a correct decision, thereby minimizing potential privacy risks. Much research on the prediction of privacy decision-making focused on online/mobile social network services (SNS). By using (semi-) supervised machine learning techniques, researchers aimed to accurately predict SNS users' sharing policies for their own contents (e.g., to allow or disallow Facebook friend John to see my photos). This is because vendor-provided privacy setting mechanisms, which ask users to manually specify their preferences, have been proven ineffective for protecting user privacy [85, 91, 16]. There also exists works regarding the prediction of people's decisions toward app permission requests on the Android platform.

2.2.1 Privacy Decision Prediction on Online SNS

Fang et al. proposed a framework named Privacy Wizard that infers and recommends each user's access control policies for his/her personal information on Facebook [36]. Each policy specifies who can access specific personal information. The authors adopted a supervised machine learning approach to learn each user's policies by asking him/her a number of

questions (e.g., would you like to share your birthday with Facebook friend John?). Users' answers were considered intended policies (labels) for their friends. Regarding input features for machine learning models, the authors utilized both *demographic information* and *community membership* which characterized each user. Most importantly, the framework asked users about policies that machine learning models are most uncertain about (namely, active learning with uncertainty sampling). By selectively asking users to label the most informative data first, active learning can effectively reduce manual labeling efforts. In the experiments with active Facebook users ($N = 45$), Privacy Wizard showed 90% accuracy in predicting individual's privacy policies with a small amount of labeled training data (25 out of 200 friends with privileges). The authors utilized a decision tree as a machine learning algorithm.

Shehab et al. proposed a framework called Policy Manager that predicts users' sharing policies on SNS [87]. Like [36], Policy Manager was designed to infer binary access control policies on a specific personal object posted on each user's SNS account. The authors used both *user profiles* (that included, e.g., gender) and *social network structure* (e.g., closeness among users) as input features, and asked users to manually determine policies for a subset of their friends (labeling of training data). They tried nine different machine learning algorithms for each user, and chose the best algorithm showing the highest cross validation accuracy on his/her data. In addition, Policy Manager selected models (i.e., classifiers) trained by other users, based on their accuracy in predicting a target user's training data. It then utilized these classifiers with the target user's own classifier to produce a final classification (i.e., classifier fusion via group voting). The authors tested their framework with Last.FM users ($N = 200$), and confirmed that the proposed classifier fusion approach was effective in improving predictive performance compared to using an individual classifier only: 83% and 70% best accuracy when using an alternating decision tree (ADTree) machine learning algorithm, respectively. The authors extended their work by applying an active learning paradigm into the semi-supervised learning framework [88]. Similar to [36], they aimed at

minimizing user burden in labeling training data by selecting the most informative data points to label. Sinha et al. also developed automated tools to assist users in correctly configuring privacy policies for text-based content on Facebook [91]. The authors utilized diverse features such as the *text of posts*, *time of creation*, *n-grams*, *the previous policy*, and *attachments* to predict future policies (e.g., visible to only me) for Facebook posts. They also adopted a supervised learning approach based on the MaxEnt algorithm. Through an online survey study with real Facebook users, they found out that the system could predict policies with a maximum accuracy of 81%, leading to a 14% increase in accuracy compared to when users used Facebook's privacy setting mechanism for a new post.

2.2.2 Privacy Decision Prediction on Mobile SNS

Sadeh et al. proposed a mobile social networking application, PeopleFinder, which recommends optimized location privacy policies to users [85]. To this end, the authors adopted Random Forests, an ensemble supervised learning method, to build a classifier predicting sharing policies for each user's current location based on his/her previous decisions. The authors proved that these machine-generated policies have better accuracy than the user-defined policies: 91% and 79% success rate in matching users' actual behavior, respectively. Recently, Bilogrevic et al. proposed a personal information sharing platform named SPISM that semi-automatically determines whether to disclose users' personal information on SNS and at what level of granularity [16]. Like other previous works, the authors used a supervised learning approach to predict SNS users' privacy decision-making. For each user, SPISM constructed a multi-level classifier based on Naïve Bayes or support vector machine (SVM) by using his/her past behavior as training data. Training data are composed of diverse personal and contextual factors, such as the *identity of requester*, *type of information requested*, *user location*, *co-presence of others*, and *time* (features) and each user's past privacy decision (label). Like [36, 88], SPISM also adopted an active learning paradigm to minimize users'

labeling efforts. SPISM made decisions automatically whenever the confidence (probability) in the classification result was high enough; otherwise it explicitly asked users' decisions then added them to the preexisting training data. With the updated training data, SPISM continuously learns and adapts to users' privacy behaviors. Therefore, it will require less and less user input over time. A user study with human subjects ($N = 70$) indicated that SPISM outperforms user-defined policies; it showed a median prediction accuracy of 72% when each user provided 40 manual decisions. Even if the authors focused on building a personalized machine learning model for each user (i.e., individual modeling), they also assessed potentials of one-size-fits-all modeling. A universal model trained by all users' data showed a reasonable performance with a median accuracy of 67%. The authors claimed that one-size-fits-all modeling could be suitable for building an initial predictive model that produces default privacy settings for the new user.

2.2.3 Privacy Decision Prediction on Android Mobile Platform

Liu et al. developed and evaluated a personalized privacy assistant (PPA) that proactively produces permission settings for Android applications on behalf of users [65]. 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*) 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 real Android users ($N = 72$), the authors confirmed that 78.7% of the recommendations made by PPA was accepted by the users.

2.3 Gained Insights

Regarding the modeling of privacy decisions, we find little research that comprehensively investigates a wide range of contextual factors influencing users' attitudes and preferences toward their privacy, particularly in an IoT environment. Furthermore, there has been a lack of research that studies the effects of users' perceptions or understanding of privacy implications regarding IoT services on their actual privacy decision-making processes. One of the primary aims of this dissertation is to fill this gap. To this end, we aim to collect people's various reactions toward hypothetical IoT service scenarios which are composed of multiple contextual factors. Through the quantitative analysis of the collected user responses, we construct each user's characteristics in terms of privacy (i.e., personal privacy propensity). After that, we reveal the latent relationships among contextual factors, privacy propensity, and decision-making behaviors via statistical modeling.

With regard to the prediction of privacy decisions, all previous works not only confirm the necessity of decision support systems for protecting user privacy in diverse computing environments, but also provide practical guidelines for learning people's privacy behavior which often evolves over time. However, most previous works utilized conventional machine learning algorithms to predict people's privacy decisions in online/mobile SNS. Considering the ever-increasing complexity and diversity of IoT systems, these algorithms might not be optimal for capturing high-dimensional privacy behavioral patterns can be observed in using IoT services. Also, researchers did not take personal privacy propensity into account when building and evaluating their predictive models. To tackle these difficulties, in this dissertation, we adopt a novel machine learning methodology (that included, e.g., deep learning algorithm) while using people's privacy propensity as both input features and criteria for training machine learning models. We also uncover the importance of the quality of training data on the predictive performance of the trained models via extensive experimentation.

Chapter 3

Effects of Contextual Factors on Privacy Preferences in IoT

In this chapter, we conduct a series of studies to uncover the relationships between contextual information and privacy preferences in IoT-enabled spaces. To begin with, we perform an online survey via Amazon Mechanical Turk (MTurk) to collect people's privacy preferences toward hypothetical IoT service scenarios, which are composed of combinations of diverse contextual factors. We then perform unsupervised cluster analysis on the collected dataset for drawing useful implications regarding user privacy in IoT. To gain additional insights, we also conduct a location-based survey to better capture and analyze the situated privacy preferences of the user who interacts with IoT.

3.1 Introduction

With the widespread use of artificial intelligence technologies like machine learning, tech firms are developing new products aimed at making our lives more convenient and productive. For

example, Apple, Google and Microsoft are developing intelligent personal assistants, such as Apple Siri, Google Assistant and Microsoft Cortana. These products provide services tailored to each individual user by proactively predicting their needs. To better understand the user, service providers and device manufacturers aggressively collect personally identifiable information (e.g., location data, photos of users' faces, voice recordings) and use it as training data for their intelligent services. These industry practices will become even more powerful in future ubiquitous computing environments such as IoT, given that nearly all devices are networked and can collectively gather a wide range of sensor data about users.

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 [8, 72, 37, 94]. 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.

As such, 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 sensor data and/or personal information without informing users, let alone asking for their permission [70, 27, 97, 62]. Without doubt, greater and more detailed volumes of user-related data gathered from IoT devices can enable intelligent IoT services to better understand their users and thus provide more accurately personalized services to each individual user. At the same time, however, the uninformed collection (or inference) of personal information can lead to serious violations of privacy expectations that may harm

both the user and the reputation of the firms providing these services [37]. Hence, providing IoT services with minimized privacy risks is very important for both protecting users' privacy expectations and keeping intelligent IoT services sustainable. In order to achieve these objectives, researchers and developers should understand how different contextual factors influence people's privacy perceptions in an IoT environment. This understanding will enable them to better design and develop privacy-preserving IoT systems and services. Therefore, we performed both online and situated survey studies to uncover meaningful relationships between users' current context and their privacy preferences in IoT.

3.2 Collection and Analysis of Privacy Preferences in IoT

3.2.1 Online Survey Study

In earlier work [26], we conducted an interview study with the goal of qualitatively assessing users' privacy perceptions regarding different IoT scenarios. We interviewed 10 participants about 9 IoT scenarios, to gather their opinions on information monitoring activities which they may encounter in their daily lives. These scenarios differed from each other in terms of five contextual parameters: place (*where*), type of collected personal information (*what*), agent (*who*), purpose (*reason*) and frequency (*persistence*) of the monitoring. We then asked participants for their thoughts on each scenario in terms of several reaction parameters: willingness to be notified (*_notification*), willingness to allow tracking (*_permission*) and evaluations of comfort, risk and appropriateness of the monitoring (*_comfort*, *_risk* and *_appropriateness*).

Even though our previous work provided us with useful insights to extract contextual param-

eters that might induce privacy violations in IoT environments, it was still unclear how these parameters actually affect people's concerns about device tracking in such environments. To address this issue, we performed a cluster analysis on online survey data ($N = 200$), composed of 2,800 IoT scenarios (i.e., contextual parameters) and user responses thereto (i.e., reaction parameters). Because all parameters have either categorical or ordinal values, we utilized K-modes, a variant of the K-means clustering algorithm. We determined that four clusters ($K = 4$) are optimal, each of which is associated with three unique reaction parameters. We compared the identified clusters with respect to each contextual parameter, and discovered some latent relationships between the given contextual information and people's privacy preferences in IoT.

3.2.1.1 Data Collection

We recruited 200 participants on Amazon Mechanical Turk (MTurk), educated them about IoT and asked for their opinions on 14 IoT scenarios one by one. To assure the quality of survey responses, we restricted participants to adults who live in the United States, are proficient in English and have a high worker reputation (above 95% approval ratings). 100 females and 99 males participated (one person did not disclose their gender), and the majority (57.5%) are aged 25-40.

We generated IoT scenario descriptions through random combinations of the five contextual parameters. Table 3.1 shows the possible parameter values. Since the scenarios were unique for each participant, 2,800 scenarios[1] were created in total. Like in the interview study, we enquired participants about their privacy concerns on the presented scenarios. Table 3.2 shows the possible values of the reaction parameters[2]. In addition, we asked participants to describe their opinions in a free text field (i.e., qualitative feedback). We then performed

[1]Sample scenario description: A device of a *friend* ($\mathtt{C}_3 = 3$) records your *voice to check your presence* ($\mathtt{C}_2 = 9$). This happens *once* ($\mathtt{C}_5 = 0$), while you are at *semi-public place* ($\mathtt{C}_1 = 2$), for your *safety* ($\mathtt{C}_4 = 1$).

[2]Sample question: If this situation happens, would you want to be *notified* ($\mathtt{R}_1$) about it?

a cluster analysis to determine in what way these contextual parameters affect people's reactions to being tracked in IoT environments.

3.2.1.2 Cluster Analysis

K-modes clustering The K-means clustering algorithm is the most popular data mining technique to partition observations into K clusters. Each observation is assigned to the cluster with the nearest mean, which itself serves as a representative value of the cluster. However, the applicability of K-means is restricted to continuous numeric values. A variant of K-means, the K-modes clustering algorithm, aims to utilize the K-means paradigm for clustering categorical data without the need for data conversion. The K-modes clustering algorithm makes the following extensions to K-means: (1) replacing cluster means with modes, (2) using the simple matching dissimilarity function in place of the Euclidean distance function to compute the distance between categorical objects and (3) updating modes with the most frequent categorical attributes in each iteration of the clustering [43, 44]. To be specific, it 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 determine the dissimilarity of two objects. It is computed by counting the number of mismatches in all variables. This distance is weighted by the frequencies of the categories (modes) in data. We used `klaR`, an R implementation of the K-modes clustering algorithm, on our Amazon MTurk survey data to find cluster modes and assign each data point to the corresponding cluster based on its dissimilarity function through an iterative clustering process.

Selecting the number of clusters Determining the correct number of clusters is an important step in unsupervised data clustering like K-modes. 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

Parameter (`id`)	Values
where ($\mathtt{C}_1$)	(0) your place (1) someone else's place (2) semi-public space (e.g., restaurant) (3) public space (e.g., street)
what ($\mathtt{C}_2$)	(1) phoneID (2) phoneID⇒identity (3) location (4) location⇒presence (5) voice (6) voice⇒gender (7) voice⇒age (8) voice⇒identity (9) voice⇒presence (10) voice⇒mood (11) photo (12) photo⇒gender (13) photo⇒age (14) photo⇒identity (15) photo⇒presence (16) photo⇒mood (17) video (18) video⇒gender (19) video⇒age (20) video⇒presence (21) video⇒mood (22) video⇒lookingAt (23) gaze (24) gaze⇒lookingAt
who ($\mathtt{C}_3$)	(1) unknown (2) colleague/fellow (3) friend (4) own device (5) business (6) employer/school (7) government
reason ($\mathtt{C}_4$)	(1) safety (2) commercial (3) social (4) convenience (5) health (6) none
persistence ($\mathtt{C}_5$)	(0) once (1) continuously

Table 3.1: Contextual Parameters of IoT Scenarios

Parameter (id)	Values
_*notification* (R1)	(1) notify me, always (2) notify me, just this time (3) don't notify me
_*permission* (R2)	(1) don't allow, always (2) don't allow, just this time (3) allow, just this time (4) allow, always
_*comfort* (R3)	(1) very uncomfortable (2) uncomfortable (3) somewhat uncomfortable (4) neutral (5) somewhat comfortable (6) comfortable (7) very comfortable
_*risk* (R4)	(1) very risky (2) risky (3) somewhat risky (4) neutral (5) somewhat safe (6) safe (7) very safe
_*appropriateness* (R5)	(1) very inappropriate (2) inappropriate (3) somewhat inappropriate (4) neutral (5) somewhat appropriate (6) appropriate (7) very appropriate

Table 3.2: Reaction Parameters of Privacy Preferences

($K = N$). As a priori knowledge of the appropriate value of K does not exist for our data set, we heuristically searched for the optimal K by applying the Elbow method [68]. First, we computed the sum of errors (SE) of the K-modes clustering with a limit of 50 iterations, while increasing K from 2 to 10. The SE is defined as the sum of the distance between each member of the cluster and the cluster's centroid (mode):

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

where x is a data point belonging to the i^{th} cluster and c_i is the mode of the i^{th} cluster. Then, we calculated the difference values between SE_K and SE_{K-1}, and found that the largest decrease in errors occurs when we increase K from 3 to 4 (Table 3.3). Therefore, we chose 4 as the appropriate number of clusters, and used it as a parameter (`modes`) for initializing the K-modes clustering algorithm. The algorithm then randomly choses 4 rows from the dataset as the initial modes, and updates the modes through iterative clustering. As we did not configure the maximum number of iterations allowed (`iter.max`), the algorithm continued until the clustering error was minimized.

Num of Clusters (K)	Sum of Errors (SE_K)	Error Diff. ($SE_{K-1} - SE_K$)
2	15765	
3	15075	-690
4	**14170**	**-905**
5	13655	-515
6	13129	-526
7	12917	-212
8	12562	-355
9	12329	-233
10	12311	-18

Table 3.3: Errors in Clustering IoT Scenarios

Interpretation of clusters Table 3.4 shows the modes that the clustering algorithm generated as the centroids of the four clusters. Interestingly, the clusters differ from each other primarily in the contextual parameters *what* ($\texttt{C}_2$) and *who* ($\texttt{C}_3$): each mode has a unique categorical value for these parameters. This means that *what* and *who* define clusters more than the remaining contextual parameters *where* ($\texttt{C}_1$), *reason* ($\texttt{C}_4$) and *persistence* ($\texttt{C}_5$). In addition, each mode has identical and unique values for the reaction parameters *_comfort* ($\texttt{R}_3$), *_risk* ($\texttt{R}_4$, reverse-coded) and *_appropriateness* ($\texttt{R}_5$). These three parameters indicate respondents' attitudes about a scenario on a scale of 1 to 7. For instance, $\texttt{R}_3 = 1$, $\texttt{R}_4 = 1$ and $\texttt{R}_5 = 1$ indicate that the scenario is perceived as *very uncomfortable*, *very risky* and *very inappropriate*, respectively (see Table 3.2).

Mode (Cluster)	Contextual Param. $\{C_1, C_2, C_3, C_4, C_5\}$	Reaction Param. $\{R_1, R_2, R_3, R_4, R_5\}$	Label (Number of Instances)
$M_1(CL_1)$	0, **8**, **4**, 6, 0	3, 4, **6**, **6**, **6**	*Acceptable* (352/2,800)
$M_2(CL_2)$	2, **9**, **3**, 1, 0	1, 1, **3**, **3**, **3**	*Smwht. Unacceptable* (466/2,800)
$M_3(CL_3)$	3, **22**, **5**, 6, 1	1, 1, **2**, **2**, **2**	*Unacceptable* (840/2,800)
$M_4(CL_4)$	0, **24**, **1**, 6, 0	1, 1, **1**, **1**, **1**	*Very Unacceptable* (1,142/2,800)

Table 3.4: Modes of Clustered IoT Scenarios (Online Survey)

In contrast, the remaining reaction parameters *_notification* (R_1) and *_permission* (R_2) do not have unique values per cluster.

Since the reaction parameters R_3, R_4 and R_5 pertaining to each mode are unique, we can characterize each cluster along these parameters. We will label scenarios that belong to the cluster CL_1 as *acceptable* to the study participants because its mode M_1 has the second highest value for R_3, R_4 and R_5 (namely 6 on a 7-item scale). We label CL_2 scenarios as *somewhat unacceptable* (since the value of its reaction parameters (3) falls slightly below the scale average), CL_3 scenarios as *unacceptable* and CL_4 scenarios as *very unacceptable*. As it can be seen, only 12.6% of the scenario descriptions that we presented to participants fall into the *acceptable* cluster, while 40.8% fall into the *very unacceptable* cluster.

Verification of clustering results To verify the distinctiveness of the clusters, we first conducted 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 of these clusters is statistically significant ($p < 0.016$, Bonferroni-corrected for three comparisons). Then, we also performed Welch's t-tests on the R_4 and R_5 parameters, and reached the same conclusion. Thus, we find that the clusters are sufficiently distinct from each other in terms of user reactions to the scenario descriptions pertaining to each cluster.

Next, we visually inspect clusters so that we can confirm our cluster labeling is rea-

sonable. In doing this, we utilized the reaction parameter $\mathtt{R}_5$ since *appropriateness* is a crucial element for assessing potential privacy risks in a given context. This concept has been proposed by Helen Nissenbaum's Contextual Integrity theory [74], which provides a systematic way of determining when and why people perceive certain usage and disclosure of personal information as appropriate, or as a privacy violation.

We first assign colors to clusters: green for $\mathtt{CL}_1$, yellow for $\mathtt{CL}_2$, red for $\mathtt{CL}_3$ and black for $\mathtt{CL}_4$. We then project all scenario descriptions from all clusters onto a 2-dimensional space, using their $\mathtt{R}_5$ reaction parameter as both their x and y values and their cluster color as their surface code. We add a small amount of random noise to the coordinates of each data point to make them visible. Figure 3.1 shows that situation descriptions that respondents deemed very inappropriate ($\mathtt{R}_5 = 1$) mostly became clustered into $\mathtt{CL}_4$ (black). In contrast, situation descriptions that respondents deemed appropriate or very appropriate ($\mathtt{R}_5 = 6, 7$) became clustered into $\mathtt{CL}_1$ (green).

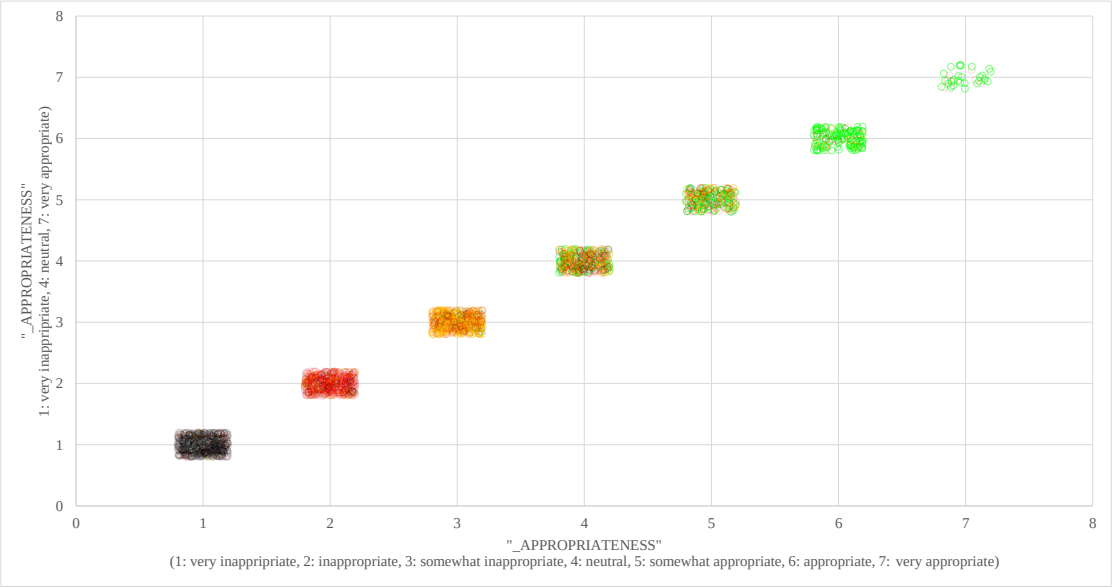

Figure 3.1: Visualized Clustering Results of IoT Scenarios

3.2.1.3 Comparison of Clusters based on Contextual Parameters

In this section, we compared the clusters in terms of the five contextual parameter values pertaining to their modes, to understand how contextual information influences people's reactions toward their privacy in IoT environments.

***where* parameter** Regarding the *where* parameter (see Figure 3.2), participants consider monitoring that occurs at personal places like their homes as very unacceptable (*where* = 0, see $\texttt{CL}_4$; $p < .0001$, Cohen's $d = .479$[3]). This is probably because people do not exercise self-control in such private spaces, and therefore do not want to be monitored. In addition, many participants also have privacy concerns if the monitoring is performed in a public space (*where* = 3, $\texttt{CL}_3$; $p < .0001$, $d = .4921$). A participant commented:

> "Serious invasion of privacy (yes, even in a public place). If the data is stored, a profile could be created as to what I am doing or where I am going." [P46]

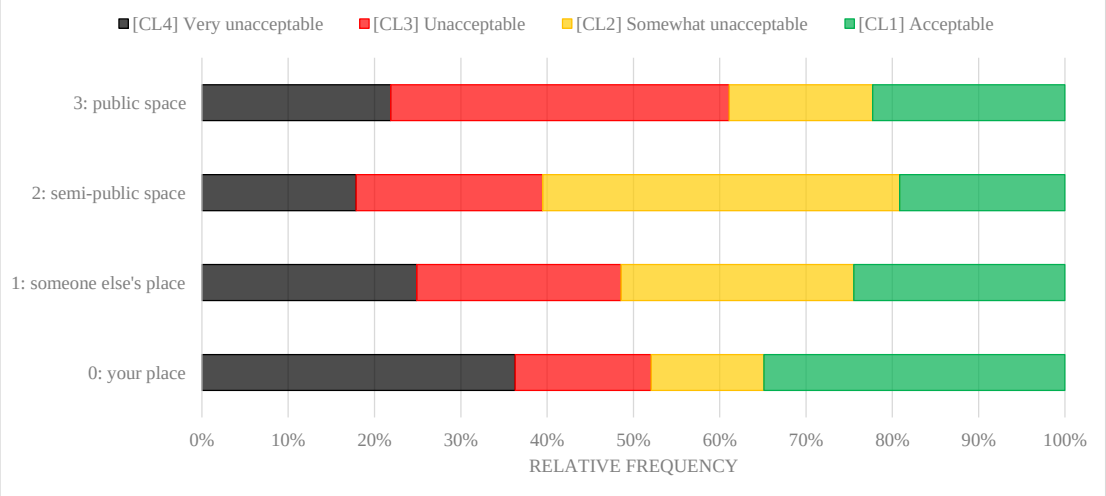

Figure 3.2: Relative Distribution of *where* Parameter per Cluster (Online Survey)

[3]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 [28].

As for semi-public spaces such as restaurants, participants feel that monitoring is somewhat unacceptable (*where* = 2, see $\texttt{CL}_2$; $p < .0001$, $d = .6109$). Interestingly, personal place is a dominant factor for making scenarios acceptable (*where* = 0, see $\texttt{CL}_1$). Therefore, we need to further investigate other contextual factors like the *what* and *who* parameters to fully understand this cluster.

***what* parameter** With regard to the *what* parameter (see Figure 3.3), participants do not accept situations in which someone is monitoring them to figure out what they are looking at (*what* = 23, 24, see $\texttt{CL}_4$; $p = .0001$, $d = .3041$). Participants also considered photo-taking and/or video monitoring unacceptable for various purposes (*what* = 12, 16, 22, see $\texttt{CL}_3$; $p < .0001$, $d = .319$).

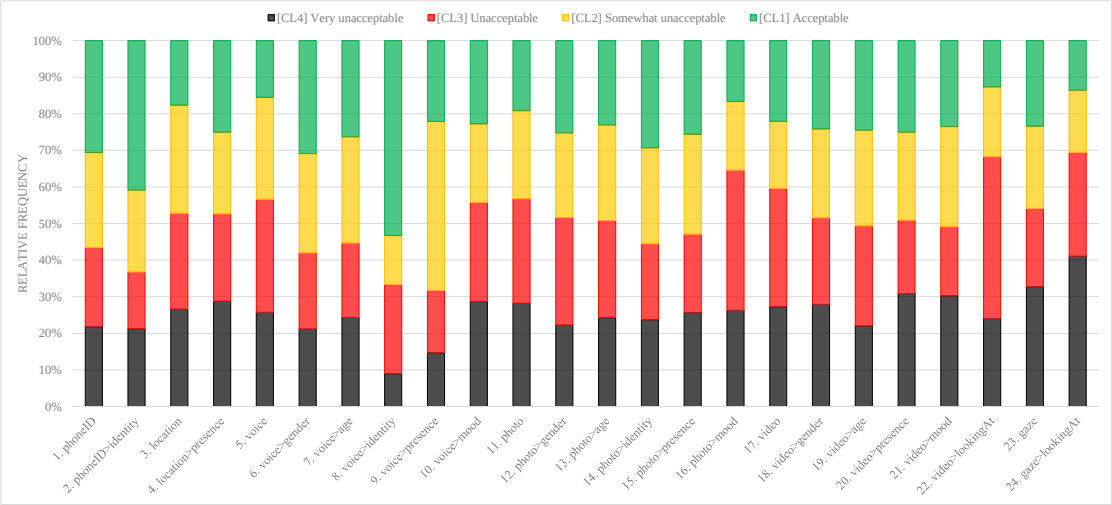

Figure 3.3: Relative Distribution of *what* Parameter per Cluster (Online Survey)

Participants' reactions toward voice monitoring were generally positive compared to the above cases. For instance, many participants are likely to allow their voice to be monitored for gender identification and location awareness (*what* = 6, 9, see $\texttt{CL}_2$; $p = .0006$, $d = .2713$). They are also very open to giving their personally identifiable information such as unique phone ID or voiceprint (*what* = 2, 8, see $\texttt{CL}_1$; $p < .0001$, $d = .6237$), presumably due to convenience or habituation. For instance, participant

P19 deemed voice-based authentication and personalization as acceptable:

> "Maybe a voice recording could be used in place of credit/debit cards for transaction purposes. The system analyzes the recording and knows what you want from the business and prepares your order or services." [P19]

***who* parameter** In our previous interview study, some interviewees stated that *who* is an important parameter affecting their privacy preferences regarding IoT services. The results of our present study clarify its impact (see Figure 3.4). If the monitoring entity is unknown to online survey participants, their responses on the given scenarios are very conservative ($who = 1$, see $\texttt{CL}_4$; $p < .0001$, $d = .7268$). They also do not trust the government ($who = 7$, see $\texttt{CL}_4$; $p < .0001$, $d = .2603$). Participants also have privacy concerns if a nearby business tracks their personal information ($who = 5$, see $\texttt{CL}_3$; $p < .0001$, $d = .5845$). This may be because they doubt that the company safeguards their information. People feel safe if the monitoring is performed by either their friends ($who = 3$, see $\texttt{CL}_2$; $p < .0001$, $d = .6305$) or own devices such as their smartphone ($who = 4$, see $\texttt{CL}_1$; $p < .0001$, $d = .9989$).

***reason* parameter** We specified the purpose of each IoT scenario using the *reason* parameter. This parameter has six possible values, namely *safety*, *commercial*, *social*, *convenience*, *health* and *not specified* (see Figure 3.5). The absence of a purpose causes the greatest number of unacceptable scenarios ($reason = 6$, see $\texttt{CL}_3$ and $\texttt{CL}_4$; $p < .0001$, $d = .3221$). On the other hand, a considerable amount of scenarios was still considered acceptable ($reason = 6$, see $\texttt{CL}_1$), even though the purpose of monitoring was not indicated. This suggests that participants have a tendency to base a privacy decision mainly on concrete contextual factors like the *who* and *what* parameters, as we explained in the previous sections. In fact, several participants reacted to purposeless scenarios by imagining the possible purposes by themselves:

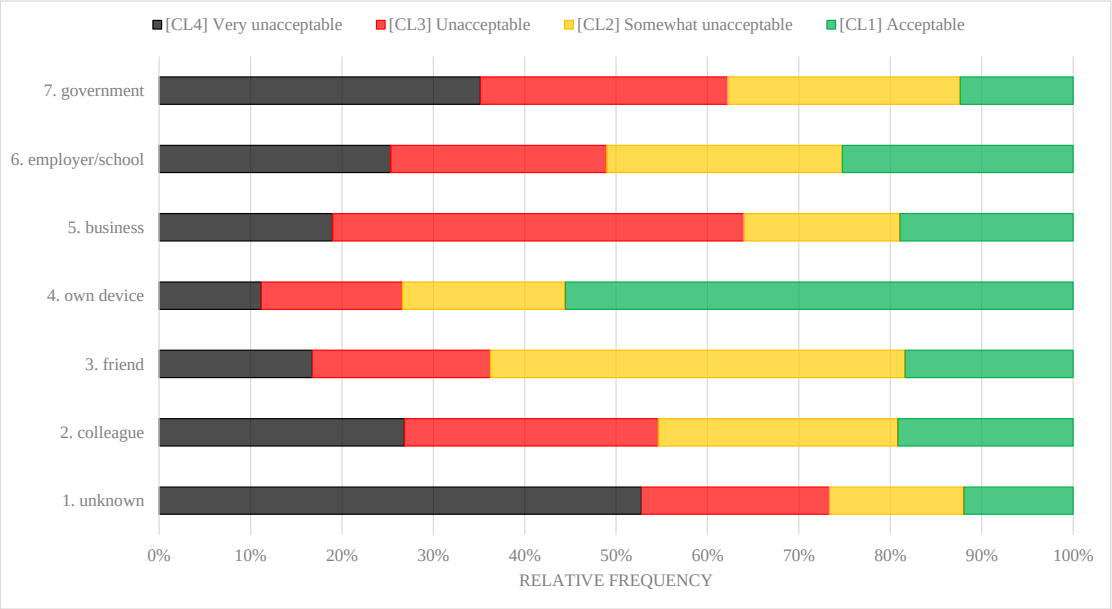

Figure 3.4: Relative Distribution of *who* Parameter per Cluster (Online Survey)

> "It would be able to find criminals and catch criminal behavior on tape so the benefit could possibly improve public safety." [P10]

> "If I was on a trip to get to my friend's house, they could see how far I am from them without having to call or text me." [P105]

> "If I have a health related accident or injury then the person watching can come assist me immediately." [P112]

Other than *not specified*, *convenience* is the most significant purpose that participants found acceptable ($reason = 4$, see $\texttt{CL}_1$). Additionally, *safety* is a reasonable justification for participants to generally accept the situation ($reason = 1$, see $\texttt{CL}_2$).

***persistence* parameter** We assumed that participants will have strong privacy concerns if information monitoring happens continuously ($persistence = 1$) rather than just once ($persistence = 0$). However, our analysis results are inconclusive: no clear tendency toward the one or the other can be seen in Figure 3.6.

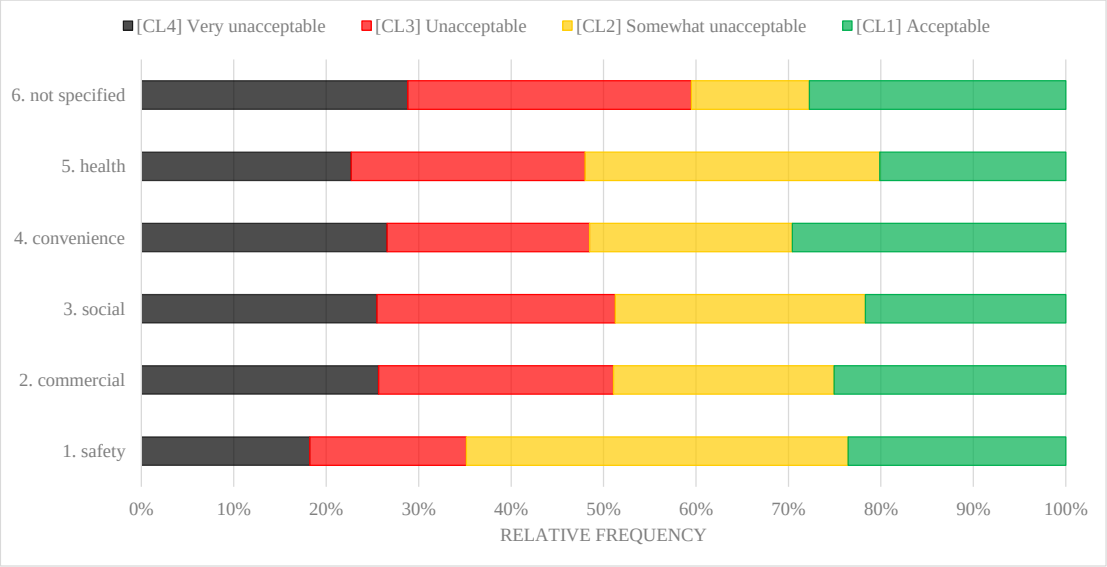

Figure 3.5: Relative Distribution of *reason* Parameter per Cluster (Online Survey)

3.2.1.4 Limitations

Our analysis showed how each contextual factor impacts people's privacy preferences in IoT environments. Yet, this study still has some limitations that need to be considered. We notice that some contextual parameters were defined with coarse granularity. For example, *someone else's place* of the *where* parameter might be interpreted differently by different participants because the meaning of *someone else* is too broad. Furthermore, given that participants took this online survey at a location that has no association with the IoT scenarios described in the survey, there could have been a decreased sense of realism to the scenarios. Privacy research has repeatedly shown that people's stated attitudes with regard to privacy often differ from their actual behaviors in a concrete situation [3, 45, 30]. For these reasons, *out-of-context* attitudinal studies like our online survey must be viewed with some caution and be verified.

We plan to tackle these limitations by developing a location-based (i.e., situated) survey system on a mobile/wearable device which presents scenarios that are specifically related to

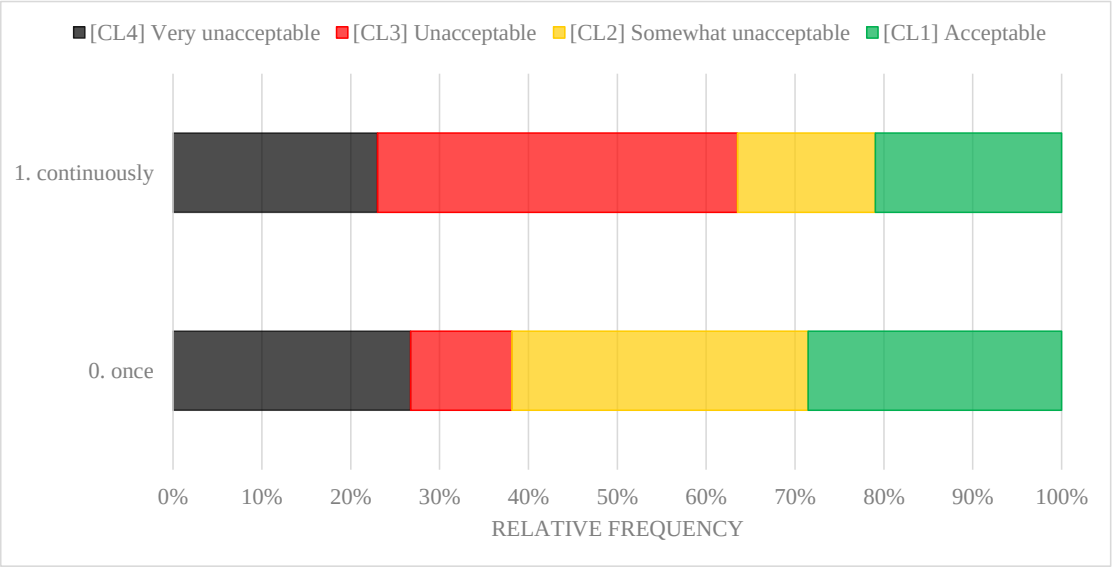

Figure 3.6: Relative Distribution of *persistence* Parameter per Cluster (Online Survey)

participants' current locations. The aim of this system is to simulate user experience in a *real* IoT environment as closely as possible. To that end, our research team will collaboratively create realistic scenarios that are mapped to real locations (e.g., through embedded GPS information), and feed the scenarios into the survey system. Participants will then be asked to walk around in designated areas while carrying the device. As participants move toward a certain location, the device will display an IoT scenario description related to this location. Then, participants will answer questions about their privacy preferences for the presented scenario. We will use the same questionnaires from our online study to investigate whether any discrepancies in the responses of the participants exist between the online survey and the situated survey.

3.2.2 Campus-wide Situated Survey Study

In this study, we conducted a situated survey study in order to better understand people's privacy preferences in IoT. To be specific, we collected people's decisions and opinions re-

garding their privacy in diverse privacy-invasive scenarios in simulated IoT environments, based on 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 ($N = 172$) 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.

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 ($K = 4$) 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 better understand how contextual factors influence people's behaviors and perceptions toward their privacy in IoT environments.

3.2.2.1 Data Collection

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 toward 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 decision-making [39, 76]. In addition, location has been found to be a particularly critical component in understanding people's privacy preferences on diverse application scenarios [85, 17, 16, 13].

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 the previous online survey study (see Chapter 3.2.1), 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*). Each scenario can be described by an expression that includes every contextual parameter together with its respective parameter value for this scenario (see Table 3.1). 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*).

In this study, we shared the same definitions of contextual and reaction parameters with Chapter 3.2.1 while only revising the values of *_notification* and *_permission* reaction parameters to improve the readability of survey questions and answer options (see

Parameter (id)	Values
_notification (R1)	(1) notify me, always (2) notify me, just this time (3) don't notify me, just this time (4) don't notify me, always
_permission (R2)	(1) allow, always (2) allow, just this time (3) reject, just this time (4) reject, always
_comfort (R3)	Same as Table 3.2
_risk (R4)	Same as Table 3.2
_appropriateness (R5)	Same as Table 3.2

Table 3.5: Revised Reaction Parameters of Privacy Preferences

Table 3.5).

Scenario Generation In our earlier online survey described in Chapter 3.2.1, 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 MyMaps[4], which lets multiple users create and update a custom Google Map. As shown in Figure 3.7, we created landmarks with GPS coordinates and associated scenario descriptions containing all five contextual parameters. We aimed to

[4] https://www.google.com/mymaps

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.

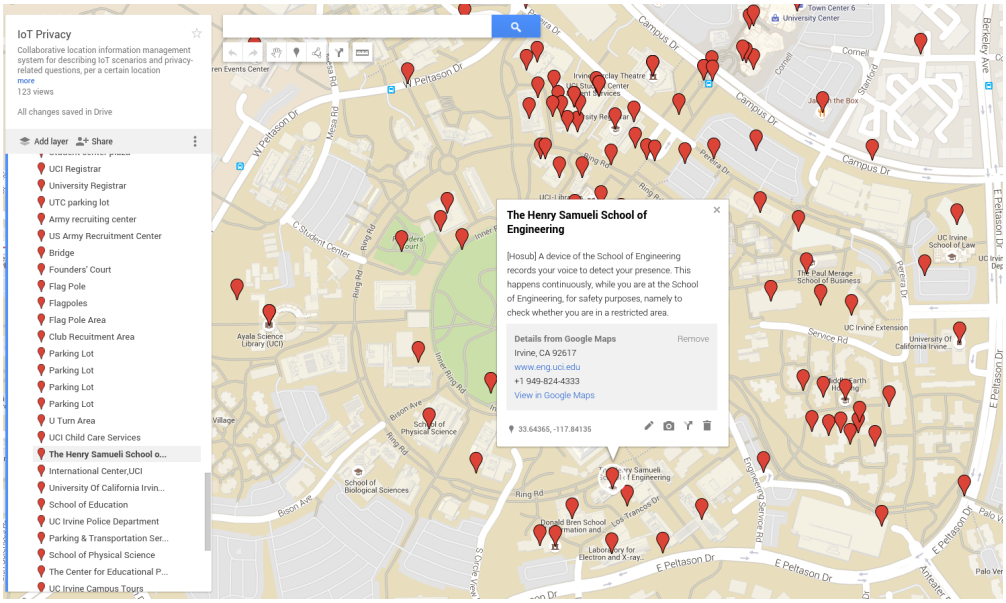

Figure 3.7: Collaborated Scenario Generation via Google MyMaps

As Google MyMaps 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 3.6). Note that we also transformed each scenario description into a specific sequence of values of the contextual parameters (see the last row of Table 3.6) to make it analyzable by data mining and machine learning algorithms. For the parameters *where* ($\mathtt{C}_1$) and *who* ($\mathtt{C}_3$), we then replaced their written values with categorical values defined in Table 3.1. For instance, the School of Engineering is mapped with $\mathtt{C}_1 = 3$ because this place is considered as a public place.

Location-based Scenario Display for Google Glass To operationalize our study, we designed and developed a novel Google Glass application named `IoT Privacy` that

Attribute	Value
Location Name	School of Engineering
GPS (Latitude)	-117.841359
GPS (Longitude)	33.643657
Scenario Description	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
Contextual Parameters	$\{C_1 = 3, C_2 = 9, C_3 = 6, C_4 = 1, C_5 = 1\}$

Table 3.6: Attributes and Values of Sample IoT Scenario

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 Figure 3.8, 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.

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 Google 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 3.5 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 Google Glass, participants would otherwise have had to navigate

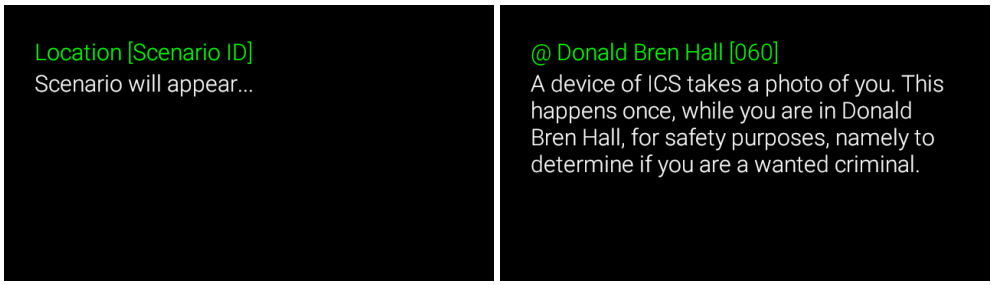

Figure 3.8: Google Glass App `IoT Privacy` Screenshot

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 five months (from May to September 2016): 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 ($SD = 14.72$). 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.

3.2.2.2 Cluster Analysis

We used the same experimental settings of K-modes clustering as described in Chapter 3.2.1.2. Regarding the optimal number of cluster, we found 4 as optimal through the Elbow method ($K = 4$).

Table 3.7 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* ($\mathtt{C_2}$) and *who* ($\mathtt{C_3}$). Each mode has a unique categorical

Mode (Cluster)	Contextual Param. $\{C_1, C_2, C_3, C_4, C_5\}$	Reaction Param. $\{R_1, R_2, R_3, R_4, R_5\}$	Label (Number of Instances)
$M_1(CL_1)$	3, **2**, **6**, 4, 0	1, 1, **6**, **6**, **6**	*Acceptable* (2,608/6,618)
$M_2(CL_2)$	2, **16**, **5**, 2, 0	1, 2, **4**, **4**, **4**	*Neutral* (1,199/6,618)
$M_3(CL_3)$	3, **20**, **3**, 4, 0	1, 4, **3**, **3**, **3**	*Smwht. Unacceptable* (1,492/6,618)
$M_4(CL_4)$	3, **17**, **7**, 3, 1	1, 4, **1**, **1**, **1**	*Very Unacceptable* (1,319/6,618)

Table 3.7: Modes of Clustered IoT Scenarios (Situated Survey)

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* (C_5). These results are consistent with our previous cluster analysis performed on online survey data (see Chapter 3.2.1.2).

We can now mark the clusters using these parameters in a similar way to our previous cluster analysis. We labeled scenarios belonging to the cluster CL_1 as *acceptable* to the study participants as its mode M_1 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. This indicates that Amazon MTurk participants (12.6% *acceptable* and 40.8% *very unacceptable* clustered scenarios) are more sensitive to their privacy compared to university students, as claimed in [48].

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 tests confirm that the difference in the means of the R_3 parameter between each pair of the clusters is statistically significant ($p < .016$, Bonferroni-corrected for three comparisons). Next, we also conducted Welch's t-tests on the remaining 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.

3.2.2.3 Comparison of Clusters based on Contextual Parameters

In this section, we compared 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.

***where* parameter** Regarding the *where* parameter (see Figure 3.9), participants consider monitoring activities as very unacceptable if they occur at their own private places like home ($where = 0$, see $\texttt{CL}_4$; $p < .0001$, Cohen's $d = .6069$). 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 [23]. In contrast, participants consider monitoring that occurs at public spaces as acceptable ($where = 3$, see $\texttt{CL}_1$; $p = .000113$, $d = .2016$). As for semi-public spaces ($where = 2$) like a restaurant, participants' attitude is somewhat neutral (see $\texttt{CL}_2$) since it can be perceived as both a personal and a public place, depending on other contextual factors like *what* and *who*.

***what* parameter** In regard to the *what* parameter (see Figure 3.10), participants do not allow situations in which someone is videotaping them without a clear purpose ($what = 17$, see $\texttt{CL}_4$; $p < .0001$, $d = .804$) or monitoring their eye movements to figure out what they are looking at ($what = 24$, see $\texttt{CL}_4$; $p < .0001$, $d = .6539$). In this context, participants also consider video monitoring as somewhat unacceptable even if it has some purpose ($what = 20, 22$, see $\texttt{CL}_3$; $p < .0001$, $d = .7449$). Photo-taking ($what = 11$, see $\texttt{CL}_1$) is relatively more acceptable to the participants than video monitoring ($what = 17$, see $\texttt{CL}_1$); however, they still worry about this activity if it aims to detect their

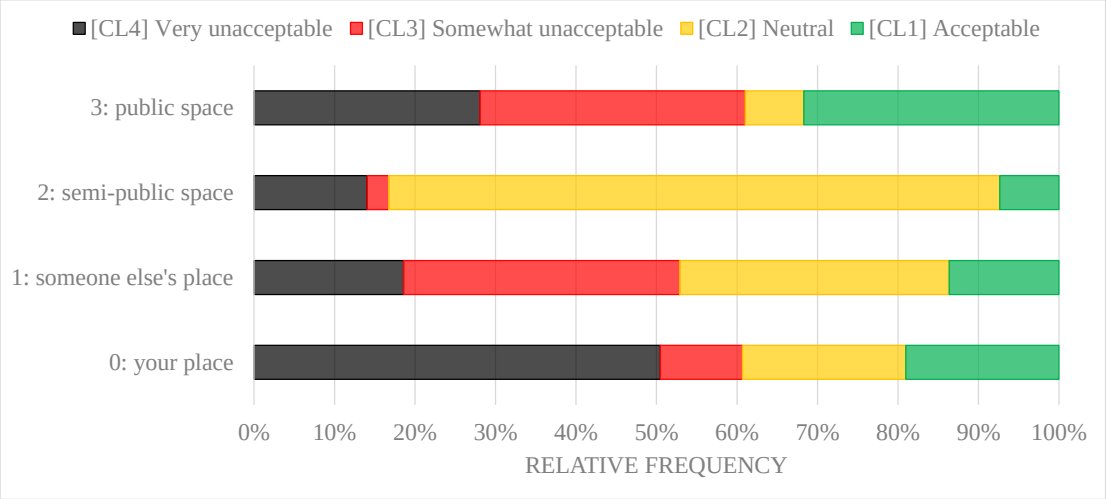

Figure 3.9: Relative Distribution of *where* Parameter per Cluster (Situated Survey)

personal information like age ($what = 13$, see $\texttt{CL}_4$). 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 $\texttt{CL}_1$; $p < .0001$, $d = .9571$), presumably because they perceive this information not to directly represent their personal behavior.

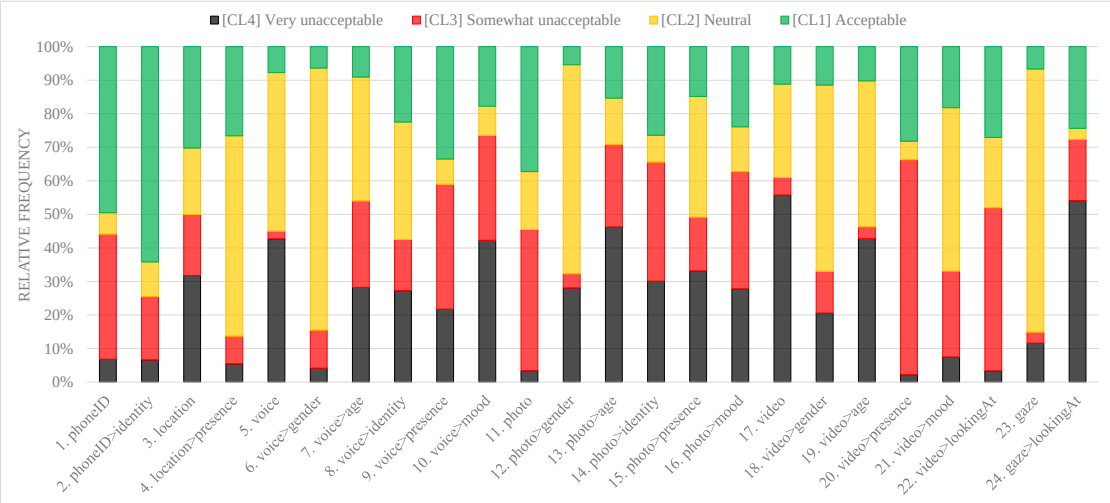

Figure 3.10: Relative Distribution of *what* Parameter per Cluster (Situated Survey)

***who* parameter** In previous studies, we found that the identity of the information requester is an important determinant of people's privacy preferences or decisions on various information monitoring activities [59, 26]. Through the present cluster analysis (see Figure 3.11), 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 $\texttt{CL}_4$; $p < .0001$, $d = 1.1071$), or if it is the government ($who = 7$, see $\texttt{CL}_4$; $p < .0001$, $d = 1.0858$). Participants also have some privacy concerns if their school/employer ($who = 6$, see $\texttt{CL}_3$; $p < .0001$, $d = .6562$) 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 $\texttt{CL}_1$; $p < .0001$, $d = .9128$). These responses are probably because most participants were students who typically trust what their school does.

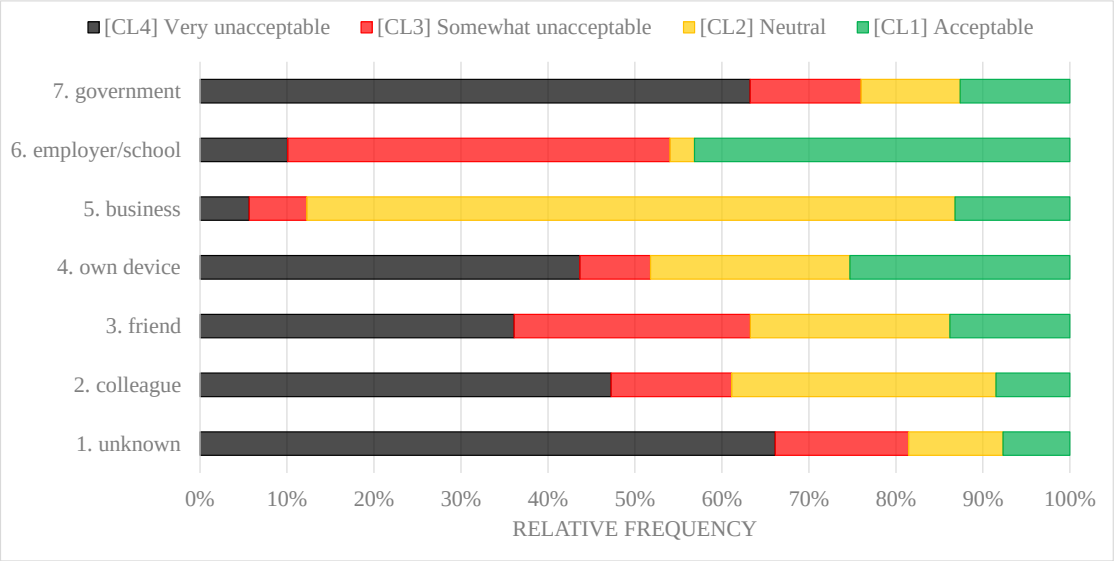

Figure 3.11: Relative Distribution of *who* Parameter per Cluster (Situated Survey)

***reason* parameter** The monitoring purpose can take on one of six values: *safety*, *commercial*, *social*, *convenience*, *health*, and *not specified* (see Figure 3.12). Participants consider monitoring as very unacceptable when it is performed for social- ($reason = 3$, see $\texttt{CL}_4$; $p < .0001$, $d = .9691$) or safety-related purposes ($reason = 1$, see $\texttt{CL}_4$; $p < .0001$,

$d = .6245$). 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 $\texttt{CL}_1$; $p < .0001$, $d = .6089$). In addition, *convenience* is also a reasonable justification (*reason* = 4, see $\texttt{CL}_1$; $p < .0001$, $d = .9004$).

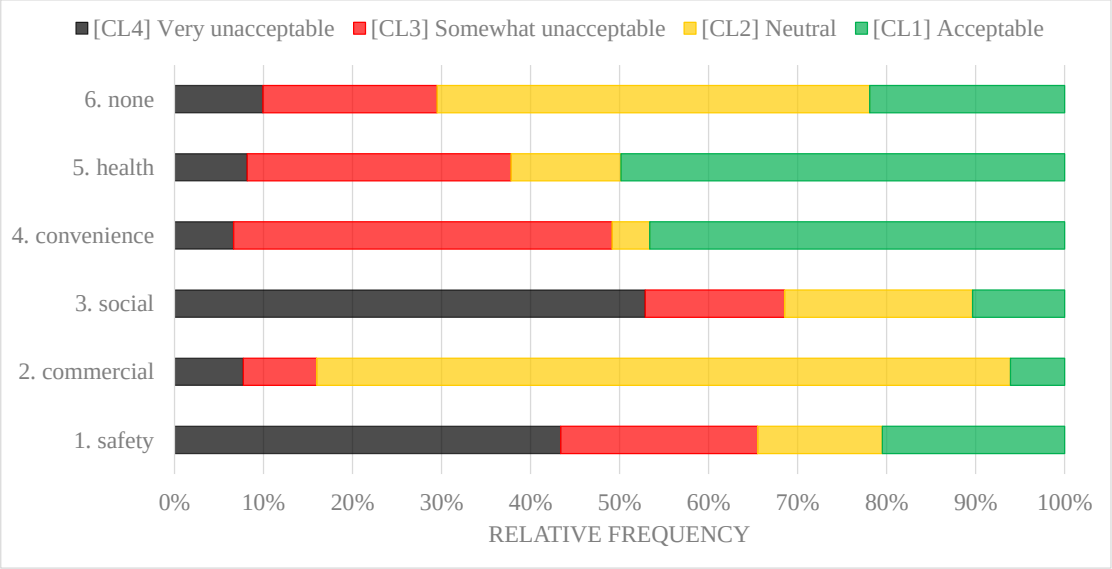

Figure 3.12: Relative Distribution of *reason* Parameter per Cluster (Situated Survey)

***persistence* parameter** 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 Figure 3.13). Participants are clearly unwilling to accept scenarios with continuous monitoring of personal data and/or information (*persistence* = 1, see $\texttt{CL}_4$; $p < .0001$, $d = .7252$). In contrast, one-time monitoring is generally acceptable to them (*persistence* = 0, see $\texttt{CL}_1$; $p < .0001$, $d = .3842$).

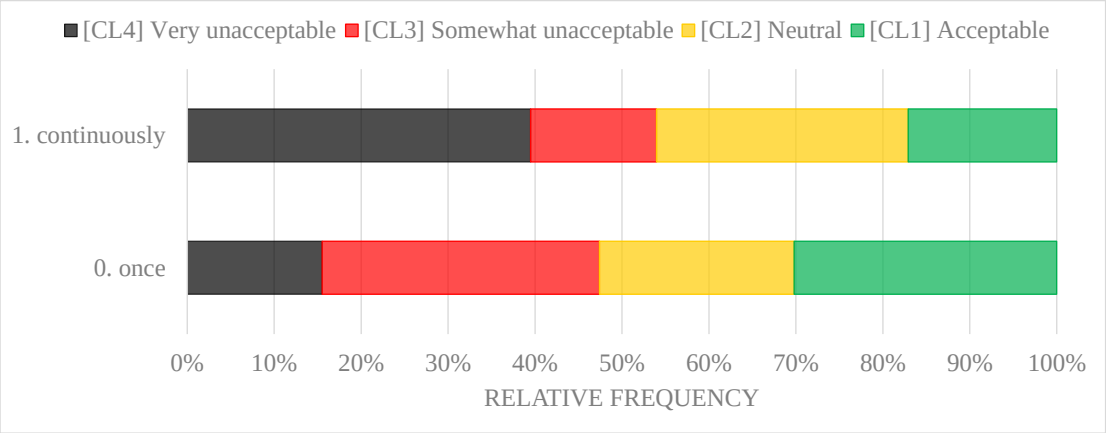

Figure 3.13: Relative Distribution of *persistence* Parameter per Cluster (Situated Survey)

3.2.2.4 Limitations

We showed that it is feasible to group privacy scenarios into clusters with distinct user reactions, thereby interpreting people's privacy perceptions regarding IoT. Yet, our work still has some issues that need to be considered and addressed.

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 Google 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.

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 [3, 45, 30] 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 behaviors in the wild.

3.3 Conclusion

In Chapter 3.2.1, we showed that IoT scenarios can be grouped into four clusters in terms of their potential privacy risks. By comparing these clusters according to the five contextual parameters, we also extracted some contextual factors that cause users' privacy concerns in IoT. To verify whether our findings are also applicable to broader IoT contexts, we planed to design and develop a location-based (situated) survey system. One of our research hypotheses was that a situated survey would be more suitable for collecting genuine responses from users than a traditional online survey, since it situates users in IoT scenarios that are comparatively more realistic than clicking through a survey in a web browser.

In Chapter 3.2.2, we also investigated how people's privacy decision-making in IoT environments can be modeled and interpreted. Specifically, 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 (i.e., situated survey). We then performed a cluster analysis on the collected preferences in order to understand users' privacy concerns toward IoT applications and services. The results of the analysis showed that IoT scenarios can be grouped into four distinct clusters in terms of their perceived privacy risks, just as we figured out in the previous online survey study. By comparing the resulting clusters, we also extracted a number of contextual factors causing privacy threats in IoT.

Chapter 4

Prediction of Privacy Decisions in IoT

In this chapter, we build machine learning (ML) models to predict users' privacy decisions, using the dataset collected in our previous situated survey study. The goal of this chapter is to explore ML mechanisms to predict users' decisions whether or not to allow the particular IoT service scenario, with reasonable predictive performance. To show the feasibility of privacy decision prediction, we first train decision tree models based on contextual information embedded in the scenarios as well as its cluster membership information (i.e., clustered *scenarios*). Next, we perform comprehensive experiments to figure out the best approach for predicting privacy decisions, in terms of ML algorithm, input feature, and model training strategy. Regarding this, we perform privacy segmentation that aims to cluster similar users according to their notion of privacy and utilize segmentation results (i.e., clustered *users*) as an additional feature for privacy decision prediction.

4.1 Introduction

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 of their available time, motivation, and their cognitive decision-making abilities [2, 92]. Therefore, many researchers proposed various mechanisms to predict users' privacy decisions via machine learning models trained on a subset of users' prior privacy behaviors [31, 85, 36, 15, 17, 105, 16, 65]. Software agents can then use these machine learning models to give users personalized privacy recommendations, thereby assisting them to better control their privacy (i.e., privacy decision support).

Privacy decision support technology is going to become more important in the Internet of Things (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 their preferred decisions to the services. It is therefore necessary to investigate whether it will be possible to learn and predict users' privacy decisions in such IoT environments as well as to recommend the predicted decisions as the optimized *default* decisions for the current IoT context. Using the dataset we already collected in our situated survey study (see Chapter 3.2.2), we performed a series of machine learning experiments to show not only the feasibility but also future directions of privacy decision prediction in IoT environments.

4.2 Related Work

We presented a literature review of privacy decision prediction in Chapter 2.2. In this section, we summarized previous research about privacy segmentation. It is understood that privacy segmentation provides useful information for understanding and predicting people's privacy

decision-making.

4.2.1 Privacy Segmentation for Privacy Decision Prediction

Researchers have investigated methodologies to segment users into several categories in terms of their privacy attitudes and behaviors. The most commonly cited methodology is Westin's privacy segmentation model [55]. Westin had conducted several surveys about privacy issues in various domains such as e-commerce, national identification systems, and e-health. To effectively summarize the survey results, Westin developed an indexing scheme that categorizes survey participants into three categories: *privacy fundamentalists*, *privacy pragmatists*, and *privacy unconcerned*. Westin treated participants' responses to several pre-defined statements (scenarios) as criteria to derive these three categories (e.g., *privacy fundamentalists* are respondents who agreed with the first statement and disagreed with the second and third statement).

Privacy segmentation has been studied in diverse contexts. Lin et al. proposed an unsupervised data clustering approach for categorizing smartphone users into distinctive groups based on their privacy preferences regarding mobile app permission (e.g., grant or deny permission to an app to access personal information) [64]. The authors utilized an agglomerative hierarchical clustering algorithm (Ward's method) on about 21,000 privacy preferences collected from Android users ($N = 725$). Each preference represents each user's willingness to grant permission to a given app for a specific purpose (i.e., *app-permission-purpose* triple). The authors identified four privacy profiles from this cluster analysis. They also presented *default* privacy settings to each user based on his/her privacy profile. This was intended to help users better control their privacy when confronted with numerous permission requests on Android platforms. In a follow-up study, the authors utilized users' privacy profile information as one of the input features for machine learning models to predict Android users'

permission settings [65].

Lankton et al. clustered SNS users into four categories based on their privacy management strategies [57]. They surveyed college students' behavior on Facebook, including the use of privacy settings, degree of content disclosure, and variety and size of friend lists. The authors then conducted a two-stage cluster analysis on this dataset. Like [64], the authors first performed hierarchical clustering to determine both the correct number of clusters and initial cluster centroids. They found that a four-cluster solution is optimal. Next, they conducted non-hierarchical (K-means) cluster analysis on the dataset, using the pre-determined cluster centroids as a starting point. After statistically comparing survey responses in each cluster, the authors confirmed that the resulting clusters are distinctive enough regarding the degree of the users' privacy concerns.

Most recently, the market research firm Forrester published a report suggesting that consumers can be divided into four privacy categories: *data-savvy digitals*, *reckless rebels*, *nervous nellies*, and *skeptical protectionists*[1]. This finding is based on large-scale online survey studies designed to capture people's behavioral reactions toward personal data collection and use by Internet companies. About 34% of the study participants were categorized as those who are not willing to share their information (*nervous nellies* and *skeptical protectionists*) since they are skeptical about corporate privacy practices.

Even though the number and type of privacy segments vary somewhat across these works, most researchers came to the same common conclusion about privacy segmentation: it is practically feasible to identify distinctive privacy segments by collecting and analyzing human behavioral data. Furthermore, privacy segment information can be used as an informative feature for understanding and predicting people's future privacy behavior because it represents users' overall perception of privacy [64, 65].

[1] `https://go.forrester.com/blogs/its-here-forresters-consumer-privacy-segmentation/`

4.3 Context-based Privacy Decision Prediction

Using the dataset described in Chapter 3.2.2, we tried to predict participants' privacy decisions by learning conditional inference trees, using the gathered survey responses as training data. We utilized the five contextual factors as well as clustering results as input features for predicting how participants will have privacy decisions in the presented scenarios. The final trained model has a reasonable accuracy in predicting whether or not participants will allow personal information monitoring in a given IoT scenario.

4.3.1 Experimental Setup

Here, we aimed to predict participants' response to the following question:

> "If this situation [= scenario] happens, would you want to *allow* it?" ($\mathtt{R}_2$)

using contextual parameter values and cluster membership of the scenarios as input features. We focus on the reaction parameter *_permission* ($\mathtt{R}_2$) because it may reflect people's substantive privacy decisions in IoT environments. We utilize a conditional inference tree algorithm for building machine learning models.

Conditional inference tree (CTree) is a statistics-based decision tree learning algorithm that uses non-parametric tests as splitting criteria [41]. 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 ($p < .05$) 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 over-fitting but guarantee unbiased predictor selection. We used `party`, an R implementation of the CTree algorithm, for training CTree decision tree models on our dataset.

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* ($\mathtt{R}_2$) for the presented IoT scenarios. Among the attributes of our dataset, we chose the five contextual parameters, *where* ($\mathtt{C}_1$), *what* ($\mathtt{C}_2$), *who* ($\mathtt{C}_3$), *reason* ($\mathtt{C}_4$), and *persistence* ($\mathtt{C}_5$), for specifying a basic feature vector for the classifiers. We saw in Chapter 3.2.1.3 and 3.2.2.3 that all these parameters influence people's privacy decision-making. We then added cluster membership ($\mathtt{CL}_\mathtt{K}$), 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 $\mathtt{R}_2$ (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 [17, 16, 65]. Therefore, we also build and evaluate CTree models as binary classifiers by converting $\mathtt{R}_2 = 1, 2$ into *allow* and $\mathtt{R}_2 = 3, 4$ into *reject*.

4.3.2 Experiment 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 [4]. For the binary classification, we also measured the F1 score to consider both precision and recall for the classification results.

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**

Table 4.1: Privacy Prediction Performance (CTree)

Table 4.1 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 [17, 16, 65] are not directly comparable because each work uses different datasets 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 decisions of users in IoT if we can extract and model privacy-related contexts from the the user's environments.

4.4 Context- and User-based Privacy Decision Prediction

We performed an in-depth study to predict privacy decisions of users in IoT environments, through data mining and machine learning techniques. To construct predictive models, we tested several different machine learning models, combinations of features, and model train-

ing strategies on human behavioral data collected from our situated survey study. Experimental results showed that a machine learning model called linear model and deep neural networks (LMDNN) outperforms conventional methods for predicting users' privacy decisions for various IoT services. We also found that a feature vector, composed of both contextual parameters and privacy segment information, provides LMDNN models with the best predictive performance. Lastly, we proposed a novel approach called *one-size-fits-segment* modeling, which provides a common predictive model to a segment of users who share a similar notion of privacy. We confirmed that one-size-fits-segment modeling outperforms previous approaches, namely individual and one-size-fits-all modeling. From a user perspective, our prediction mechanism takes contextual factors embedded in IoT services into account and only utilizes a small amount of information polled from the users. It is therefore less burdensome and privacy-invasive than the other mechanisms. We also discussed practical implications for building predictive models that make privacy decisions on behalf of users in IoT.

4.4.1 Summary

In this study, we proposed a novel machine learning mechanism for predicting privacy decisions of users in IoT environments. The aim of this mechanism is to correctly predict users' decisions whether or not to allow the given personal information monitoring based on both the user's current context and personal attitudes on privacy. We tested the proposed mechanism on a privacy-related behavioral dataset collected from human subjects ($N = 172$) who were presented with descriptions of personal information tracking scenarios relating to their physical location on a university campus (see Chapter 3.2.2).

We treated the five contextual parameters as basic features which collaboratively represent the current context in which information monitoring is performed by IoT devices. Addition-

ally, we assigned each user to a specific privacy perception segment based on a small portion of their data (i.e., prior privacy decisions). We then used this privacy segment information as an additional feature. We used the reaction parameter *_permission* ($\mathtt{R}_2$) as target value (or label), since it best reflects users' substantive privacy decisions in IoT environments (namely, allow or reject the monitoring).

First, we utilized a state-of-the-art machine learning model, called linear model and deep neural networks (LMDNN, [22]), to make privacy decisions on behalf of users. LMDNN, which is also known as Wide & Deep Learning, jointly trains wide linear models and deep neural networks. By doing so, it can take the benefits of memorization (linear models) and generalization (neural networks) at the same time. LMDNN showed a remarkable performance on binary classification problems with sparse input features [22, 89]. Because our dataset is also composed of categorical data with many possible feature values (e.g., 24 values for the contextual parameter *what*), we decided to use LMDNN to build predictive models for privacy decision support in IoT. We also selected machine learning models that have been widely used in the literature (e.g., decision trees) and compared them with LMDNN in terms of predictive performance on the dataset. Experimental results indicated that LMDNN outperforms all the conventional models.

Next, we explored the most suitable combination of features for building LMDNN models with a reasonable predictive performance. We chose the five contextual parameters as basic features because these parameters are known to be related to users' privacy decisions in IoT (see Chapter 3). In addition, we considered each user's privacy segment information as an additional feature. This is because previous research indicates that privacy segment information is helpful for machine learning models to better predict users' privacy decisions [64, 65]. We applied an unsupervised data clustering algorithm, K-modes clustering, on a subset of users' privacy decisions, in order to segment the users by their perceptions of privacy (i.e., privacy segmentation). Therefore, we tested the following feature combinations:

(1) contextual parameters only (basic features), (2) contextual parameters with interactions between them, and (3) contextual parameters with interactions between them and privacy segment information. Experimental results showed that the feature combination (3) gives LMDNN models with the highest predictive performance. It also means that both interactions between contextual factors (e.g., *who* by *what*) and privacy segment information are useful for LMDNN models to predict privacy decisions of the users.

Lastly, we investigated different approaches for training machine learning models. There exist two traditional approaches in the literature: individual and one-size-fits-all modeling. Individual modeling is a process of building a user-specific predictive model based on each user's data only. This approach is known to be effective for modeling each user's unique characteristics (e.g., habits and personality). A model with a reasonable predictive performance, however, typically requires a considerable amount of training data from each individual user. In contrast, one-size-fits-all modeling utilizes multiple users' data as a single training data and constructs a universal model for all of them (including new users). This approach enables predictive models to make general predictions, therefore it can be useful for new users who did not provide data to the system yet, but want to get recommendations immediately. However, prediction results may not be personalized to each individual user. We presented another approach called *one-size-fits-segment* modeling, a variant of one-size-fits-all, taking privacy segment information into account in building predictive models. The basic idea is to serve each user with a machine learning model trained by data collected from others who share the same notion of privacy with this user. We divided the dataset based on the results of privacy segmentation, then trained an LMDNN model for each segment of the users. By using both contextual and privacy segment information as input features, we compared the predictive performance of individual, one-size-fits-all, and one-size-fits-segment modeling. Experimental results confirmed that the proposed approach performs the best in the dataset. Final LMDNN models trained via one-size-fits-segment modeling showed an average area under curve (AUC) of 0.6782 across all users. We noticed that one-size-fits-segment

modeling performs much better than individual modeling for about 80% of the users. However, it does not work well for about 20% of the users who have highly accurate individual models ($AUC > 0.7$).

To sum up, our proposed prediction mechanism not only showed a reasonable performance for most users, but also can cause less burden and privacy risks to users since it mainly utilizes non-personal contextual information which can often be automatically collected from the IoT environment, and only prompts each user for a small amount of privacy decisions (answers to five reaction parameters about a single scenario) to determine his/her privacy segment. We also presented some practical implications for designing and developing machine learning-based privacy decision support systems for IoT.

In summary, our work makes the following contributions to the field of privacy decision support:

- We adopted a state-of-the-art machine learning model called linear model and deep neural networks (LMDNN) to make privacy decisions on behalf of IoT users, and verified that LMDNN outperforms conventional methods that have been widely used in the literature.
- We investigated the best approach for training machine learning models in terms of input feature and training strategy, and verified that the proposed approach called *one-size-fits-segment* modeling outperforms preexisting methods such as individual and one-size-fits-all modeling.
- We reported some practical implications in predicting users' privacy decisions on personal information monitoring in IoT.

4.4.2 Experiment — Overview

By using the dataset described in Chapter 3.2.2, we investigated mechanisms to learn and predict users' privacy decision-making in IoT environments. Specifically, we aimed to predict the value of the reaction parameter *_permission* ($\mathtt{R}_2$) based on the current context and privacy segment of the user. Even though this parameter denotes four possible privacy decisions (see Table 3.2), we focused on binary decisions by converting $\mathtt{R}_2 = 1, 2$ into $\mathtt{R}_2 = 1$ (*allow*) and $\mathtt{R}_2 = 3, 4$ into $\mathtt{R}_2 = 0$ (*reject*). Therefore, prediction for this parameter can be formalized into binary classification problems.

In this vein, we conducted a series of machine learning experiments, varying the models (algorithms), features, and training strategies. First, we tested multiple machine learning algorithms to find the most suitable for making predictions about privacy decisions in IoT. We first tried LMDNN since it is known to be very effective for processing categorical data with high sparsity. We then compared the performance of LMDNN with several machine learning algorithms that have been extensively used in earlier research. Thereafter we also assessed the impact of input features, consisting of contextual and privacy segment information, on the predictive power of the trained machine learning model. This assessment allowed us to determine which features should be used for building classifiers. Last, and most importantly, we conducted a comparative evaluation of the predictive performance of well-known model training strategies, such as individual and one-size-fits-all modeling, and our proposed approach called *one-size-fits-segment* modeling. The results informed us of practical implications for developing a privacy decision support system for IoT environments.

4.4.3 Experiment — Machine Learning Algorithm

In this section, we explained in detail why we chose the LMDNN machine learning model for realizing privacy decision support in IoT. We also presented experimental results showing

that LMDNN can provide the most reasonable predictive performance compared to conventional machine learning algorithms used in the literature.

4.4.3.1 LMDNN

Linear model and deep neural networks (LMDNN), also known as Wide & Deep Learning, has been proposed by [22] to solve the problem of recommending apps on Google Play. Generalized linear models like logistic regression are widely used for large-scale regression and classification problems as they are simple, scalable, and interpretable. The models are often trained on binarized sparse input features with one-hot encoding. Memorization of diverse feature interactions can be efficiently achieved by feeding a wide set of cross-product feature transformations with a target value (or label) into the model. While linear models are effective for learning relationships between categorical features and a target value, they cannot generalize the relationships to identify feature-target patterns that do not exist in training data. In contrast, deep neural networks (DNN) can better generalize to the previously unseen patterns by using low-dimensional dense embedding vectors learned from the sparse input features (i.e., transforming a categorical feature value into a vector of continuous values). This means that DNN can make a reasonable prediction for new observations based on preexisting training data. At the same time, DNN also can over-generalize when the underlying feature-target matrix is too sparse (e.g., rare interactions between features and target value). LMDNN is a mixture of logistic regression (wide learning) and DNN (deep learning) to achieve both memorization and generalization in a single model. By jointly performing wide and deep learning, LMDNN complements the weakness of deep learning (i.e., over-generalization) by letting the wide learning take some cross-product feature transformations into account in producing final classifications. This Wide & Deep Learning paradigm shows a remarkable predictive performance on diverse classification problems [22, 19, 89].

Figure 4.1 shows the LMDNN model structure we used in this study. The bottom (input)

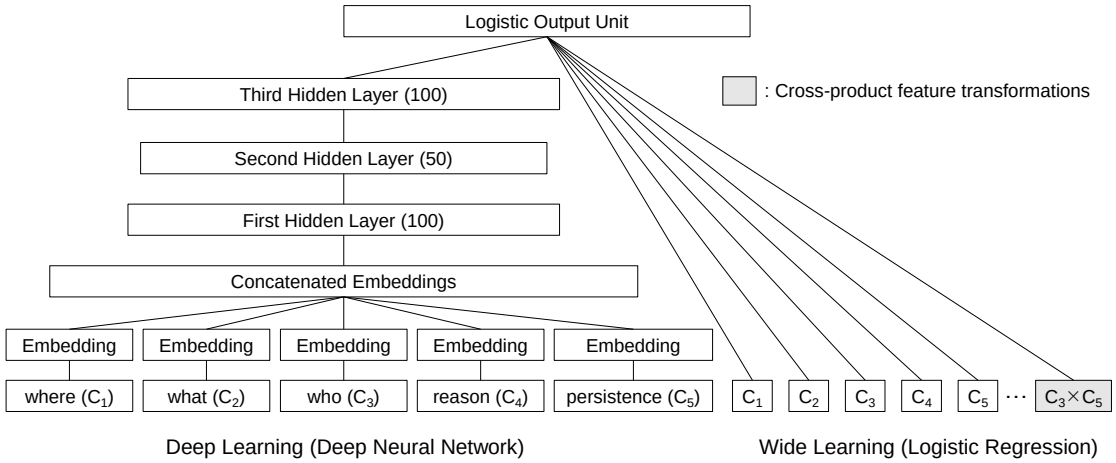

Figure 4.1: Architecture of Wide & Deep Machine Learning Model

layer receives training data composed of categorical features (contextual parameters) together with a target value (binarized reaction parameter *_permission*). Next, the model can generate cross-product feature transformations and/or dense embedding vectors from the inputted data. As explained above, feature transformations and embedding vectors are used by wide and deep learning, respectively. For the deep learning, an 8-dimensional embedding vector is learned from each categorical feature. The model also combines all embedding vectors into a single dense embedding vector. The resulting concatenated embedding vector is then fed into three hidden layers with the ReLU activation function, and finally the logistic output unit (here, the sigmoid function). We configured 100, 50, and 100 units in consecutive hidden layers, respectively. We determined this network structure based on internal performance benchmarking on our dataset. In training LMDNN models, we followed the default mechanism (e.g., backpropagation, mini-batch stochastic optimization, AdaGrad regularization) described in [22]. We utilized a `TensorFlow` [1] implementation of LMDNN for all our experiments.

4.4.3.2 Machine Learning Algorithm Benchmark

To verify whether the LMDNN is suitable for making predictions based on many categorical input features, we compared its predictive performance on our dataset against other conventional machine learning algorithms: recursive partitioning tree [95], conditional inference tree (CTree) [41], random forests and bagging ensemble based on CTree (conditional random forests), SVM, Naïve Bayes, and logistic regression. We chose these models because they are widely used in privacy decision support systems based on machine learning (e.g., decision trees [36, 87], random forests [85], and SVM [16]). For each machine learning model, we measured its predictive performance on the dataset, through 10-fold cross validation (CV). We only utilized the five contextual parameters as input features in these experiments, without cross-product feature transformations for memorization. This is because we first needed to assess the generalization capability of these models since the dataset is small (6,618 rows), thereby potentially leading to over-fitting. Therefore, we evaluated a deep model of LMDNN (see the left side of Figure 4.1) in this experiment[2]. Regarding a performance metric, we primarily used the area under the ROC curve (AUC) because it is unaffected by the class imbalance problem (there were 64% allow and 36% reject decisions in the dataset) and is independent of the threshold applied to compute the probability of the binary classification results. Additionally, AUC itself is comprehensible; a random classifier has an AUC score of 0.5 while a perfect classifier has an AUC score of 1.0.

Experimental results indicated that LMDNN outperforms all other models (see Table 4.2). However, the performance difference is not large (up to 3%). Conditional random forests yield competitive performance because it has proven effective at generalizing to variants not seen in the training set [12], just as deep neural networks do. As mentioned before, however, LMDNN is known to further enhance the deep model by efficiently memorizing feature-target

[2]The `TensorFlow` implementation of LMDNN provides an API that enables programmers to selectively configure a wide, deep, or wide and deep model.

Machine Learning Algorithm	AUC
Recursive Partitioning Tree	0.6142
Conditional Inference Tree	0.6293
Conditional Random Forests	0.6353
Support Vector Machine	0.6107
Naïve Bayes	0.6161
Logistic Regression	0.6208
Deep Neural Networks (of LMDNN)	**0.6421**

Table 4.2: Comparison of Machine Learning Algorithms in Privacy Prediction

patterns that are rarely observed in a sparse dataset (wide learning). For this distinctive feature, we decided to utilize LMDNN as the machine learning model for this study.

4.4.4 Experiment — Feature Engineering

Here we explained how we identified the most useful features for training LMDNN models. Specifically, we presented the reasons why we chose the five contextual parameters with interactions between them as underlying input features. In addition, we described how we conducted privacy segmentation on the human subjects in a previous situated survey study and why we utilized the resulting privacy segment information as an additional feature.

4.4.4.1 Contextual Information

We decided to use the five contextual parameters as basic features for the following reasons: (1) our previous research indicated that all these contextual parameters impact people's privacy decision-making (see Chapter 3), and (2) we aimed to make predictions based on contextual information which can be automatically collected by IoT environments, thereby avoiding as much as possible asking users to manually enter additional information. We also considered interactions between the contextual parameters because people's privacy decisions about specific contextual information could be influenced by other factors. For instance, the

monitoring of personal information (e.g., face photos: $\mathtt{C}_2 = 11$) can be perceived differently depending on who is performing this monitoring (e.g., unknown vs. employer/school: $\mathtt{C}_3 = 1$ vs. 6). Furthermore, LMDNN models can make predictions for unusual feature-target patterns by considering cross-product feature transformations (i.e., memorization). Because the contextual parameters *what* ($\mathtt{C}_2$) and *who* ($\mathtt{C}_3$) have the most significant influence on people's privacy decisions, we determined that we should feed feature transformations based on *what* and *who* parameters (e.g., $\mathtt{C}_2 \times \mathtt{C}_1$, $\mathtt{C}_2 \times \mathtt{C}_3$, $\mathtt{C}_2 \times \mathtt{C}_4$, $\mathtt{C}_2 \times \mathtt{C}_5$) into the LMDNN models.

4.4.4.2 Privacy Segment

As discussed before, privacy segmentation is known to provide useful information for understanding and predicting people's privacy behavior. By applying the K-modes clustering algorithm on our dataset, we identified distinctive privacy segments and then assigned each user into one of the segments. We had already performed cluster analysis on the same dataset (e.g., see Chapter 3.2.2.2). However, in that study, we clustered IoT *scenarios* in terms of users' privacy concerns about contextual factors (e.g., *what* and *who*). In this study, we aimed to cluster *users* into privacy segments based on each user's expectation of privacy in a single IoT scenario.

To that end, we selected scenario #60[3] as a base scenario that all participants had responded to, and filtered the dataset by considering the base scenario only. All participants responded to scenario #60 because it is related to the location where each experiment started. By analyzing this partial dataset, we expected to understand how individual users perceive and react differently in the same scenario. To determine the optimal number of clusters (K) for this new dataset, we used the well-known Elbow method [68]. First, we computed the sum of errors (SE) of the K-modes clustering with a maximum of 50 iterations, while increasing

[3]A device of ICS ($\mathtt{C}_3 = 6$) takes a photo of you ($\mathtt{C}_2 = 11$). This happens once ($\mathtt{C}_5 = 0$), while you are in DBH ($\mathtt{C}_1 = 3$), for safety purposes ($\mathtt{C}_4 = 1$), namely to determine if you are a wanted criminal.

K from 2 to 10. We repeated this procedure 10 times, and took average values of SE for each value of K. Next, we calculated the values for the mean difference between SE_{K} and $SE_{\mathrm{K\text{-}1}}$, and found that the largest decrease in errors occurred when we increased K from 2 to 3 (see Figure 4.2). Therefore, we chose 3 as a suitable number of clusters ($K = 3$), and used it as a parameter (`modes`) for running the K-modes clustering algorithm on the dataset.

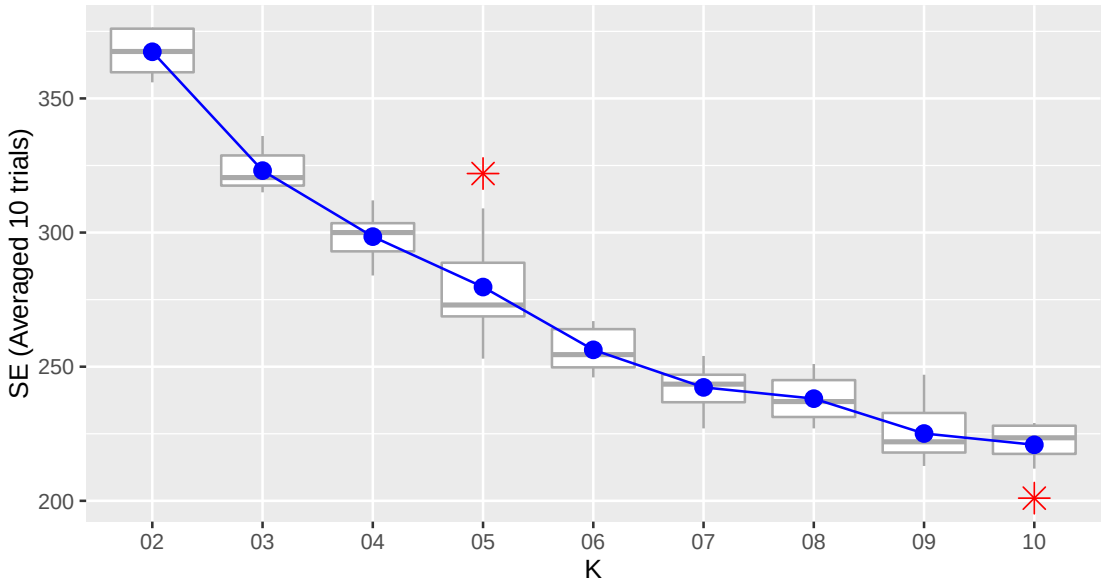

Figure 4.2: Errors in Clustering Users

Table 4.3 summarizes the resulting cluster modes, which are composed of both contextual and reaction parameter values. Note that all contextual parameter values are identical across the clusters because we fixed them to describe the base scenario. The clusters are quite distinct from each other, primarily in the reaction parameters *_comfort* ($\mathtt{R}_3$), *_risk* ($\mathtt{R}_4$, reverse-coded), and *_appropriateness* ($\mathtt{R}_5$): each mode has a unique combination of values for $\mathtt{R}_3$, $\mathtt{R}_4$, and $\mathtt{R}_5$. As shown in Table 3.2, these parameters represent people's privacy attitudes about IoT scenarios on a scale of 1 to 7. For example, cluster mode 3 ($\mathtt{M}_3$) contains $\mathtt{R}_3 = 1$, $\mathtt{R}_4 = 1$, and $\mathtt{R}_5 = 1$, indicating that the given scenario is perceived by participants as *very uncomfortable*, *very risky*, and *very inappropriate*, respectively. For $\mathtt{M}_3$, the value of the reaction parameter *_permission* ($\mathtt{R}_2$) is zero. This means that users belonging to this cluster are likely to reject

Mode	Contextual Param. $\{C_1, C_2, C_3, C_4, C_5\}$	Reaction Param. $\{R_1, R_2, R_3, R_4, R_5\}$	Label (Privacy Segment)
M_1	3, 11, 6, 1, 0	1, **1**, 4, 6, 6	*Indifferent*
M_2	3, 11, 6, 1, 0	1, **0**, 3, 3, 3	*Somewhat Sensitive*
M_3	3, 11, 6, 1, 0	1, **0**, 1, 1, 1	*Sensitive*

Table 4.3: Modes of Clustered Users

the base scenario because they have negative views on privacy in this scenario. Therefore, we marked the clusters (privacy segments) using these three reaction parameters. We labeled privacy segment 1 (PS_1) as *indifferent to privacy* since its mode contains the second highest value for R_4 and R_5 (namely, 6 on a 7-item scale). Likewise, we labeled privacy segment 2 (PS_2) as *somewhat sensitive to privacy* (the values of R_3, R_4, and R_5 fall slightly below the scale average), and privacy segment 3 (PS_3) as *sensitive to privacy.* As a result, 74%, 15%, and 11% of the participants were assigned into *indifferent to privacy*, *somewhat sensitive to privacy*, and *sensitive to privacy* segments, respectively. Finally, we repeated this clustering on additional scenarios (#20, #73, #93, #111)[4] which were the next most frequently visited scenarios after #60. We arrived at the same conclusions regarding the number ($K = 3$) and labels of the resulting privacy segments.

To validate the distinctiveness of the resulting privacy segments, we performed two Welch's t-tests on the R_3 parameter between the following pairs of privacy segments: (PS_1, PS_2) and (PS_2, PS_3). The reason for using Welch's t-test is that all privacy segments 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 segments is statistically significant ($p < .025$, Bonferroni-corrected for two comparisons). Next, we also conducted Welch's t-tests on the R_4 and R_5 parameters and drew the same conclusion. Thereby, we verified that the privacy segments are sufficiently distinct from each other in terms of participants' reactions to the given scenario.

Because privacy segment information for all users was available from such a cluster analysis,

[4]Number of respondents (scenario ID): 140 (#20), 138 (#73), 136 (#93), 162 (#111)

we then utilized it as an additional feature for building predictive models. This is because privacy segment information is known to be useful for quantifying an individual's judgement about privacy and utility in various circumstances [51, 57].

4.4.4.3 Predictive Performance Evaluation

As described before, we treated the five contextual parameters as basic features for training the deep learning part of LMDNN. We then considered interrelated parameters, especially based on the *what* ($\mathtt{C}_2$) and *who* ($\mathtt{C}_3$) parameters, in conducting both wide and deep learning via LMDNN. This is not only because both contextual parameters have a significant impact on people's privacy decisions, but because the memorization of some interactions between parameters (i.e., cross-product feature transformations) can improve the performance of the deep model of LMDNN. As an additional feature, we adopted privacy segment information because it differentiates users according to their perceptions of privacy. We tested the following combinations of features to assess their influence on the predictive performance of LMDNN models.

1. Contextual parameters (*deep* learning)
2. Contextual parameters with interactions (*wide* and *deep* learning)
3. Contextual parameters with interactions and privacy segment information (*wide* and *deep* learning)

Using the whole dataset, we performed 10-fold CV on the LMDNN model trained with each of these feature combinations. As expected, the AUC score gradually improves as we added cross-product feature transformations and privacy segment information to the five basic contextual features (see Table 4.4). We then concluded that both the contextual parameters

Feature Combination	AUC
(1) $\mathtt{C}_1, \mathtt{C}_2, \mathtt{C}_3, \mathtt{C}_4, \mathtt{C}_5$	0.6421
(2) $\mathtt{C}_1, \mathtt{C}_2, \mathtt{C}_3, \mathtt{C}_4, \mathtt{C}_5, \mathtt{C}_2 \times \mathtt{C}_{\{1,3,4,5\}}, \mathtt{C}_3 \times \mathtt{C}_{\{1,4,5\}}$	0.6528
(3) $\mathtt{C}_1, \mathtt{C}_2, \mathtt{C}_3, \mathtt{C}_4, \mathtt{C}_5, \mathtt{C}_2 \times \mathtt{C}_{\{1,3,4,5\}}, \mathtt{C}_3 \times \mathtt{C}_{\{1,4,5\}}, \mathtt{PS}$	**0.6725**

Table 4.4: Privacy Prediction Performance (LMDNN)

(including interactions) and privacy segment information can act as informative features for predicting the binary value of the reaction parameter *_permission* ($\mathtt{R}_2$) through LMDNN.

4.4.5 Experiment — Training Strategy

Based on the selected machine learning model (LMDNN) and features (contextual and privacy segment information), we investigated the best strategy to build a predictive model (i.e., classifier) for each individual user. First, we reviewed two commonly used strategies, individual and one-size-fits-all modeling. Individual modeling typically utilizes a single user's instances as training data, and will therefore result in a highly personalized user-specific model if a sufficient amount of training data is available. One-size-fits-all modeling, in contrast, trains a single model based on all users' data, so that reasonable predictions can be made about new users for whom insufficient data is available. Next, we proposed our strategy that we dub *one-size-fits-segment* modeling. It was designed to utilize the one-size-fits-all paradigm for making predictions for users grouped by privacy segmentation. We analyzed the overall and per-user performance of the predictive models trained through these three different strategies, and then draw some practical implications.

4.4.5.1 Individual Modeling

This is the most popular and straightforward approach for building user-specific machine learning models for privacy decision support systems, especially targeted at predicting users'

decisions about the disclosure of personal information on social network services [85, 36, 87, 88, 91, 16]. It constructs a distinctive predictive model per each user by using his/her data only. This assumes that each user has a very different point of view regarding privacy, therefore the others' data are not useful for modeling and predicting his/her own privacy decision-making. As most previous works have adopted supervised machine learning approaches, each user will need to provide a certain amount of labeled training data (e.g., historical decisions for privacy-invasive scenarios). For instance, Bilogrevic et al. [16] stated that they needed 40 manual decisions from each user to build personalized models with a reasonable predictive performance, which is a quite burdensome amount. The more training data a user provides, the more the performance of the individual predictive model tends to improve. For these reasons, an individual modeling strategy is suitable for situations in which service providers could acquire enough training data from each individual user. Like [16], the study participants in our dataset made about 40 privacy decisions (reaction parameter *_permission*) on average. Therefore, we first applied individual modeling for constructing personalized LMDNN models to check whether this strategy is adequate for solving our own problem. Specifically, we trained and evaluated a single LMDNN model for a specific user based on his/her data. We repeated it for all users in the dataset. As input features, we utilized solely the five contextual parameters with their interactions. Since each individual LMDNN model is exclusively trained using each user's data, we do not use privacy segment information in this experimental setup.

4.4.5.2 One-size-fits-all Modeling

One-size-fits-all contrasts with an individual modeling strategy. Instead of building multiple user-specific predictive models, it generates a single universal model based on data collected from a crowd of users. Because a one-size-fits-all predictive model is trained using a larger dataset (multiple users' data), it typically represents a wider range of common feature-target

patterns than an individual model. Therefore, researchers have often utilized a one-size-fits-all modeling strategy for building a predictive model that makes initial predictions (e.g., *default* privacy settings) for new users who did not provide training data [16]. Recently published works [93, 73] also adopted one-size-fits-all modeling to make their privacy decision support systems generalizable to a wide range of users. The authors in [73] collected IoT-related privacy decisions from Amazon MTurkers ($N = 1,007$) through an online survey, and built one-size-fits-all predictive models for randomly chosen 50 participants. They reported that the overall accuracy of these trained models ranges from 76% to 80%, depending on whether most (75%) or all (100%) of the other participants' responses are used as training data. To verify the applicability of the one-size-fits-all modeling strategy to privacy decision support for *real-world* IoT environments (using the dataset from our *situated* survey), we built a one-size-fits-all LMDNN model for each user based on data collected from all other users, utilizing both contextual and privacy segment information as input features. Each user's data was used thereafter to assess the predictive power of his/her LMDNN model.

4.4.5.3 One-size-fits-segment Modeling

We proposed an approach called *one-size-fits-segment* modeling. This is a modified version of one-size-fits-all modeling. It builds multiple universal models for several groups of users rather than for the entire user population. Here, we utilized the result of privacy segmentation (see Chapter 4.4.4.2) as a criterion to cluster users. We intended to improve the performance of one-size-fits-all modeling by constructing per-user predictive models based on data collected from *same-minded* users in terms of privacy. We believe that each individual user will be benefitted if a predictive model is trained on large volumes of data provided by others similar to him/her. To verify the proposed approach, we conducted experiments as follows. First, we divided users in the dataset into three groups according to our privacy segmentation. For each single user, we determined the privacy segment to which he/she

belongs, and built a one-size-fits-segment LMDNN model by utilizing data collected from other users in the corresponding privacy segment. Each user's data was used as test data for measuring the predictive performance of his/her LMDNN model. Unlike one-size-fits-all modeling, we only utilized contextual parameters (including interactions between them) as input features. As explained, privacy segment information was utilized to split users with regard to the subset of their stated privacy decisions, thereby constructing one-size-fits-segment LMDNN models for every single user. In comparison with individual modeling and one-size-fits-all modeling, the proposed mechanism does not burden new users because it just asks five questions (reaction parameters) about a single base scenario to perform privacy segmentation.

4.4.5.4 Predictive Performance Evaluation

We compared the predictive power of these three different training strategies as follows. Regarding individual modeling, we trained and assessed the predictive performance of user-specific LMDNN models via 10-fold CV. Each of these models was trained on all data instances from a different user in our dataset. We repeated this procedure for all users and took the average of their AUC scores as the overall predictive performance of individual modeling. For one-size-fits-all or one-size-fits-segment modeling, we constructed a target user's LMDNN model based on all other users' data and tested the trained model using the target user's data. Unlike [73], we repeated this for all users and calculated the average AUC score for each strategy. Finally, we compared these mean AUC scores for three training strategies to choose the best.

Training Strategy	Mean AUC	Std. Dev. (SD)
Individual	0.4806	0.2179
One-size-fits-all	0.6699	0.1721
One-size-fits-segment	**0.6782**	**0.1668**

Table 4.5: Privacy Prediction Performance per Training Strategy (LMDNN)

Table 4.5 summarizes the experimental results showing that both the one-size-fits-all and the one-size-fits-segment modeling strategy significantly outperform the individual modeling strategy. Furthermore, one-size-fits-segment shows a slightly better performance than one-size-fits-all. This is because one-size-fits-segment LMDNN models are trained by the data of same-minded users, thereby predicting each user's privacy decisions more accurately. This finding was unexpected since most previous research about privacy decision support has reported that individual modeling is better than one-size-fits-all modeling for inferring users' privacy decision-making. One possible explanation is that the size of per-user training data was not large enough to learn sparse high-dimensional feature spaces. Participants in the dataset responded to an average of 39.58 scenarios ($SD = 14.65$), or about 33% of all available scenarios. Each scenario received 50.9 responses on average ($SD = 40.16$). When we transformed the dataset into a user-scenario matrix (i.e., each cell represents an observed feature-target pattern), the sparsity[5] of this matrix was 0.6566. This low response rate may also be found in future real-world situations since the number of applications and services in IoT environments is likely to increase over time. As a result, it could be difficult to collect enough training data from each IoT user to build individual predictive models from scratch.

Figure 4.3 presents the per-user predictive performance of two model training strategies: individual and one-size-fits-segment modeling. Results have been filtered out by the threshold of $AUC > 0.5$ and sorted by the AUC score of individual LMDNN models in ascending order. As can be seen, one-size-fits-segment LMDNN models show a better predictive performance than individual LMDNN models for about 80% of the users (left side of the chart). However, there is also the opposite effect for about 20% of users (far right side of the chart) who have highly accurate individual LMDNN models ($AUC > 0.7$) trained solely on their data. To check the causes of this per-user difference, we performed Pearson's chi-square tests for independence between personal attributes (gender, age, type of university affiliation)[6] and

[5] 1 - (nonzero entries/total entries in a user-scenario matrix)

[6] Mode values of these attributes are male, 18-25, and undergraduate students, respectively.

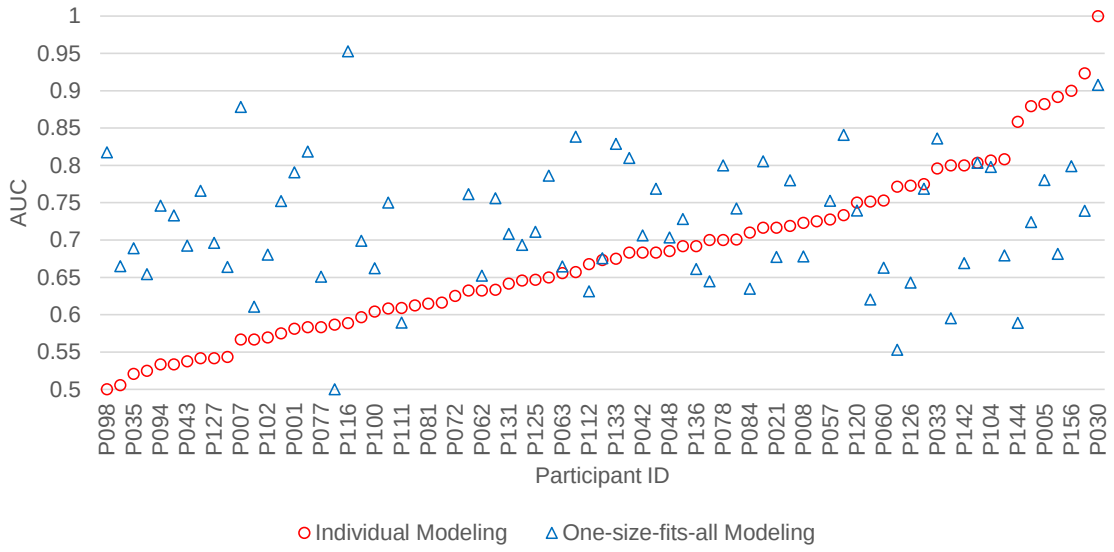

Figure 4.3: Privacy Prediction Performance per Individual User (LMDNN)

the fitted training strategy. For instance, we constructed a 2×2 contingency table using two binary variables, age (18-25 or older than 25) and training strategy (one-size-fits-segment or individual modeling), and assessed the significance of the difference between the two proportions (i.e., users aged 18-25 and older users, both have superior one-size-fits-segment LMDNN models). The test confirms that there is no statistical evidence of an association between these two variables at the .05 significance level. We iterated this test for other personal attributes such as gender (male or others) and the type of university affiliation (undergraduate students or others), and reached the same conclusion. Identifying latent factors (e.g., cultural background in privacy decision-making [61]) that cause this difference is one possible future research direction. If we were to find these factors, we could determine the most appropriate modeling strategy for each individual user in advance, and then accordingly utilize it for better predictive performance.

4.4.6 Implications

We confirmed that privacy segment information is helpful for understanding and predicting people's privacy decision-making about various IoT services. Specifically, one-size-fits-segment modeling shows the best predictive performance (mean AUC of 0.6782; $N = 172$) compared with the previously proposed model training strategies. However, even a single *false positive* may cause undesired information monitoring in IoT, making users reluctant to use privacy decision support systems. At its current accuracy, we should therefore use the prediction results only to give recommendations to users that they can still inspect and override; we should not use the prediction results for automated disclosure decisions. Moreover, we should strive for further improvement of the predictive performance. In the following, we explained how the current predictive performance could be improved.

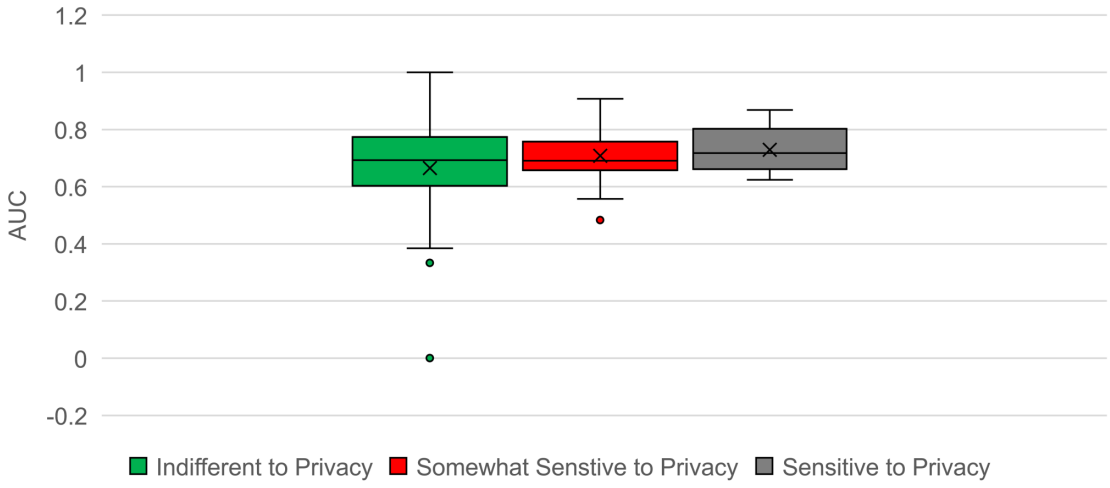

Figure 4.4: Privacy Prediction Performance per Privacy Segment (LMDNN)

4.4.6.1 Role of Privacy Segment Information

To better interpret the performance of one-size-fits-segment modeling, we measured the predictive performance of the trained LMDNN models for each privacy segment. Figure 4.4 indicates that users who are assigned to the *sensitive to privacy* segment have the most

accurate LMDNN models (mean AUC of 0.729, $SD = 0.074$). Likewise, users belonging to the *somewhat sensitive to privacy* segment also tend to have LMDNN models with a reasonable performance (mean AUC of 0.708, $SD = 0.101$). Users in the *indifferent to privacy* segment however have relatively inaccurate LMDNN models compared to those in other privacy segments (mean AUC of 0.665, $SD = 0.184$). Here, we also could not identify any inter-segment differences in terms of personal characteristics such as gender, age, and the type of university affiliation. One possible explanation is that participants in both the *sensitive to privacy* and *somewhat sensitive to privacy* segments presumably responded to IoT monitoring scenarios more carefully than participants in the *indifferent to privacy* segment, thereby providing more consistent privacy decisions (i.e., high-quality training data). We might need to consider strategies for improving the data collection process for the *indifferent to privacy* segment. For instance, users who are indifferent about privacy could be asked for more training data than the other users. This is because a higher amount of training data would reduce uncertainty in predictive models, and therefore likely improve performance. It might also be possible to utilize privacy nudges [65], which can make privacy-insensitive users aware of unexpected outcomes of personal information monitoring, for steering users toward more confident and consistent privacy behavior.

4.4.6.2 Size of Individual Training Data

28 users had highly accurate ($AUC > 0.7$) individual LMDNN models trained on their own data. Furthermore, individual models can be continuously improved as each user provides additional training data. To verify this claim, we built a linear regression model with the number of user responses (i.e., size of training data) as the independent variable and AUC scores (i.e., predictive performance) as the dependent variable. According to the fitted regression model, a statistically significant positive relationship exists between these two variables ($p < .005$). Therefore, the collection of extra training data might further improve

the performance of individual machine learning models. Yet, asking for additional data can also be burdensome and privacy-invasive.

4.4.6.3 Hybrid Modeling

For these reasons, it might be desirable to gradually morph one-size-fits-segment models into individual models. Like [16, 93, 73], we verified that our one-size-fits-segment LMDNN models can make reasonably accurate predictions for two segments, even for new users. Therefore, we can utilize them to produce *default* privacy settings that are applicable to the general population (grouped by privacy segmentation). As mentioned before, however, these general predictive models would need to be tailored to each user to make more accurate predictions. One possible avenue would be to gradually transform one-size-fits-segment LMDNN models into *personalized* individual LMDNN models for each user. Specifically, it is necessary to continuously retrain (update) the base one-size-fits-segment LMDNN model with user-provided training data while increasing the weights of this user-provided data (i.e., transfer learning [80, 50, 60]). This should then better fit the updated LMDNN model with the user's unique behavioral patterns. Active learning paradigms can also be applied to reduce users' labeling efforts as much as possible.

4.5 Discussion

We studied machine learning mechanisms for modeling and predicting users' privacy decisions regarding personal information tracking in IoT environments. After investigating various machine learning algorithms, input features, and model training strategies, we proposed a novel mechanism called one-size-fits-segment modeling for privacy decision support systems in IoT. We showed that the proposed mechanism exhibits a reasonable predictive

performance on privacy behavioral data captured in the field. Our work goes far beyond the limited machine learning approach in our earlier study (see Chapter 4.3), where contextual parameters (without consideration for interaction effects) and cluster membership information of IoT service scenarios were used as features for building a decision tree-based predictive model. We believe this approach is not very practical because it requires users to specify reaction parameter values for dozens of scenarios in advance. Another difference between Chapter 4.3 and this work lies in the examination of the model training strategies; we only tested one-size-fits-all modeling (i.e., single classifier) in Chapter 4.3, and hence the results may provide limited implications (e.g., lack of consideration for per-user predictive performance). Yet, our work still has some shortcomings that will be addressed in the next few sections.

4.5.1 Representability of Data

First, we need to consider the representativeness of the dataset we used. The study participants were predominantly university students aged 18-25 (82%), since we recruited them on campus. This may induce a sampling bias that makes our results less generalizable. In other words, we do not know how the proposed mechanism will work on datasets collected from other populations (e.g., older users who are not familiar with IoT). In this regard, we plan to validate our mechanism with more representative samples, thereby confirming a future direction of this research (e.g., continuous updates of one-size-fits-segment LMDNN models with user-provided data). It will also be necessary to apply the proposed mechanism to different domains (e.g., privacy decisions in healthcare settings), and verify its expandability. To that end, we need to collect or get access to users' privacy behavioral data regarding such a domain.

4.5.2 Reliability of Privacy Segmentation

We utilized the results of privacy segmentation to improve the performance of our predictive models. As discussed, we performed privacy segmentation by clustering users into several groups based on their responses toward a single scenario. Here, we utilized scenario #60 as a base scenario since all users responded to this scenario; this enabled us to assign privacy segment information to each of the users. We also tried four different base scenarios to validate whether our clustering methodology (i.e., determining the correct number and labels of the clusters) is sound, and the result was positive. However, it is also important to check the invariance of the clustering results. As future work, we plan to collect all users' privacy decisions regarding more than one base scenario and conduct pair-wise comparisons on the clustering results so as to confirm the resulting privacy segments are stable no matter which base scenario is used for privacy segmentation. If a user was assigned to the same privacy segment regardless of the base scenario, we can consider the results of privacy segmentation as ground truth. Otherwise, we need to devise a way to decide on the most accurate privacy segment for this user. One possible approach is a majority vote among the segmentation results from all base scenarios.

4.5.3 Privacy Paradox

The privacy paradox is the phenomenon that people's stated privacy preferences or decisions often seem inconsistent with their actual behaviors [3, 45, 30]. As explained before, we utilized stated privacy decisions collected through ESM in a simulated IoT environment as training data, and not actual behavior observed in a working IoT environment. Although we had tried to make participants perceive they were in a real situation, we do not know how they would actually behave in real-world situations. Therefore, we need to construct a working IoT environment, let participants freely interact with the environment, and then

collect the corresponding privacy decisions, possibly with sensor data. To that end, we plan to conduct an experimental study for collecting (large-scale) privacy behavioral data from real users in an operational IoT system based on open protocols for IoT discovery and interaction, such as the Open Connectivity Foundation (OCF), Web of Things, or Physical Web. By using this dataset, we need to confirm the effectiveness of the proposed mechanism.

4.6 Conclusion

In Chapter 4.3, 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 for the specific IoT scenario.

In Chapter 4.4, we proposed a novel machine learning mechanism for predicting people's privacy decisions in IoT environments. We aimed to predict binary privacy decisions for each user, namely whether to allow or reject a given personal information monitoring scenario in IoT. To begin with, we adopted linear model and deep neural networks (LMDNN) as the machine learning model for our study. Using a privacy behavioral dataset ($N = 172$) collected from our earlier situated survey study, we confirmed that LMDNN provides better predictive performance than conventional machine learning models that have been widely used in the literature. Next, we utilized both contextual and privacy segment information as input features for training LMDNN models. We adopted a wide range of contextual factors comprising diverse IoT scenarios. We then generated users' privacy segment information by clustering their privacy decisions about a single selected IoT scenario. Lastly, we proposed a new model training strategy called *one-size-fits-segment* modeling, and compared its performance with two commonly used strategies: individual and one-size-fits-all modeling. Experimental results indicated that one-size-fits-segment outperforms other modeling strate-

gies. We also presented some practical implications regarding the design and development of privacy decision support systems for IoT environments. Future work will focus on collecting privacy-related human behavioral data from more representative samples of users interacting with working IoT systems, and validating the proposed mechanism on this new dataset.

Chapter 5

Modeling and Prediction of Informed Privacy Decision-Making in IoT

Researchers are building IoT systems that aim to raise users' privacy awareness, so that these users can make informed privacy decisions. However, there is a lack of empirical research on the practical implications of informed privacy decision-making in IoT. To gain deeper insights into this question, we conducted an online study ($N = 488$) of people's privacy decision-making as well as their levels of privacy awareness toward diverse IoT service scenarios. Statistical analysis on the collected data confirmed that people who are well aware of potential privacy risks in a scenario tend to make more conservative and confident privacy decisions. Machine learning experiments also revealed that an individual's overall privacy awareness is the most important feature when predicting their privacy decisions. We verified that machine learning models trained on privacy decisions made with confidence can produce highly accurate privacy recommendations for users (AUC of 87%). Based on these findings, we propose functional requirements for privacy-aware systems to facilitate well-informed privacy decision-making in IoT, which results in conservative and *confident* decisions that enjoy high consistency.

5.1 Introduction

The Internet of Things (IoT) is a system of interconnected sensors, computing devices and physical objects with unique identifiers, capable of transferring data and information over a network while minimizing human-computer interaction as much as possible [8]. IoT service providers increasingly deploy intelligent services that collect and analyze various types of available sensor data. For instance, a company can operate an IoT service allowing employees to freely enter the building without badges by automatically verifying their identity based on face photos (captured by smart security cameras)[1]. Such kinds of IoT services can make people's lives more convenient and efficient. At the same time, however, users may have privacy concerns about information that can be inferred from sensor data. Researchers therefore argued that users' privacy should be respected when designing and developing IoT systems, and that it is necessary to provide IoT users with technical means of controlling their privacy [70, 101, 106, 90, 82, 40], such as a feature that let them decide whether or not to use the IoT service (coarse-grained control) or to opt-out of the collection of specific sensor data (fine-grained control). Such a fine-grained privacy control in IoT is being actively investigated [71, 81], but it has not yet been deployed in a real-world setting.

However, even when controling over the use of IoT services is possible, users need to fully understand both the utility benefits and privacy risks in using IoT services, evaluate the balance between them (i.e., perform a utility-privacy tradeoff), and make a final privacy decision based on that evaluation. Yet in reality, it is difficult for them to do so because they are provided with no or only limited information to properly assess potential privacy risks of each IoT service. This lack of privacy risk awareness has been identified in previous research as one of the causes of the privacy paradox[2], which can result in undesirable privacy violations in ubiquitous computing environments [34]. For instance, IoT service providers tend to focus

[1] See, e.g., Baidu's face-enabled entrance (`https://youtu.be/wr4rx0Spihs`)

[2] The privacy paradox is the phenomenon that people's stated privacy preferences or intentions often seem inconsistent with their actual behaviors [3, 45, 76].

on promoting their services (e.g., automatic attendance check) while not informing the users about the sensor data (e.g., face photos) they are collecting and analyzing to provide those services. Informing users in detail about sensor data collection and analysis is important since, due to the rapid advancements in data mining and machine learning, all kinds of data can be used to infer potentially sensitive personal information (e.g., sexual orientation [100]). It is also possible to aggregate multiple types of data from different sensors, so as to infer personal information that previously seemed impossible to figure out.

A meaningful approach to help users make better privacy decisions therefore seems to reside in making them aware of all privacy implications of an IoT service, specifically the types of personal information that can be inferred on the basis of the collected sensor data [40, 25]. However, it is still unclear what consequences greater user awareness of inferable personal information will have on their privacy decisions. We are therefore specifically interested in the following research questions:

RQ1: If people have sufficient awareness of potentially inferable personal information when using an IoT service, will they make their privacy decisions more conservatively, or will they rather disclose more openly, or will there be no difference?

RQ2: If people have sufficient awareness of potentially inferable personal information when using an IoT service, will they make their privacy decisions more confidently or less confidently, or will there be no difference?

RQ3: Will privacy decisions made with confidence form a better basis for accurately predicting people's preferred privacy decisions toward unseen IoT services (effective privacy decision support)?

Regarding **RQ2** & **RQ3**, we had observed in a previous situated survey study that people's privacy decisions toward hypothetical IoT scenarios often seem unreliable (i.e., random

choice). We assume that they lacked understanding of the hypothetical situation and thus made best-guess (unconfident) decisions. We also suspect that these decisions would not be useful for predicting their future decisions due to the lack of consistent behavioral patterns. For these reasons, we believe that the answers to the three research questions will provide useful insights for realizing privacy-preserving IoT environments. To the best of our knowledge, this is the first attempt to verify the practical effects of confident privacy decision-making on protecting user privacy in IoT.

To answer our research questions, we first conducted an online survey on Amazon Mechanical Turk ($N = 488$) to collect and analyze people's privacy decision-making about diverse IoT service scenarios that they may encounter in their everyday lives. We created 180 realistic scenarios while varying underlying contextual factors of the IoT service, such as location, purpose, relationships between collectible sensor data and inferable personal information, and data privacy policies (e.g., how data will be protected, retained, and shared). Each participant was successively presented with 15 randomly chosen scenarios and asked whether or not he or she would use the given IoT service (privacy decision). Along with each privacy decision, participants were prompted for their self-reported level of awareness about the inference of personal information (privacy awareness), perceived balance between utility benefits and privacy risks (utility-privacy tradeoffs), and their level of confidence in their privacy decision (decision confidence). We also constructed each individual's personal privacy propensity which is composed of privacy segment, overall privacy awareness, and privacy self-efficacy. Regarding privacy segment and privacy self-efficacy, which have proven useful for predicting [65] and interpreting [52] people's privacy decision-making, we performed unsupervised K-modes clustering and confirmatory factor analysis (CFA) on survey responses separately collected before and after the abovementioned survey procedures. Lastly, we asked participants for their demographics in order to perform more in-depth analysis.

Next, we conducted statistical analysis to understand which factors influence people's privacy

decision-making in IoT environments. We analyzed the survey data with two different aims: (1) to understand how people make privacy decisions based upon the factors embedded in the scenarios, and (2) to understand how people have different levels of confidence in making such decisions. Regarding statistical models for analyzing people's privacy decisions (binary response) and their decision confidence (ordinal response), we used a generalized linear mixed model (GLMM) and cumulative link mixed model (CLMM), respectively. Through GLMM analysis, (1) we confirmed the fact that several contextual factors, such as service location, the type of inference of personal information, and data sharing practice, significantly impact people's privacy decisions. The fitted GLMM model also confirmed that people who have a higher level of privacy awareness are more likely to make conservative decisions compared to those who do not (**RQ1**). Both privacy segment and self-efficacy turned out to be important privacy decision factors in the GLMM model. Through CLMM analysis, (2) we verified that a number of contextual, privacy-attitudinal and demographic factors are correlated with privacy decision confidence. Furthermore, we verified that the level of privacy awareness is positively associated with the level of decision confidence, regardless of the type of decisions made (**RQ2**). We therefore argue that guiding people to make conservative and confident decisions via the enhancement of privacy awareness is important for protecting their privacy in IoT. Conservative privacy decisions may imply people's willingness to minimize potential privacy risks by opting out of the service, thereby lowering the possibility of privacy violations. Confident privacy decision-making is also desirable since it may be the case that people are then more likely to make their IoT decisions reliably and consistently.

To answer **RQ3**, we performed machine learning (ML) experiments with a Random Forest model, to predict people's privacy decisions based on their past decisions (privacy decision support; [105, 103, 16, 65]). We specifically aimed to investigate whether the level of users' confidence in privacy decisions impacts the predictive performance of the ML model. The experimental results indeed indicate that models exclusively trained on confident privacy decisions show superior accuracy in predicting people's decisions compared to models based

on unconfident decisions (AUC of 87% vs. 67%; statistically significant). Thus, we conclude that collecting confident privacy decisions is crucial for realizing privacy decision support that can give users meaningful privacy recommendations.

Based on the abovementioned findings, we propose functional requirements for privacy-aware systems (PAS) in IoT environments. PAS should be able to maximize users' awareness of potential privacy risks posed by IoT services (e.g., undesired inferences of personal information), thereby not only helping them make conservative decisions on their own but also allowing the PAS to accumulate confident decision samples which can be used to later give users accurate privacy recommendations. We articulate specifically how PAS can better assist users in grasping privacy implications of using the IoT services, leading to their better-informed privacy decision-making.

In summary, our work makes the following contributions to the field of ubiquitous computing, with the ultimate aim of preserving user privacy in IoT:

- We proposed methodologies to model each individual's personal privacy propensity, consisting of privacy segment, overall privacy awareness, and privacy self-efficacy, based on collected survey data regarding people's privacy decision-making behaviors toward IoT services.
- We extracted several context- and user-specific privacy decision factors (including personal privacy propensity) that are significantly associated with the type and confidence of decisions made in IoT, and confirmed that a higher level of privacy awareness is related to the likelihood of privacy decisions being made more conservatively (**RQ1**) and confidently (**RQ2**).
- We demonstrated the feasibility of predicting people's privacy decisions via a machine learning method (AUC of up to 87%) and verified that confident decision samples yield more accurate machine learning models for privacy decision support in IoT (**RQ3**).

- We propose functional requirements for PAS to maximize people's privacy awareness in IoT, thereby nudging them to make more informed privacy decisions on increasingly diverse IoT services.

5.2 Related Work

In this section, we summarize the previous research aimed at understanding and predicting people's privacy decisions in IoT environments. Then we describe research efforts articulating the importance of privacy awareness in making informed privacy decisions. Lastly, we enumerate a set of different systems/methodologies that maximize people's privacy awareness to safeguard their privacy as much as possible in diverse ubiquitous computing environments, including IoT.

5.2.1 Modeling and Prediction of Privacy Decision-Making in IoT

To alleviate the cognitive burden in users' privacy decision-making, there have been attempts to computationally analyze users' privacy perceptions in the context of IoT and make privacy recommendations to users based thereon via machine learning.

In our previous studies, we collected end users' privacy attitudes and decisions toward hypothetical IoT service scenarios through both an online survey ($N = 200$; see Chapter 3.2.1) and a situated survey ($N = 172$; see Chapter 3.2.2). We performed an unsupervised K-modes cluster analysis on the collected datasets and confirmed the fact that contextual parameters, such as the identity of an agent collecting sensor data (parameter *who*) and the type of personal information inferred from the collected sensor data (parameter *what*), significantly impact users' privacy decisions. We also trained diverse machine learning models that predict users' binary privacy decisions based on their past decision-making behaviors

(see Chapter 4). Using the dataset collected in Chapter 3.2.1, Bahirat et al. conducted a statistical analysis (linear mixed model; LMM) in order to determine underlying factors influencing users' privacy decisions toward the presented IoT scenarios [9]. Regarding this, the authors reached the same conclusions with ours, namely that parameters *who* and *what* have significant impacts on user-stated privacy decisions. Through the combination of K-means clustering and decision tree learning, Bahirat et al. also showed that it is possible to predict the users' future privacy decisions with an accuracy as high as 82%. Researchers have continued to study people's privacy decision-making in IoT. Emami-Naeini et al. conducted an online survey study in an experimental setting similar to that used by our own study (see Chapter 3.2.1) but with larger sample populations ($N = 1,007$) [73]. The authors found that privacy decisions not only vary from person to person but are also highly context-dependent. According to their GLMM analysis on the survey responses, the physical location where sensor data collection occurs, user-perceived benefit, and the type of collectible sensor data, significantly impact users' privacy decisions. To validate the feasibility of predicting privacy decisions, Emami-Naeini et al. trained logistic regression-based AdaBoost classifiers capable of predicting users' decisions of whether to accept or reject the IoT service. The average accuracy of their predictions was as high as 86%.

All of the works mentioned above showed that people make different privacy decisions in IoT environments according to underlying contextual factors (e.g., type of inferable personal information) and that it is practically feasible to predict people's decisions using these factors as input features for training machine learning models. All studied factors are characteristics of the IoT scenario. However, it is also possible that people may have internal characteristics that influence the outcome of privacy decisions.

5.2.2 Privacy Awareness and Privacy Decision

Privacy awareness is considered one of the key factors determining the type and quality of people's privacy decisions since it allows them to form more accurate mental models of the current situation [5, 79]. Privacy awareness is the level of user perception of possible privacy risks from interacting with an application or service that can gather personal user data. For example, in online social network services (e.g., Facebook) users post photos and share them with their friends. However, some users are not aware of the fact that if they did not configure privacy settings properly, their photos might be shown to people they do not know. Without a full understanding of this kind of privacy risk, the user may make *uninformed* privacy decisions.

Pötzsch stated that the reinforcement of privacy awareness is one possible solution for remedying the negative outcomes of uninformed privacy decisions [83], namely the privacy paradox. The author not only defined the attributes of privacy awareness in e-commerce and web environments but also argued for the necessity of tools which enhance users' privacy awareness (e.g., an easily-understandable written privacy policy), thereby minimizing the inconsistency between privacy decisions and behaviors. The author also proposed several requirements for realizing such tools. In the same vein, Bergmann performed an online survey study to empirically show how the user's privacy awareness can be changed according to the user interface (UI) presenting privacy-related information on the web [14]. The author quantified users' levels of understanding of possible consequences of disclosing personal data online by showing them two different versions of data entry UIs: (1) a conventional login form collecting user ID and password and (2) a privacy-enhanced login form with clickable links for informing users about the detailed privacy policy (e.g., how the entered data will be protected by the server). Statistical analysis on the survey responses indicated that users who interacted with the privacy-enhanced form indeed have higher privacy awareness than the users of the conventional form. In this context, Deuker also proposed theoretical proposi-

tions that raising the users' privacy awareness should be closely aligned with their knowledge about privacy-enhancing technologies (e.g., differential privacy [35]) applied to application or services they are using [34]. Hong recently stressed the importance of increasing privacy awareness to safeguard user privacy in a ubiquitously connected world [40].

5.2.3 Enhancement of Privacy Awareness

Due to the importance of privacy awareness in making informed decisions, many researchers have designed, developed, and tested computer systems that aim to maximize people's privacy awareness in diverse computing environments.

Malandrino et al. developed and evaluated their privacy protection tool called NoTrace, which provides users with the ability to monitor what kind of personal information is being collected, aggregated, and consumed by the web services (e.g., nytimes.com) [69]. NoTrace has been developed as a web browser plug-in so as to visually summarize possible personal information leakage by websites the user has visited. The authors insisted that NoTrace can help users to make an informed decision about feasible countermeasures (e.g., stop browsing the website), thereby achieving better privacy protection compared to the conventional tools (e.g., AdBlock).

There also have been attempts to augment the user's privacy awareness on SNS. Kang & Kagal proposed a privacy-enhancing framework named Respect My Privacy (RMP) which is designed to let SNS users not only easily configure privacy settings associated with their personal data (e.g., do not use my photos for marketing purposes), but also publicize their preferences through some intuitive icons [49]. The authors stated that this mechanism enables other users to easily grasp the potential consequences of posting/sharing personal data on SNS. Cetto et al. proposed a web-based application named Friend Inspector which lets Facebook users simulate diverse information disclosure scenarios using their own personal

data before actually sharing it [20]. The main objective of Friend Inspector is to decrease the gap between perceived and actual visibility of the shared items (e.g., photos) on Facebook. Using Friend Inspector, the authors stated that users are allowed to effectively learn about the implications of their privacy settings on Facebook, thereby achieving enhanced awareness of the privacy risks.

Regarding ubiquitous computing, Langheinrich proposed a privacy-aware system (PAS) whose aim was not only to allow service providers to announce the details of how sensor data is collected but also to provide users with the technical means of managing how their personal information is inferred and stored by the service [56]. In order to achieve this goal, the author designed a networked system that enables the user to communicate with nearby services by exchanging machine-readable P3P privacy policies[3]. Winkler et al. proposed a location-based PAS targeted at video surveillance systems [102]. The authors built trustworthy surveillance cameras which are capable of interacting with users through their smartphones, to inform them how their sensitive personal information (e.g., identity) is being handled by the cameras. Könings et al. proposed a theoretical framework explaining how the user's privacy awareness can be modeled and enhanced in ubiquitous computing environments [54]. Using a graph-based model, the authors proposed a methodology to represent the data flow and information usage of a specific application (e.g., heart-rate monitoring) and to inform the user about potential privacy violations during its usage. More recently, researchers are embedding privacy awareness in real-world IoT environments. Mehrotra et al. are developing a privacy-aware IoT framework named TIPPERS and deploying it in a university building [71, 81]. TIPPERS is designed to allow its users to specify their preferred privacy settings (e.g., opt-out of specific data collection) and then systematically enforces that the IoT apps to follow these user preferences. To augment users' privacy awareness in IoT environments such as TIPPERS, Das et al. proposed a PAS that allows the users to easily discover the privacy properties of IoT services using a smartphone app called IoT

[3] `https://www.w3.org/TR/P3P/`

Assistants (IoTA) [33, 98, 32]. IoTA users can either configure privacy settings by themselves or get personalized recommendations based on their historical behaviors (i.e., privacy decision support).

As it can be seen, there are a lot of research efforts in designing and developing systems with the aim of increasing people's privacy awareness in IoT environments. However, we find little empirical research that not only measures users' privacy awareness but also investigates the effects of the measured privacy awareness on actual privacy decision-making in IoT. One of the aims of our work is to fill this gap.

5.3 Collection of Privacy Behavioral Dataset in IoT

In this section, we describe how we designed and performed an online survey asking the user's privacy attitudes and decisions about hypothetical IoT services. We also present our approach to characterize each user in three privacy-related dimensions such as privacy segment, overall privacy awareness, and privacy self-efficacy. We then explain how we prepared a dataset for performing statistical analysis and machine learning experiments.

5.3.1 Data Collection Procedure

5.3.1.1 IoT Service Scenario Generation

To gather users' opinions and reactions about their privacy in IoT, we need to present them with real IoT applications/services or hypothetical scenarios. Due to the lack of operational IoT environments, researchers choose the latter approach. One method involves generating IoT scenarios based on random permutations of contextual parameter values (e.g., *who*, *what*, *when*, *where*) [73]. We extracted factors that construct IoT scenarios by referring to the IoT

ID	Location	Purpose	Core Inference[a]	# Possible Inferences
S01	Private	Health	*Vital* ⇒ *Nothing*	1
S02	Private	Saving	*Electricity Usage* ⇒ *Energy Consumption Pattern*	4
S03	Private	Convenience	*Photo* ⇒ *Identity*	4
S04	Private	Convenience	*Voice* ⇒ *Device Control Intention*	4
S05	Private	Saving	*Motion* ⇒ *Presence*	3
S06	Private	Safety	*Voice* ⇒ *Identity*	3
S07	Private	Safety	*OBD*[b] ⇒ *Nothing*	1
S08	Private	Safety	*Photo* ⇒ *Emotion*	4
S09	Work	Convenience	*Photo* ⇒ *Identity*	5
S10	Work	Convenience	*Device ID* ⇒ *User Location*	2
S11	Work	Health	*Video* ⇒ *Physical Activity*	5
S12	Work	Convenience	*Video* ⇒ *Presence*	4
S13	Public	Convenience	*Photo* ⇒ *Identity*	5
S14	Public	Convenience	*Device ID* ⇒ *User Location*	2
S15	Public	Safety	*Video* ⇒ *Identity*	5

[a] *Sensor Data* ⇒ *Personal Information*
[b] On-board diagnostics data (e.g., RPM, speed, pedal position) of the connected car

Table 5.1: Contextual Factors of Base IoT Scenarios

use cases collaboratively created by the IoT community [77, 29]. The factors include the service location (3 values), service purpose (4 values), and personal information inferable from the raw sensor data (12 values). We selected 15 representative base scenarios from the set of possible factor combinations (see Table 5.1). For each of them, we created a list of all possible and some impossible inferences of personal information from the collected sensor data. Regarding this, we first defined 34 possible and 22 impossible inferences based on the literature [66, 6]. Then, we convened a panel of three outside experts in the field of machine learning and ubiquitous computing who decided individually if each inference is currently possible or not. Inter-rater reliability (Fleiss's kappa) was 0.57 at the significance level .05. If there was no unison agreement between raters, a majority rule was applied.

We then divided the possible inferences into a *core* inference and the remaining *non-core* inferences based on whether the inference is articulated in the scenario. For the base scenario S02[4], for instance, *Electricity Usage* ⇒ *Energy Consumption Pattern* would be regarded as a core inference since the scenario explicitly describes this inference. The inference *Electricity*

[4] At your home, your smart electricity meter collects your **electricity usage** to infer your **energy consumption patterns**, thereby suggesting energy saving methods, for your savings.

Usage ⇒ *Number of Household Members* is also technically possible [104], but this is a non-core inference because it is not used to construct S02. On the other hand, *Electricity Usage* ⇒ *Emotion* is currently an impossible inference to date. By considering both core and non-core inferences, we were able to determine the number of all possible inferences for each scenario. Appendix A lists the textual descriptions and factor values of all base scenarios.

Lastly, we further diversified the base scenarios by adding factors describing data privacy policies, namely: whether the collected data will be protected (protected, unprotected), how long the data will be preserved (a week, a month, a year), and whether the data will be shared with third parties (shared, unshared). These factors are known to have a significant impact on people's privacy concerns in IoT [73]. Merging these additional contextual factors with the base scenarios yielded a total of 180 IoT scenarios.

5.3.1.2 Survey Design

Our online survey will be conducted in the following sequence: (1) verifying respondents' comprehension of IoT and example services (IoT comprehension), (2) determination of each user's privacy segment (privacy segmentation), (3) modeling of privacy decision-making behaviors (privacy decision modeling), (4) verifying the attentiveness in taking the survey (attention check), (5) privacy self-efficacy measurement, and (6) collection of demographic information of the survey population.

Step 1 To begin with, we presented survey participants with a brief definition of IoT, and some examples of real-world application (e.g., smart meter), in order to let them better understand the IoT context and then situate themselves in hypothetical but realistic scenarios we created. Participants are advised to read the information carefully and then asked three multiple-choice questions. We excluded participants from our data analysis who entered a wrong answer more than twice.

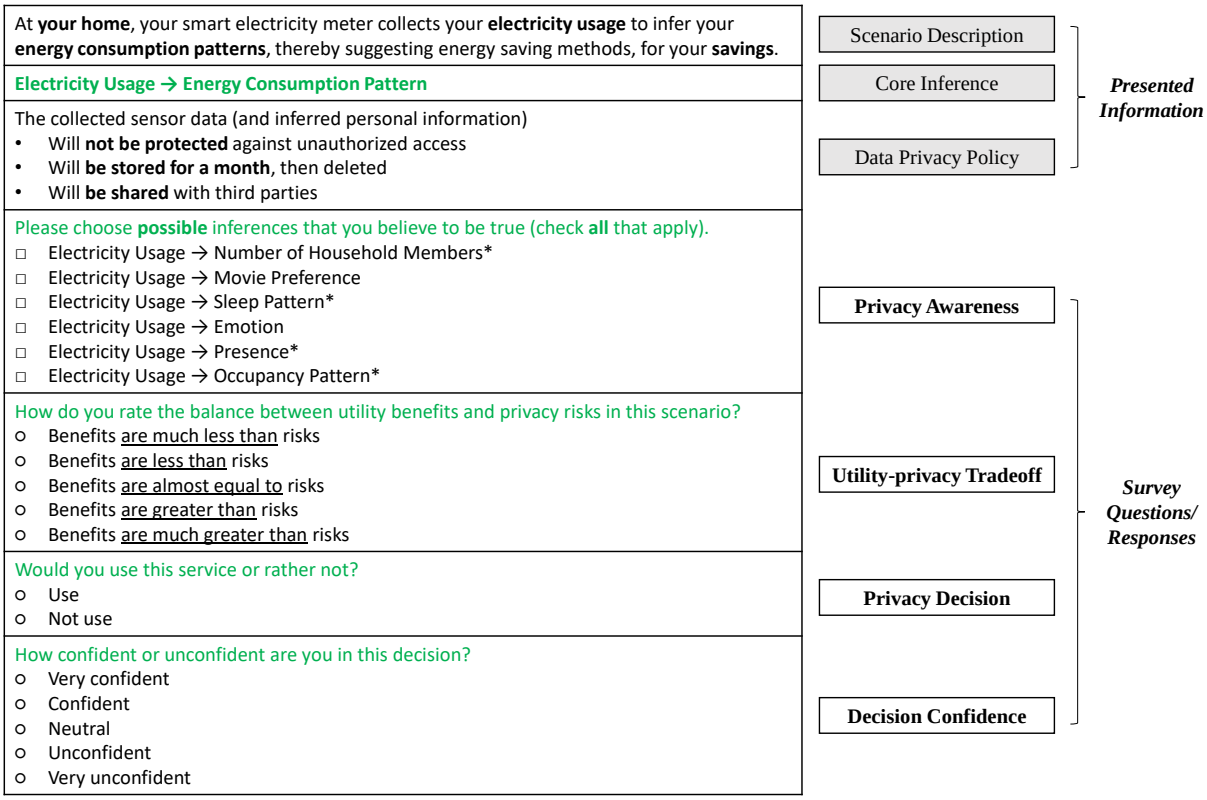

Figure 5.1: Main Survey Screenshot (Privacy Decision Modeling)
Note: $*$ denotes correct possible inference (not shown to participants)

Step 2 To determine each participant's privacy segment, we created three hypothetical IoT services and a questionnaire similar to those described in Chapter 3.2.2. We first generated service descriptions with varying underlying contextual information. Then we asked participants to answer five questions about their reaction attitudes toward the presented services.

Step 3 We then gathered people's decisions of whether or not to use the IoT service, along with their pre/post-decisional factors, namely privacy awareness, utility-privacy trade-off, and decision confidence. A screenshot of a sample scenario (derived from base scenario S02, see Table 5.1) with all questions and answer options is depicted in Figure 5.1. We sequentially presented 15 random IoT scenarios from the available 180 scenarios. Each scenario is composed of the textual service description, the core infer-

ence of personal information as a form of *if-then* rule, and its data privacy policy. All contextual factors defining the scenario (see Section 5.3.1.1) are boldfaced to enable the participant to better understand it. To gauge the privacy awareness of participants, we asked them to mark additionally possible inferences of personal information in the given scenario. We presented all the pre-defined possible and impossible inferences as answer options. Next, we measured participants' utility-privacy tradeoff using a 5-pt Likert scale (e.g., 5 indicates that benefits *are much less than* risks). Participants make privacy decisions about whether or not to use the presented IoT scenario. Lastly, we inquired about their level of self-confidence regarding the decisions they made in the previous step on a 5-pt Likert scale (e.g., 5 indicates privacy decisions with *high confidence*).

Step 4 We aimed to check whether survey participants are paying enough attention to survey content. To that end, we informed them of the estimated remaining time and asked them on the next page to enter this value. We discarded participants who answered incorrectly.

Step 5 We measured privacy self-efficacy using a scale developed by LaRose & Rifon [58]. This scale consists of 10 items that gauge participants' familiarity and confidence regarding practices for safeguarding online privacy (e.g., I know how to change the privacy settings of my browser), on a 5-pt Likert scale.

Step 6 Lastly, we asked participants about their age, gender, length of residence in the US, education level, and annual household income, to understand the survey population and identify additional factors that potentially impact people's privacy decision-making in IoT. Also, we allowed participants to choose *Prefer not to disclose* (n/a: no answer) if they wanted.

We summarize the survey responses collectible through the abovementioned procedures in

Table 5.2. We implemented and deployed our survey on SurveyMonkey[5].

5.3.1.3 Survey Administration

We recruited 507 survey participants on Amazon Mechanical Turk (MTurk) in late March 2018. To assure the quality of survey responses, we restricted participants to adults who live in the US, are proficient in English, and have a high reputation on MTurk (above 97% task approval rate and above 5,000 approved tasks). The recruited participants are provided with a link to our survey deployed at SurveyMonkey. All of them received $2 as compensation if they completed the survey. As discussed before, we discarded participants if they failed to pass IoT comprehension or the attention check question. We also eliminated some participants who finished the survey too rapidly (in less than 5 minutes, compared to the average response time of about 15 minutes). Thereby, we secured 488 qualified survey participants. Table 5.3 displays their demographics.

5.3.2 Privacy Propensity Modeling

Based on the collected survey responses, we extracted additional factors that represent people's personal privacy propensity, namely privacy segment, overall privacy awareness, and privacy self-efficacy.

5.3.2.1 Privacy Segment

Privacy segmentation, which groups users into specific segments according to the similarity of their privacy perception and behavior, can yield insights for understanding and predicting people's privacy decision-making in various circumstances [51, 65, 57]. The most common

[5] `https://www.surveymonkey.com/`

Purpose	Response	Description	Data Type
IoT Comprehension	**Understanding Level**	Ratio of correct answers to all IoT-related questions	**Numerical** (0.0∼1.0)
Privacy Segmentation	**Notification Preference**	Intention to be notified about data collection being performed by IoT service	**Categorical** (4-class)
	Permission Preference	Intention to allow or reject data collection being performed by IoT service	**Categorical** (4-class)
	Perceived Comfort	Perceived level of comfort about data collection being performed by IoT service	**Ordinal** (7pt-scale)
	Perceived Risk	Perceived level of risk about data collection being performed by IoT service	**Ordinal** (7pt-scale)
	Perceived Appropriateness	Perceived level of appropriateness about data collection being performed by IoT service	**Ordinal** (7pt-scale)
Privacy Decision Modeling	**Privacy Awareness**	Ratio of correct answers to all questions about possible inferences in given IoT service	**Numerical** (0.0∼1.0)
	Utility-privacy Tradeoff	Perceived balance between utility benefits and privacy risks in given IoT service	**Ordinal** (5pt-scale)
	Privacy Decision	Intention to use or not to use given IoT service	**Categorical** (binary)
	Decision Confidence	Perceived level of confidence in making a privacy decision for given IoT service	**Ordinal** (5pt-scale)
Attention Check	**User Attention**	Indication of whether the participant is paying attention while participating	**Categorical** (binary)
Privacy Self-efficacy Measurement	**Agreement Level**	Level of agreement with presented statements on online privacy self-management	**Ordinal** (5pt-scale)
Demographic	**Age**	Age group	**Categorical** (6-class)
	Gender	Gender identity	**Categorical** (3-class)
	Residence	Length of residence in the US	**Categorical** (5-class)
	Education	Highest level of education	**Categorical** (8-class)
	Income	Annual household income level	**Categorical** (7-class)

Note: Categorical/ordinal values of the survey response are as follows.

- **Utility-privacy Tradeoff**: (1) Benefits $\gg$ Risks, ..., (5) Benefits $\ll$ Risks
- **Decision Confidence**: (1) Very unconfident, ..., (5) Very confident
- **Agreement Level**: (1) Strongly disagree, ..., (5) Strongly agree

Table 5.2: Summary of Survey Responses

Age		Gender		US Residence (yr)		Education		Income	
18-24	4.7%	Male	48.6%	0-4	0.2%	<High School	0.8%	<$20,000	10.9%
25-34	**36.9%**	Female	51.2%	5-10	0.8%	High School	25.2%	$20,000-$34,999	19.3%
35-44	28.9%	n/a	0.2%	11-20	1.2%	Associate	21.1%	$35,000-$49,999	18.6%
45-54	14.8%			**>20**	**96.9%**	**Bachelor**	**41.2%**	**$50,000-$74,999**	**24.0%**
55-64	11.3%			n/a	0.8%	Master	8.0%	$75,000-$99,999	14.3%
>65	3.5%					Professional	1.4%	>$100,000	10.7%
						Doctorate	1.4%	n/a	2.3%
						n/a	0.8%		

Table 5.3: Demographic Breakdown of Survey Population

approach for privacy segmentation is running an unsupervised clustering algorithm on user-provided data (e.g., survey responses).

We performed privacy segmentation as follows. Based on methodologies described in Chapter 3.2.2, we first created three representative IoT scenarios and collected the reaction attitudes of our survey participants for each scenario (see Appendix B). Then we performed K-modes[6] cluster analysis to segment users based on the commonality of their responses. To begin with, we determined the optimal number of clusters (K) through the Elbow method; we found out that for a specific scenario, the largest decrease in clustering errors occurred when we increased K from 2 to 3. Therefore, we chose 3 as a suitable number of clusters for this scenario. We repeated this procedure against the remaining two scenarios and drew the same conclusion ($K = 3$). Next, we ran the K-modes clustering algorithm[7] on our data in order to determine each participant's cluster membership (privacy segment). Table 5.4 presents the cluster centroids learned from the user responses to the specific scenario. These centroids are formulated with the representative reaction attitude values for the corresponding cluster. As it can be seen, each cluster is quite distinct, primarily in the perceived level of comfort, risk, and appropriateness: each centroid has a unique combination of these reaction values. Therefore, we marked the clusters based on these three reactions (see Privacy Segment column in Table 5.4). We also statistically validated the distinctiveness of the

[6]As a variant of K-means, K-modes clustering [43, 44] is designed to cluster categorical (or ordinal) values without data conversions.

[7]We used `klaR` package, an R implementation of K-modes, for the task of privacy segmentation.

Notification Preference	Permission Preference	Perceived Comfort	Perceived Risk	Perceived Appropriateness	Privacy Segment
Notify, always	**Allow**, just this time	Somewhat **comfortable**	**Neutral**	Somewhat **appropriate**	***Somewhat Insensitive***
Notify, always	**Reject**, just this time	Somewhat **uncomfortable**	Somewhat **risky**	Somewhat **inappropriate**	***Somewhat Sensitive***
Notify, always	**Reject**, always	Very **uncomfortable**	Very **risky**	Very **inappropriate**	***Very Sensitive***

Table 5.4: User Cluster Centroids for Sample IoT Scenario

clustering results by conducting Welch's t-tests[8]. The tests confirm that the difference in the means of each reaction value between each pair of the resulting segments is statistically significant ($p < 0.025$, Bonferroni-corrected for two comparisons). Again, we repeated this procedure for the remaining scenarios, and came up with similar results.

Ideally, the results of privacy segmentation should be independent of the scenario used for clustering. However, we noticed that some participants (about 13%) get clustered into three different privacy segments. It is possible that these participants reacted differently to each scenario. Accordingly, we marked their segments as *undecided*. Otherwise, we adopted the majority voting approach to determine an individual's privacy segment. As a result, 26% of participants were classified as *somewhat privacy insensitive* ($N = 127$), 17% as *somewhat sensitive* ($N = 85$), and 43% as *very sensitive* ($N = 211$).

5.3.2.2 Overall Privacy Awareness

As discussed, we measured the privacy awareness of the participants by asking them to choose additional inferences of personal information that they believed to be true. For each of the 15 presented scenarios, we calculated a participant's privacy awareness as the ratio of correct answers. We used a modified mean of the measured scores after excluding the highest

[8]The reason for using Welch's t-test is that all privacy segments have different variances in all reaction values.

and lowest score as a measure of overall privacy awareness of the participant (population means 0.72, $SD = 0.15$).

5.3.2.3 Privacy Self-efficacy

Privacy self-efficacy is defined as people's self-confidence or belief of abilities in protecting their privacy by themselves. We adopted LaRose & Rifon's scale [58] for quantifying the level of privacy self-efficacy of our survey participants (see Table 5.5) and tested whether our survey data fit this hypothesized privacy self-efficacy measurement model via confirmatory factor analysis (CFA) [18]. To be specific, we fitted a CFA model[9] on our survey responses toward all the presented indicators (i.e., full model). In doing this, we used a weighted least squares estimator (e.g., WLSMV) since it provides the optimal way for modeling categorical/ordinal data without an assumption of data normality. The full CFA model showed that 4 indicators should be dropped since they have unsatisfactory factor loadings < 0.7.

ID[a]	Item (from LaRose & Rifon 2007)	FL[b]	R^2
PSE1	It's easy to figure out which sites you can trust on the Internet.	0.68	0.47
PSE2	**I am confident I know how to protect my credit card information online.**	**0.84**	**0.71**
PSE3	**I know how to identify sites with secure servers.**	**0.76**	**0.57**
PSE4	**I know how to evaluate online privacy policies.**	**0.76**	**0.58**
PSE5	It's easy to set up dummy email account to shield my identity.	0.60	0.37
PSE6	**I know how to change the security settings of my browser to increase privacy.**	**0.71**	**0.51**
PSE7	I know how to use a virus-scanning program.	0.63	0.40
PSE8	**I am able to protect myself against the release of personal information.**	**0.75**	**0.57**
PSE9	I know how to block unwanted E-mails.	0.64	0.40
PSE10	**Overall, I am confident that I can protect my privacy online.**	**0.85**	**0.72**

[a]Indicator
[b]Factor Loading

Table 5.5: Confirmatory Factor Analysis for Privacy Self-efficacy

Considering only the indicators with satisfactory factor loadings, we fitted a reduced CFA model with a reasonable convergent validity ($AVE = 0.61$). Based on the fitted reduced model, we computed the factor score (standardized weighted average value of the latent

[9]We used `lavaan` package, an R implementation of CFA, for the task of measurement of privacy self-efficacy.

variable based on factor loadings) for each participant. We then treated the computed factor score as the measurement of his/her privacy self-efficacy. The measured values of privacy self-efficacy range from -1.7 to 2.4 ($M = 0.00$, $SD = 0.79$).

5.3.3 Dataset Summary

Through the abovementioned procedures of online survey administration and privacy propensity modeling, we created a dataset with 19 attributes that characterize both the IoT-enabled spaces and users who interact with them (see Table 5.6). There exist 7 attributes with contextual information about IoT services, 4 attributes for users' privacy decisions and pre/post-decisional factors (privacy decision-making behavior), 3 attributes for indicating their privacy propensity, and 5 attributes for demographic information. 15 attributes are categorical/ordinal and 4 numerical. As 488 survey participants individually responded to 15 scenarios, we came up with 7,320 user-scenario instances. There were 2,916 accept decisions (*use* the service) and 4,404 rejects.

The dataset is now fed into our statistical (GLMM & CLMM) and machine learning (Random Forest; RF) models to determine underlying factors that yield conservative and confident privacy decisions, and to investigate the effects of the decision confidence on the prediction of future decisions. Most of the attributes will be treated as independent variables (IVs) for statistical models or input features for machine learning models. A *privacy decision* attribute, in contrast, will be used as a dependent variable (DV) or a label to be predicted (see the last three columns in Table 5.6). Note that we excluded *utility-privacy tradeoff* from the input features for building RF models since it only constitutes a semi-final privacy decision. Regarding the numerical attributes, we performed min-max normalization in order to rescale the range of original values to $[0, 1]$, thereby achieving better fitted predictive models.

Category	Attribute	Type	Range	GLMM	CLMM	RF
N/A	Subject ID	Cate.	488-class	Random IV	Random IV	*Excluded*
Contextual Information	Location	Cate.	3-class	Fixed IV	Fixed IV	Feature
	Purpose	Cate.	4-class	Fixed IV	Fixed IV	Feature
	Core Inference	Cate.	12-class	Fixed IV	Fixed IV	Feature
	Data Protection	Cate.	2-class	Fixed IV	Fixed IV	Feature
	Data Retention	Cate.	3-class	Fixed IV	Fixed IV	Feature
	Data Sharing	Cate.	2-class	Fixed IV	Fixed IV	Feature
	# Possible Inferences	Nume.	[0,1]	Fixed IV	Fixed IV	Feature
Privacy Decision-making Behavior	Privacy Awareness	Nume.	[0,1]	Fixed IV	Fixed IV	Feature
	Utility-privacy Tradeoff	Ordinal	5-class	Fixed IV	Fixed IV	*Excluded*
	Privacy Decision	Binary	2-class	**DV**	Fixed IV	**Label**
	Decision Confidence	Ordinal	5-class	*Excluded*	**DV**	*Excluded*
Personal Privacy Propensity	Overall Privacy Awareness	Nume.	[0,1]	Fixed IV	Fixed IV	Feature
	Privacy Segment	Cate.	4-class	Fixed IV	Fixed IV	Feature
	Privacy Self-efficacy	Nume.	[0,1]	Fixed IV	Fixed IV	Feature
Demographic Information	Age	Cate.	6-class	Fixed IV	Fixed IV	Feature
	Gender	Cate.	3-class	Fixed IV	Fixed IV	Feature
	Residence	Cate.	5-class	Fixed IV	Fixed IV	Feature
	Education	Cate.	8-class	Fixed IV	Fixed IV	Feature
	Income	Cate.	7-class	Fixed IV	Fixed IV	Feature

Table 5.6: Dataset Summary

5.4 Modeling of Informed Privacy Decision-Making in IoT

In this section, we explain how we statistically analyzed people's privacy decision-making in more detail. First, we not only identify factors that have a meaningful effect on binary privacy decisions in IoT environments but also interpret how these factors impact the type of decisions (e.g., permissive vs. conservative; **RQ1**). In addition, we investigate which factors influence the user-perceived confidence about the decisions made (e.g., confident vs. unconfident; **RQ2**), and how.

5.4.1 Factors Impacting Privacy Decision-Making

In order to model privacy decision factors in IoT, we adopted a generalized linear mixed model (GLMM) regression [10] with a random intercept per participant. GLMM is an extension of the generalized linear model (GLM) in which the linear covariates contain both fixed and random effects. For the analysis of the grouped data, GLMM can model the differences between groups as a random effect. GLMM is suited for analyzing our survey data since each of our participants responded to multiple parallel scenarios (grouped by subject); it enables us to take the within-subject associations into account when interpreting the analysis results.

5.4.1.1 Problem Definition

Modeling privacy decision factors can be defined as follows. Consider a GLMM model with p covariates for the probability of the dichotomous response Y_{ij} of the subject $i\,(1,\ldots,N)$ for the item $j\,(1,\ldots,n_i)$ being equal to one:

$$Logit(P(Y_{ij}=1)) = x_{ij}^T\beta + v_i, \tag{5.1}$$

where Y_{ij} is the binary privacy decision of the i^{th} subject toward the j^{th} scenario (1: use, 0: not use); x_{ij} is the $(p+1)$-dimensional vector of independent variables (i.e., fixed effects); β is the $(p+1)$-dimensinal vector of regression coefficients; v_i is the random subject effects distributed iid-normal: $v_i \sim \mathcal{N}(0,\sigma_v^2)$; N is the number of subjects (488); and n_i is the number of scenarios shown to the i^{th} subject (15). We specified a logit link function since we assumed that the probability distribution of $Y_{ij}=1$ is binomial.

5.4.1.2 Model Selection

We performed model selection to find the best combination of privacy decision factors through a backward elimination approach [46]. Our model selection routine starts with fitting a model with the most complex structure possible given the specified combination of fixed effects and their interactions. It then performs backward stepwise selection to obtain the minimum adequate model based on likelihood ratio test (LRT). Interaction terms are tested first, and then removed to test each fixed effect. All fixed effects that are part of significant interaction terms are retained in the final model regardless of their significance level. The specified random effects are fixed. A maximum likelihood (ML) estimator is used for selecting privacy decision factors and then restricted maximum likelihood (REML) estimator is used for fitting the final reduced model.

Though the abovementioned procedures, we fitted the reduced GLMM model[10] as shown below:

$$Logit(P(Y_{ij}=1)) = \beta_0 + \beta_1(\texttt{location}_{ij}) + \beta_2(\texttt{infer}_{ij}) + \beta_3(\texttt{protect}_{ij}) + \beta_4(\texttt{share}_{ij}) + \beta_5(\texttt{aware}_{ij}) + \beta_6(\texttt{tradeoff}_{ij}) + \beta_7(\texttt{segment}_i) + \beta_8(\texttt{efficacy}_i) + v_i, \tag{5.2}$$

where $\{\texttt{location}_{ij}, \texttt{infer}_{ij}, \texttt{protect}_{ij}, \texttt{share}_{ij}\}$ represent contextual information of the j^{th} scenario shown to the i^{th} subject, namely its service location, core inference, data protection and data sharing; $\{\texttt{aware}_{ij}, \texttt{tradeoff}_{ij}\}$ is the measured privacy awareness and utility-privacy tradeoff[11] of the i^{th} subject toward the j^{th} scenario (i.e., pre-decisional factors); $\{\texttt{segment}_i, \texttt{efficacy}_i\}$ are the determined privacy segment and self-efficacy of the i^{th} subject (i.e., personal privacy propensity); β_0 is an intercept; and β_p is a regression coefficient

[10] We used `lme4` package, an R implementation of GLMM, for modeling privacy decision factors.

[11] We treated the ordinal values of utility-privacy tradeoff as continuous $[1, \ldots, 5]$ in the model-fitting procedures for easy interpretation.

of the p^{th} independent variable ($p = 1, \ldots, 8$). Y_{ij} and v_i are the same as denoted by Equation (5.1).

5.4.1.3 Analysis Results

We present the regression results of the fitted GLMM model in Table 5.7. As can be seen, most factors are associated with the dependent privacy decision variable. However, it is also true that not all levels of the categorical factors are statistically significant. Below we interpret the results based on the odds ratio (OR) computed for each factor category with $p < .05$ (marked out in bold).

Contextual Information Regarding contextual information, we confirmed that the current location, type of inferable personal information, and data sharing practice significantly impact people's privacy decisions.

- Previous research [9] indicates that, in general, people consider IoT-based sensor data collection and analysis as very unacceptable if they occur in private places, like their home. Thus, we interpreted our regression results by considering `location:Private` as a baseline category. We confirmed that our findings are consistent with the literature. For the IoT service being operated in the workplace (`location:Work`) and public spaces (`location:Public`), the odds of using the service (i.e., accept decision) are predicted to be 4 times and 3 times higher than the service at the private space, respectively ($p < .001$ for both). It also indicates that the US respondents are more used to giving up privacy at their workplaces than public area.
- The type of inference of personal information was also a dominant factor influencing privacy decisions in IoT, as addressed by [73]. Regarding this, we specified a baseline category as users' indoor location inferred from unique identifiers of

Factor Impacting Privacy Decision	Coef β	SE(β)	z	p	OR (95% CI)
`Intercept`	5.71	0.40	14.37	***	301.16 (138.25-656.04)
`location:Private`				*baseline*	
`location:Work`	1.64	0.22	7.61	***	**5.13** (3.37-7.82)
`location:Public`	1.40	0.21	6.53	***	**4.04** (2.66-6.15)
`infer:DeviceID⇒UserLocation`				*baseline*	
`infer:Electricity⇒EnergyConsumptionPattern`	2.89	0.28	10.35	***	**18.05** (10.43-31.22)
`infer:Motion⇒Presence`	2.65	0.28	9.57	***	**14.22** (8.26-24.51)
`infer:Video⇒Presence`	-0.08	0.21	-0.39		0.92 (0.61-1.39)
`infer:Voice⇒DeviceControlIntention`	2.26	0.28	8.08	***	**9.62** (5.56-16.67)
`infer:Voice⇒Identity`	1.87	0.27	6.80	***	**6.48** (3.78-11.11)
`infer:Photo⇒Identity`	0.94	0.17	5.60	***	**2.56** (1.84-3.56)
`infer:Video⇒Identity`	0.83	0.21	4.00	***	**2.29** (1.52-3.43)
`infer:Video⇒PhysicalActivity`	-0.71	0.22	-3.28	**	**0.49** (0.32-0.75)
`infer:Photo⇒Emotion`	0.12	0.31	0.40		1.13 (0.62-2.08)
`infer:OBD⇒Nothing`	2.77	0.28	10.03	***	**15.88** (9.25-27.27)
`infer:Vital⇒Nothing`	2.36	0.27	8.61	***	**10.63** (6.2-18.2)
`protect:Protected`				*baseline*	
`protect:Unprotected`	-0.14	0.08	-1.67		0.87 (0.74-1.02)
`share:Unshared`				*baseline*	
`share:Shared`	-0.19	0.08	-2.35	*	**0.82** (0.7-0.97)
`aware`	-0.81	0.21	-3.93	***	**0.45** (0.3-0.67)
`tradeoff`	-2.06	0.06	-34.86	***	**0.13** (0.11-0.14)
`segment:SomewhatInsensitive`				*baseline*	
`segment:SomewhatSensitive`	-0.45	0.24	-1.91		0.64 (0.4-1.01)
`segment:VerySensitive`	-1.51	0.20	-7.62	***	**0.22** (0.15-0.32)
`segment:Undecided`	-0.19	0.26	-0.72		0.83 (0.5-1.37)
`efficacy`	0.95	0.42	2.25	*	**2.59** (1.13-5.93)

Note: $*$ $p < 0.05$, $**$ $p < 0.01$, $***$ $p < 0.001$

Table 5.7: Regression Results of GLMM model

their mobile devices (`infer:DeviceID⇒UserLocation`) as this is the most common inference possible in currently available IoT environments (e.g., Wi-Fi-based occupancy monitoring).

- We found that people tend to accept an IoT service if it anonymously infers simple behavioral patterns of theirs. The probabilities of accepting the inference of electricity usage patterns (`infer:Electricity⇒EnergyConsumptionPattern`) and presence (`infer:Motion⇒Presence`) at a certain location are 17 times and 13 times higher than the inference of user location mapped with device IDs, respec-

tively ($p < .001$ for both). Similarly, the odds of allowing the voice-based inference of intentions to control consumer electronics (e.g., Turn on the TV) without user identification (`infer:Voice⇒DeviceControlIntention`) are more than 8 times higher than the baseline inference ($p < .001$).

- In contrast, people are relatively less receptive toward services involving user identification. The odds of utilizing the service that infers (and verifies) the user's identity are 5.5, 1.6, and 1.3 times higher than the baseline, depending on the type of raw sensor data used: voice (`infer:Voice⇒Identity`), photo (`infer:Photo⇒Identity`), and video (`infer:Video⇒Identity`), respectively ($p < .001$ for all). We also found that people are more reluctant to admit a video-based identification mechanism than other approaches (photo and voice). Note that inferring users' identity is still considered more acceptable than inferring their current location. This is probably because people are more comfortable with user identification since they already have been using it in non-IoT applications/services (e.g., face unlock on smartphone). However, the inference of their physical activities (e.g., running) based on the captured video data (`infer:Video⇒PhysicalActivity`) was unacceptable; the odds to accept this kind of inference are predicted to be 0.51 times lower than the baseline ($p < .01$).
- Lastly, we intentionally defined IoT services that collect certain sensor data for a specific purpose but does not infer any additional information from it, to check people's reactions to virtually privacy-safe IoT scenarios. As expected, people are more likely to allow the service if it declares not to infer personal information, even though its collectible sensor data might allow sensitive inferences (e.g., health conditions from heart rate measurement). The odds of accepting a service that simply collects OBD data (`infer:OBD⇒Nothing`) or vital signs (`infer:Vital⇒Nothing`) without any inferences are about 15 and 10 times higher than a service that infers user location ($p < .001$ for both).

- With regard to data privacy policy, people make different privacy decisions depending on whether the collected sensor data will be shared or not. Specifically, people are more likely to refuse to use an IoT service if it shares data with third parties. The odds of using services with data sharing practice (`share:Shared`) are predicted to be 0.18 times lower compared to other services which do not share data ($p < .05$). The fact that virtually all kinds of sensor data might be used to infer higher-level user information caused people to make more conservative decisions. This finding is also aligned with [73].

Pre-decisional Factors Regression results also showed that privacy awareness and utility-privacy tradeoff significantly impact people's privacy decision-making in IoT.

- Regarding privacy awareness, we verified that people who are knowledgeable about possible inferences of personal information in IoT environments tend to be conservative when they make decisions about the use of services (**RQ1**). For one unit increase in the level of privacy awareness (`aware`), the odds of an accept decision decreases by a factor of 0.45 ($p < .001$). This indicates that a higher level of privacy awareness makes people more concerned about potential privacy violations (e.g., via undesired inferences), leading to a rejection of the service.
- Utility-privacy tradeoff was also a significant decision factor since it can act as a proxy for assessing how the IoT service would be beneficial (or risky). In this analysis, we treated this factor as a continuous variable (atomic vectors) in order to ease interpretation: (1) `Benefits`$\gg$`Risks`, (2) `Benefits`$>$`Risks`, (3) `Benefits`$\approx$`Risks`, (4) `Benefits`$<$`Risks`, (5) `Benefits`$\ll$`Risks`. As it can be seen, a higher value of utility-privacy tradeoff indicates that people consider a specific service to present more risks than benefits. We found that people decided not to use the IoT service if they felt it was risky; for one unit increase in the level of perception of privacy risks relative to utility benefits (`tradeoff`), the probabilities

of accepting the service decrease by a factor of 0.13 ($p < .001$).

Personal Privacy Propensity Regarding privacy propensity, both privacy segment and self-efficacy are shown to be significant user-specific factors impacting people's privacy decisions.

- As discussed earlier, we discovered three privacy segments and assigned each participant with one of the resulting segments. However, the regression results showed that not all segments are associated with the dependent variable at the significance level of .05. Nevertheless, we could verify that a segment labeled as very sensitive (VS) to privacy can be clearly distinguished from a somewhat insensitive (SI) segment. As expected, users who are clustered into the VS segment tend to make more conservative decisions than others in the SI segment (baseline). For people under VS segment (`segment:VerySensitive`), the odds of an accept decision are predicted to be 0.78 times lower than people under SI segment ($p < .001$).
- Privacy self-efficacy also acted as an important decision factor as described in [58, 52]. Contrary to privacy awareness, a higher level of privacy self-efficacy makes people more permissive to the IoT service; for one unit increase in the level of privacy self-efficacy (`efficacy`), the odds of using IoT services increase by a factor of 2.59 ($p < .05$). This is because people who have high privacy self-efficacy believe that they can protect their privacy themselves against the services capable of revealing personal information.

5.4.2 Factors Impacting Confidence in Privacy Decision-Making

Even though we figured out how diverse underlying factors impact the type of privacy decisions that people made, it is still unclear whether they were confident about their decisions

and what factors contribute to their confidence. Therefore we measured the level of decision confidence via a 5-pt Likert scale in our survey, along with binary privacy decisions.

Since the measured confidence level is ordinal in nature, we used a cumulative link mixed model (CLMM) [96] for identifying factors related with the user's confidence in his/her privacy decisions. CLMM, a sort of linear mixed models (LMM), is designed to handle the ordered but non-continuous ordinal response data, such as ours. Like GLMM, it can take random subject effects into account for multiple measurements taken on the same individual or across time.

5.4.2.1 Problem Definition

Modeling privacy decision confidence can be generally defined as follows. Consider a CLMM model with p covariates for the cumulative probability of the ordinal response Y_{ij} of the subject $i\,(1,\ldots,N)$ for the item $j\,(1,\ldots,n_i)$ falling in the category k or below $(k=1,\ldots,K-1)$:

$$Logit(P(Y_{ij} \leq k)) = \theta_k - x_{ij}^T\beta - v_i, \tag{5.3}$$

where Y_{ij} is the confidence level of the privacy decision made by the i^{th} subject toward the j^{th} scenario (1: very unconfident, ..., 5: very confident); θ_k is a set of threshold parameters (cut-points) for the response categories, x_{ij} is the $(p+1)$-dimensional vector of independent variables; β is the $(p+1)$-dimensinal vector of regression coefficients; v_i is the random subject effects distributed iid-normal: $v_i \sim \mathcal{N}(0,\sigma_v^2)$; N is the number of subjects (488); and n_i is the number of scenarios shown to the i^{th} subject (15); and K is the number of response categories (5). We specified a logit link function since we assumed that the probability distribution of $Y_{ij} \leq k$ is binomial.

5.4.2.2 Model Selection

Like the model selection approach described in Section 5.4.1.2, we performed a backward stepwise elimination based on LRT to find the most parsimonious CLMM model. As a result, we fitted the reduced CLMM model[12] as shown below:

$$Logit(P(Y_{ij} \leq k)) = \theta_k - \beta_1(\texttt{location}_{ij}) - \beta_2(\texttt{infer}_{ij}) - \beta_3(\texttt{aware}_{ij}) - \beta_4(\texttt{tradeoff}_{ij}) - \beta_5(\texttt{segment}_i) - \beta_6(\texttt{efficacy}_i) - \beta_7(\texttt{residence}_i) - \beta_8(\texttt{income}_i) - v_i, \tag{5.4}$$

where $\{\texttt{location}_{ij}, \texttt{infer}_{ij}\}$ are the service location and core inference of the j^{th} scenario shown to the i^{th} subject; $\{\texttt{aware}_{ij}, \texttt{tradeoff}_{ij}\}$ are the measured privacy awareness and utility-privacy tradeoff of the i^{th} subject toward the j^{th} scenario; $\{\texttt{segment}_i, \texttt{efficacy}_i\}$ are the determined privacy segment and self-efficacy of the i^{th} subject; and $\{\texttt{residence}_i, \texttt{income}_i\}$ are the length of US residence and annual household income of the i^{th} subject (i.e., demographic information). β_p is a regression coefficient of the p^{th} independent variable ($p = 1, \ldots, 8$). Y_{ij}, θ_k, and v_i are the same as denoted by Equation (5.3).

5.4.2.3 Analysis Results

We presented the regression results of the fitted CLMM model in Table 5.8. As we did in Section 5.4.1.3, we interpreted these results primarily based on factors which have a significant association ($p < .05$) with the dependent variable (decision confidence). In doing this, we used the same reference category for each factor as described before.

Contextual Information Considering contextual factors, we verified that the service location and inference of personal information significantly impact people's confidence

[12]We used `ordinal` package, an R implementation of CLMM, for modeling privacy decision confidence.

Factor Impacting Privacy Decision Confidence	Coef β	SE(β)	z	p	OR (95% CI)
location:Private			*baseline*		
location:Work	-0.36	0.13	-2.86	**	**0.70** (0.54-0.89)
location:Public	-0.39	0.13	-3.08	**	**0.68** (0.53-0.87)
infer:DeviceID⇒UserLocation			*baseline*		
infer:Electricity⇒EnergyConsumptionPattern	-0.16	0.17	-0.95		0.85 (0.61-1.19)
infer:Motion⇒Presence	-0.02	0.17	-0.13		0.98 (0.71-1.36)
infer:Video⇒Presence	0.28	0.13	2.22	*	**1.33** (1.03-1.71)
infer:Voice⇒DeviceControlIntention	-0.14	0.17	-0.82		0.87 (0.62-1.21)
infer:Voice⇒Identity	-0.49	0.17	-2.93	**	**0.61** (0.44-0.85)
infer:Photo⇒Identity	0.20	0.10	1.95		1.22 (1.00-1.50)
infer:Video⇒Identity	0.18	0.13	1.41		1.20 (0.93-1.54)
infer:Video⇒PhysicalActivity	0.40	0.13	3.19	**	**1.50** (1.17-1.92)
infer:Photo⇒Emotion	0.23	0.17	1.35		1.26 (0.90-1.76)
infer:OBD⇒Nothing	-0.38	0.17	-2.26	*	**0.68** (0.49-0.95)
infer:Vital⇒Nothing	-0.33	0.17	-1.93		0.72 (0.52-1.01)
aware	0.69	0.13	5.27	***	**2.00** (1.55-2.59)
tradeoff	0.49	0.03	17.82	***	**1.64** (1.55-1.73)
segment:SomewhatInsensitive			*baseline*		
segment:SomewhatSensitive	-0.35	0.22	-1.62		0.70 (0.46-1.08)
segment:VerySensitive	1.21	0.18	6.77	***	**3.34** (2.36-4.73)
segment:Undecided	0.41	0.24	1.72		1.50 (0.95-2.39)
efficacy	1.41	0.38	3.76	***	**4.09** (1.96-8.54)
residence:00-04			*baseline*		
residence:05-10	-5.62	1.98	-2.84	**	**0.00** (0.00-0.18)
residence:11-20	-5.85	1.94	-3.02	**	**0.00** (0.00-0.13)
residence:>20	-5.17	1.83	-2.83	**	**0.01** (0.00-0.21)
residence:n/a	-3.40	2.07	-1.64		0.03 (0.00-1.94)
income:<$20000			*baseline*		
income:$20000-$34999	-0.77	0.27	-2.89	**	**0.46** (0.27-0.78)
income:$35000-$49999	-0.61	0.27	-2.27	*	**0.54** (0.32-0.92)
income:$50000-$74999	-0.51	0.26	-1.98	*	**0.60** (0.36-1.00)
income:$75000-$99999	-0.46	0.28	-1.64		0.63 (0.36-1.10)
income:>$100000	-0.21	0.31	-0.69		0.81 (0.44-1.48)
income:n/a	-2.54	0.62	-4.11	***	**0.08** (0.02-0.27)

Note: $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$

Table 5.8: Regression Results of CLMM model

levels in their privacy decisions.

- Our GLMM analysis has revealed that people tend to reject the IoT services operating in their private spaces. Through this analysis, we further confirmed that they make such decisions with confidence. To explain, people are 30% less likely to have high confidence in their privacy decisions toward IoT services at the workplace (`location:Work`) rather than private spaces like home (`location:Private`), no matter what decisions made ($p < .01$). People showed a similar reaction to the public spaces (`location:Public`); they make 32% less confident decisions in public rather than their own places ($p < .01$). This is probably because they have limited information (e.g., available sensors) about the services deployed at locations out of their control, therefore they are not fully convinced that they could make a good decision.
- Regarding the inference of personal information, we confirmed that people's privacy decision confidence varies depending on the specific types of collectible sensor data, namely video and voice. In general, people make confident decisions about the inferences based on video data. Specifically, the cumulative odds of the level of confidence in privacy decisions being rated in a certain level or above for IoT services inferring users' presence at a specific location through video data captured by surveillance cameras (`infer:Video`⇒`Presence`) are predicted to be 33% higher than inferring user location based on their unique device IDs (`infer:DeviceID`⇒`UserLocation`), such as Wi-Fi MAC addresses captured by Wi-Fi access points ($p < .05$). This indicates that people would be more certain about their decisions if the presented services are collecting sensor data which can clearly represent their presence in some locations (i.e., video). On the other hand, it is possible that some people do not fully understand how Wi-Fi MAC addresses can be used to infer their current location. In addition, sensor devices collecting video (i.e., cameras) are more easily recognizable in real life, thereby allowing

survey participants to better situate themselves in the scenario. In a similar vein, people tend to make confident decisions about the video-based inference of their physical activities (`infer:Video⇒PhysicalActivity`); the cumulative odds of having high decision confidence toward this type of inference are 50% higher than the baseline inference ($p < .01$).

- In contrast, people are relatively unconfident about the inference based on voice data. The cumulative odds of making a confident privacy decision of whether to accept the inference of user identity via speech analysis (`infer:Voice⇒Identity`) are predicted to be 39% lower than location inference based on device IDs ($p < .01$). People may perceive voice data itself as a resource that can reveal additional sensitive personal information (e.g., intention, emotion) rather than just identity. Accordingly, it is more difficult for them to make a correct decision if the inference is based on voice data, compared with seemingly primitive data like device IDs.

Pre-decisional Factors We previously showed that both privacy awareness and utility-privacy tradeoff are significant factors determining the type of privacy decisions people make. Here, we confirmed that these factors also have influences on how confident they are on their decisions. Our regression results indeed showed that people with high privacy awareness could better assess the privacy implications of the IoT service before using it, thereby making more confident decisions (**RQ2**). For one unit increase in the level of awareness of possible inferences of personal information by IoT services (`aware`), the cumulative probabilities of making confident privacy decisions increase by a factor of 2 ($p < .001$). We obtained a similar result about utility-privacy tradeoff; each one unit increase in the perception of privacy risks relative to utility benefits (`tradeoff`) will increase the level of decision confidence by a factor of 1.64 ($p < .001$). This means that people tend to make confident decisions if they believe the presented IoT services were risky. In this case, they are probably more cautious about the corresponding service and therefore are forced to make more careful and confident

decisions.

Personal Privacy Propensity Regarding privacy propensity, both privacy segment and privacy self-efficacy have proven to be significant factors impacting the confidence level of privacy decisions. In our previous GLMM analysis, we have confirmed that privacy-conscious people tend to make conservative decisions. And now, we found out that these people make such decisions with high confidence. For people who are very sensitive to privacy issues in IoT (`segment:VerySensitive`), the cumulative odds of making confident decisions to accept (or reject) the service are predicted to be 234% higher than those who are indifferent to privacy ($p < .001$). We also realized that privacy self-efficacy can be considered as a rough estimate of people's confidence in their decisions. For one unit increase in the level of privacy self-efficacy (`efficacy`), the cumulative odds of confident privacy decision-making increase by a factor of 4.09 ($p < .001$).

Demographic Information We uncovered the fact that some demographic information, such as the residence period in the US and annual income level, impacts people's decision confidence as well. Interestingly, a minority of people are more confident about making privacy decisions. To begin with, people who have lived in the US more than at least 5 years (e.g., `residence:>20`) are about 100% less likely to make confident decisions compared to those whose residence period is less than 4 years ($p < .01$). It indicates that people who recently moved to the US are generally knowledgeable about the properties of the IoT service which possibly can invade their privacy. Next, we found that people at the higher income level tend to make less confident decisions than low-income people. There is specifically a 40% chance of making a confident decision for people whose annual household income falls between \$50,000 and \$74,999 (`income:$50000-$74999`) than people who earn less than \$20,000 annually ($p < .05$). This means that low-income US residents are more concerned about the potential

privacy risks in IoT than others, thereby trying to make decisions as confidently as possible. This finding is consistent with [67] even though this study is targeted at traditional online privacy, not IoT privacy.

Through the abovementioned statistical analysis, we presented what and how underlying factors are associated with people's privacy decisions and their confidence in such decisions. We specifically confirmed that having a higher level of privacy awareness is positively correlated with the probabilities of making more *conservative* (**RQ1**) and *confident* (**RQ2**) decisions, namely better-informed privacy decision-making in IoT.

5.5 Prediction of Informed Privacy Decisions in IoT

In this section, we address the implications of the informed decision-making in the view of privacy decision support. We first investigate whether it is feasible to build machine learning models that accurately predict people's decisions based on their observed privacy behaviors. Additionally, we also analyze the relative importance of the input features (privacy decision factors) used to train the machine learning models. Lastly, we validate whether privacy decisions made with more confidence can yield the model with better predictive performance (**RQ3**).

5.5.1 Feature Importance on Privacy Decision Prediction

We used Random Forest (RF), an ensemble learning method for classification, as an algorithm for predicting privacy decisions. We chose RF because previous research showed that it was effective in helping improve the accuracy of user-defined privacy policies [85] but also known to be effective at generalizing to variants not seen in the training data (less prone to

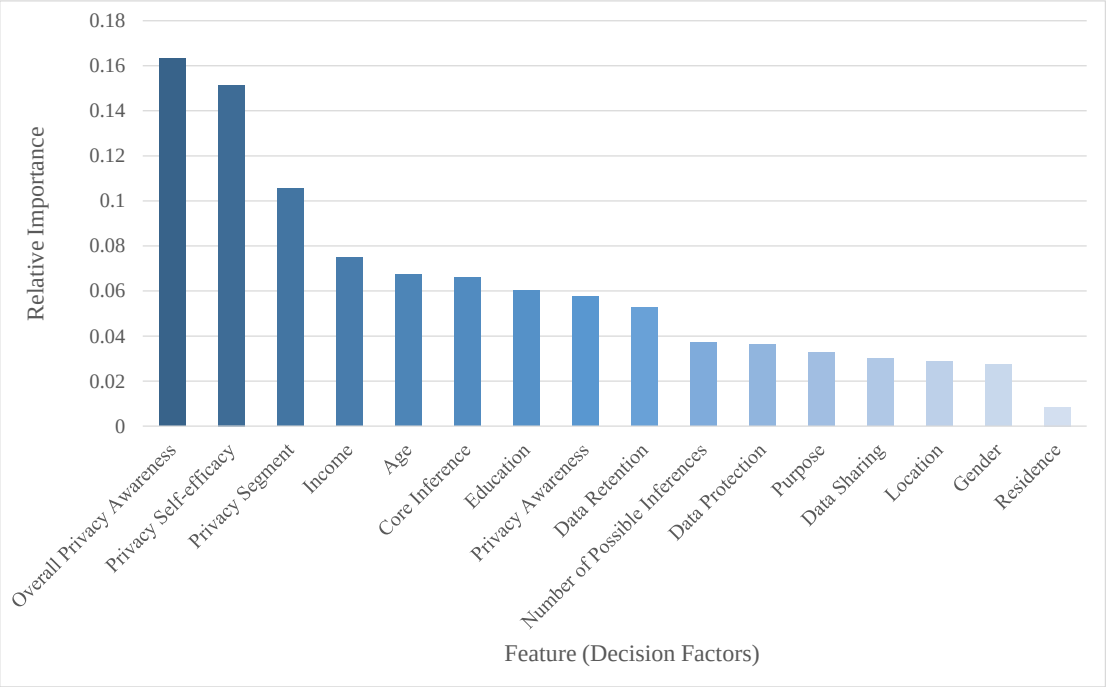

Figure 5.2: Feature Importance (Random Forest)
Note: 80%/20% random train/test split

over-fitting), unlike other popular algorithms such as decision trees [12]. Another advantage of RF is that the trained model can compute the relative importance of the input features in making a prediction. The feature importance can provide us with practical insights for constructing machine learning models for privacy decision support.

5.5.1.1 Experimental Setup

Predicting privacy decisions via an RF machine learning model can be generally defined as follows. Assume a dataset $D = \{(x_i, y_i)\}_{i=1}^{n}$ where x_i is the M-dimensional input feature vector ($x = [x_1, x_2, \ldots, x_M]$) and y_i is the decision to be predicted (1: use, 0: not use) for the i^{th} instance, respectively. Given dataset D, we train an RF classifier $f(x)$ based on a set of K decision tree classifiers $h(x) = \{h_1(x), \ldots, h_K(x)\}$ where $h(x)$ will be trained on various sub-samples of D. Here, the sub-sample size is the same as the size of original samples (D), but the samples are drawn at random with replacement (i.e., bootstrapping). During

the training, each tree ($h_k(x)$) randomly selects m features out of all possible M features ($m = \sqrt{M}$) and finds the best split on the selected m features, at each node. Thereby, each $h_k(x)$ has grown as an independent classifier with unique tree structure (e.g., list of nodes). For the prediction, each trained $h_k(x)$ casts a vote for the most plausible value (class) of y_i at input vector x_i, and then $f(x)$ finally determines the class with the most vote which wins as a final prediction result.

We trained the RF model[13] by using all attributes (features) described in Table 5.6 with the following exceptions. We excluded *subject ID* because we do not need to differentiate each participant (currently focusing on a one-size-fits-all model). We also excluded *decision confidence* as an input feature since we rather investigate the performance of the models trained on multiple datasets divided by decision confidence. In addition, we excluded *utility-privacy tradeoff* since it is strongly correlated with privacy decisions (see Section 5.4.1.3). Even though utility-privacy tradeoff could act as an important feature (*semi-final* decision), it is unrealistic to ask users to mark the ordinal values of utility-privacy tradeoffs whenever a prediction is to be generated. As a result, we came up with 16 input features for training the RF model.

Regarding a performance metric, we primarily used the area under the ROC curve (AUC) because it is unaffected by the class imbalance problem (there were about 40% accept and 60% reject decisions in our dataset); AUC is independent of the threshold applied to compute the probability of the binary classification results. Additionally, AUC itself is comprehensible; a random classifier has an AUC of 0.5 while a perfect classifier has an AUC of 1.0. Since our dataset is small (7,320 instances) and sparse (92% of the user-scenario matrix was empty[14]), we evaluated the trained models through 10-fold cross validation (CV) for deriving a more accurate estimate of their prediction performance.

[13]We used `scikit-learn`, python implementation of Random Forest, for all of our machine learning experiments.

[14]Each participant responded to about 8% of all available scenarios.

5.5.1.2 Experiment Results

The mean AUC of the trained RF model across all folds was 76.41% (± 3.03%). Even though this performance was not too bad (clearly better than random), it should be improved since even a single false prediction may cause undesired inferences of personal information, making users reluctant to follow privacy recommendations. We will discuss how we can achieve better predictive performance in a later section. Figure 5.2 displays the relative importance of the input features for the trained RF model. We used a randomly chosen 80%/20% training/test set to compute the feature importance. The results showed that personal privacy propensity (overall privacy awareness, privacy self-efficacy, privacy segment; listed in order of importance) is crucial in predicting people's privacy decisions in IoT. Furthermore, we found that user features (e.g., income, age) were more important than scenario features (e.g., core inference, data retention).

5.5.2 Decision Confidence and Predictive Performance

Since the decision-maker's confidence is a critical component in making rational judgements under uncertainty [38], we hypothesize that privacy decisions made with confidence show more clearly identifiable and predictable behavioral patterns of the user than unconfident decisions, including random choices. Accordingly, we assumed that an ML model trained on confident decisions would better capture these patterns, leading to superior predictive performance (**RQ3**).

5.5.2.1 Experimental Setup

In order to answer **RQ3**, we first divided our dataset by the five levels of decision confidence (see Table 5.9). The number of unconfident (286) and very unconfident (69) decisions is

Confidence Level	# Instances	# Accept Decisions	# Reject Decisions	Accept Ratio
Neutral and Unconfident	***1,402***	***715***	***687***	***51.00%***
Very Unconfident (1)	69	18	51	26.09%
Unconfident (2)	286	146	140	51.05%
Neutral (3)	1,047	551	496	52.63%
Confident (4)	**3,254**	**1,577**	**1,677**	**48.46%**
Very Confident (5)	**2,664**	**624**	**2,040**	**23.42%**
Total	7,320	2,916	4,404	39.84%

Table 5.9: Data Size and Class Distribution

very small compared to other types of decisions, and it is difficult to obtain a meaningful predictive performance with models trained on such a small amount of data. Therefore, we merged them with neutral (1,047) decisions into a *neutral and unconfident* (1,402) category. We finally trained three separate RF models based on these datasets (boldfaced rows in Table 5.9). Regarding the performance measurement, we used the mean AUC computed through 10-fold CV. We also presented classification accuracy, which is the number of correct predictions made divided by the total number of predictions made, as reference.

5.5.2.2 Experiment Results

Figure 5.3 presents the comparison of predictive performance of all the trained RF models; it clearly demonstrates that the performance increases with the level of decision confidence. Especially the RF model exclusively trained on very confident decisions resulted in the highest mean AUC[15] score of 86.87% (best model), which represents an increase of about 10% to the baseline model trained on all available decisions. The difference to all other models is statistically significant (paired t-test with Holms-Bonferroni correction; $p < .05$). This is an interesting finding because the size of training data used to construct the best model is approximately 64% smaller than the baseline model (see Table 5.9). This result implies that we may be able to build a highly accurate privacy recommendation engine even

[15]AUC is a performance metric which is immune to the class distribution of the given training set (23% vs. 40% accept ratio in very confident and all available decisions). AUC is considered as a more rigorous metric than accuracy.

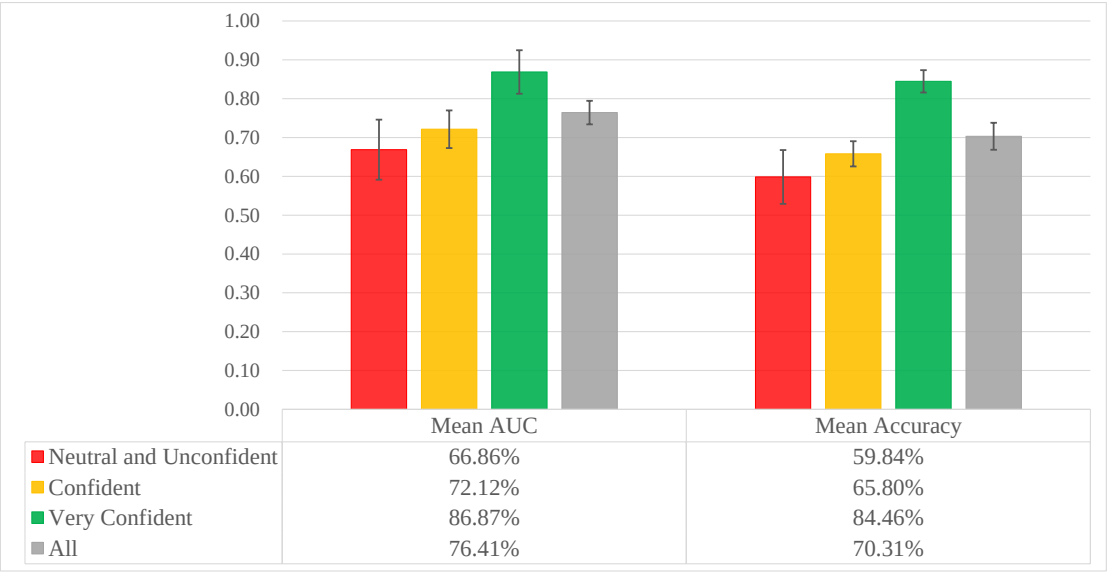

	Mean AUC	Mean Accuracy
Neutral and Unconfident	66.86%	59.84%
Confident	72.12%	65.80%
Very Confident	86.87%	84.46%
All	76.41%	70.31%

Figure 5.3: Privacy Prediction Performance per Decision Confidence (Random Forest)
Note: Error bars represent two standard deviations that account for about 95% of the mean AUC/accuracy (10-fold cross validation).

with a relatively small amount of human behavioral data, if we could ascertain that the data was collected from users who decide with confidence. On the other hand, the model based on neutral and unconfident decisions ranked the lowest (66.86%) with the largest variations ($\pm$ 7.73%) in the AUC scores calculated for each fold. The measurement of classification accuracy came up with similar results (see the right part of Figure 5.3).

To illustrate the relationships between decision confidence and model performance more clearly, we drew ROC curves for each trained RF model (see Figure 5.4). An ROC curve is constructed by plotting the trained model's true positive rate (TPR) against the false positive rate (FPR) for all the possible decision thresholds ($[0.0, 1.0]$). In other words, the ROC curve shows the tradeoff between sensitivity (TPR) and specificity ($1 - FPR$). As a baseline, a random classifier is expected to draw a 45-degree diagonal line ($TPR = FPR$). The closer the curve comes to the top left corner of the ROC space, the more accurate the classifier is, and vice versa. Note that AUC is the two-dimensional area underneath the entire ROC curve. It is also equivalent to the probability that the classifier will rank a randomly chosen positive example higher than a randomly chosen negative example.

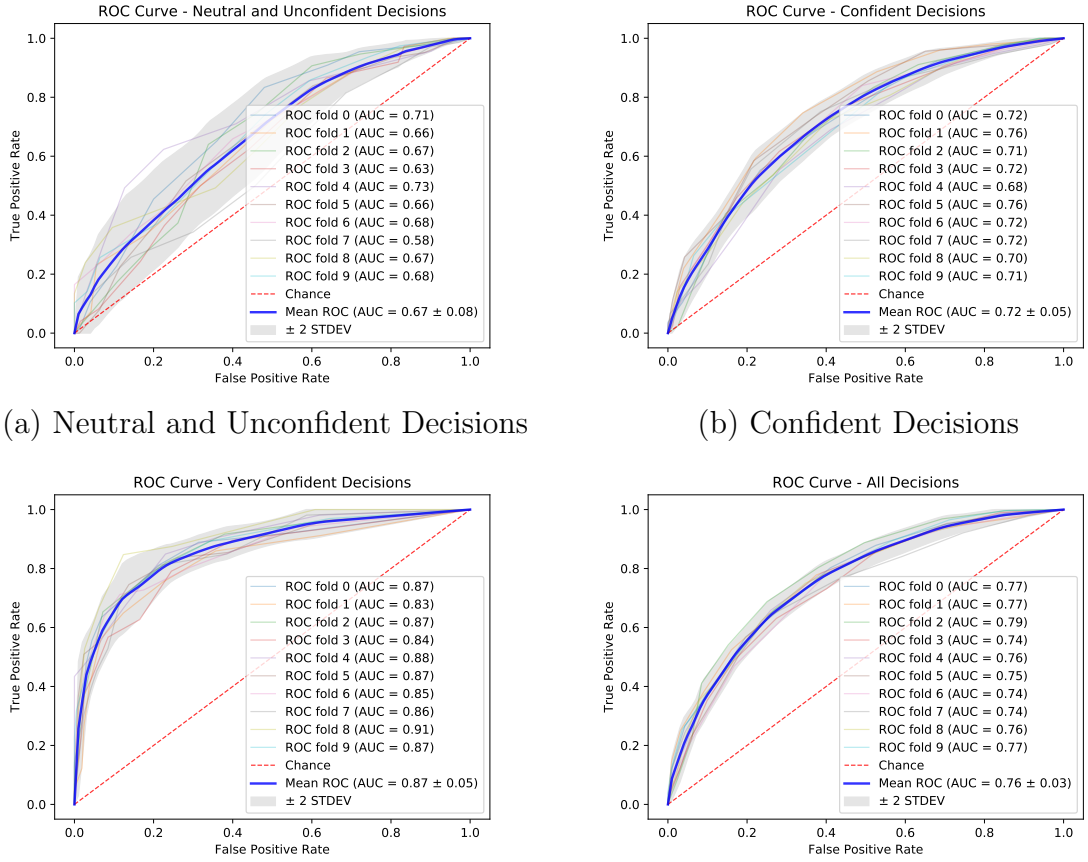

(a) Neutral and Unconfident Decisions

(b) Confident Decisions

(c) Very Confident Decisions

(d) All Decisions

Figure 5.4: Receiver Operating Characteristic (ROC) Curve

As already mentioned, the classifier trained on highly confident decision samples showed the best performance across all possible classification thresholds (Figure 4-c), not only due to their AUC but also due to their small variance. By contrast, the prediction of the classifier based on unconfident (or at least neutral) decisions was not only the most inaccurate but the most unreliable (Figure 4-a). The classifier which learned all available privacy decisions showed a reasonable performance with minimum variance (Figure 4-d). This is because the largest dataset was used to construct this prediction model.

Considering all the results of our experiments, we argue that it is important to put as many confident decision-making instances as possible into the machine learning model in order to maximize its predictive performance (**RQ3**). To that end, it is necessary to build systems

that help users in a user-friendly way to make more confident decisions. We will describe recommendations for such privacy-aware systems in the next section.

5.6 Toward Confident Privacy Decision-Making in IoT

In this section, we additionally confirm the causal relationships between privacy awareness and decision confidence, which further supports our answer to **RQ2**. As discussed in Section 5.5.2, enhancing user confidence in privacy decision-making is one of our research goals since it can enable highly effective privacy decision support. In pursuit of this goal, we then propose functional requirements for privacy-aware systems (PAS) to maximize IoT users' privacy awareness, thereby enabling them to make more confident privacy decisions. Furthermore, we describe ways to build a highly optimized privacy decision support system from the perspective of machine learning.

5.6.1 Privacy Awareness and Decision Confidence

In our statistical analysis, we confirmed that privacy awareness is positively associated with decision confidence. However, this does not necessarily mean they have explicit causal relationships. We performed a path analysis based on structural equation modeling (SEM)[16] to uncover any causal relations between (overall) privacy awareness, privacy self-efficacy, and decision confidence. We verified that there exists a strong chain of causation between privacy awareness, overall privacy awareness, and decision confidence (path coefficient $> .3$, see Figure 5.5). Note that privacy self-efficacy also has a strong causal relationship with decision confidence, but not with (overall) privacy awareness; this is probably because privacy self-efficacy measures self-perception (which might be delusional), while privacy awareness

[16]We used `lavaan` package, an R implementation of SEM, for performing path analysis.

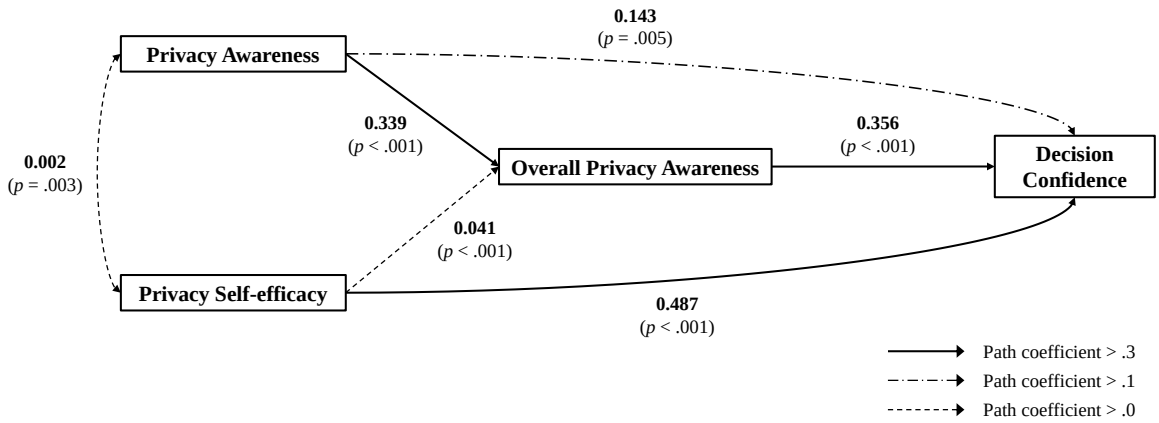

Figure 5.5: Path Analysis of Decision Confidence

measures true knowledge. Even though there exists a causal link between privacy self-efficacy and decision confidence, it is not easy to exploit this relationship since user-perceived efficacy does not change in a short period of time. Thus, we claim that increasing the level of privacy awareness of the users is a sufficient condition for nudging them to make more confident privacy decisions.

5.6.2 Advanced Privacy-Aware System

5.6.2.1 Augmentation of Privacy Awareness

To augment users' privacy awareness as much as possible, we propose functional requirements to better design and develop PAS in IoT environments. We suggest that PAS should inform users about both collectible sensor data and inferable personal information, as well as how the latter is inferred from the former. This is because we showed that both the type and confidence of privacy decisions are greatly impacted by inference information.

To that end, (1) developers of PAS first need to collaborate closely with IoT service providers to retrieve relevant information regarding the inference. For instance, a service provider of the Wi-Fi-based occupancy monitoring system can explicitly describe how they collect data

and infer personal information (here, user presence) because they have an ownership of the system.

Service providers may not describe all the possible inferences from specific sensor data. For instance, a service provider who is utilizing facial recognition software for the purpose of user authentication might be unaware of the fact that a similar technique can be used to infer users' sexual orientation from previously collected image data. Even if they are aware, the service provider is likely not to reveal this additional inference if they do not rely on it for operating the service. Nevertheless, for the sake of user privacy, the user should know the possibility of this inference [25]. Therefore, (2) developers of PAS may want to construct and update a knowledge base about the relationships between sensor data and inferable personal information. Specifically, the knowledge base should be continuously updated when a new inference technique becomes possible through recent technological advancement. This can be done by either manually (e.g., human experts) or possibly automatically (e.g., information retrieval system) analyzing literature relating to data mining and machine learning. It is also possible that different developers of PAS as well as the research community collaboratively maintain a single up-to-date knowledge base by sharing their knowledge with each other.

Next, PAS should deliver inference information to users in case they have no effective ways to communicate with an IoT service. Regarding this, (3) PAS should have user interfaces that convey factual information about all potential inferences in a user-friendly manner. For instance, an interactive visualization of the detailed procedures of both sensor data collection and personal information inference would be preferable for end users over the display of a long textual description.

5.6.2.2 Feature Extraction for Privacy Decision Support

Even though PAS can present users with information that will help them understand some privacy implications of using various IoT services, users still need to configure their privacy settings themselves. However, some users may have difficulties doing so due to limits in their available time, motivation, and cognitive decision-making abilities [92, 2]. Therefore, attempts are being made to integrate privacy decision support into PAS [32] for giving users privacy recommendations. In this context, our goal is to build a machine learning model as a core part of privacy decision support that learns and predicts users' privacy decision-making behaviors with a reasonable performance.

To achieve this objective, we would suggest that (4) PAS have some mechanisms to gauge each user's overall privacy awareness and privacy self-efficacy, since these factors were the most important features in our machine learning models. Our GLMM & CLMM analysis also confirmed that these variables are strongly associated with people's privacy decision-making. Regarding the measurement of overall privacy awareness, developers of PAS may consider utilizing the experience sampling method to repeatedly measure users' privacy awareness regarding the current IoT context, and later aggregate these measurements into a representative metric (e.g., modified mean). Given the fact that privacy self-efficacy does not frequently change, PAS may ask users to complete a one-time questionnaire (e.g. the scale by LaRose & Rifon [58]), and perform confirmatory factor analysis (CFA) to assign a level of privacy self-efficacy to each individual user.

5.7 Discussion

Thus far, we showed that making users aware of the privacy implications of using IoT services will guide them to make more conservative and confident decisions, which will be helpful for

protecting their privacy by themselves or with the help of privacy decision support systems. Yet, our work still has some issues that need to be considered and addressed.

5.7.1 Representability of Data

First, we need to consider the representativeness of the survey participants we recruited through Amazon MTurk. Previous research indicates that Amazon MTurkers are more privacy-sensitive than the general population [48]. This may result in a sampling bias that makes our results less general. Second, we analyzed participants' stated privacy attitudes and decisions on hypothetical IoT scenarios, and not their actual behavior captured in operational IoT environments. Although we tried to make the scenarios as realistic as possible, we do not definitively know how they would actually behave in real-world situations (i.e., intention-behavior gap). Lastly, we did not collect qualitative feedback which is often useful for gaining insights into underrepresented reasons, opinions, and motivations that people may have of their privacy decision-making. It may reveal other research questions or hypotheses that can inform future research directions. Therefore, it is necessary to develop experimental, but operational, IoT systems/services, recruit participants from the general population (e.g., including people who are not familiar with IoT), and gather their privacy decision-making behaviors as well as qualitative feedback about the IoT services that they interact with. We believe that this will provide us with a more comprehensive privacy decision datasets to understand and predict people's privacy decision-making in IoT.

5.7.2 Effective Measurement of Privacy Awareness

We consulted existing literature to define a set of *if-then* rules specifying possible and impossible inferences of personal information that can be drawn from sensor data available in IoT, and then used these rules to gauge each survey participant's level of privacy awareness.

However, we recognize that in practice, our claimed inferences could be open to misinterpretation. To be specific, one of our expert panel members insisted that it was hard to decide whether a specific inference is possible or not since there was no concrete definition of the quality of the inference (e.g., the desired accuracy of the inference results). Therefore, some participants may not accurately respond to our privacy awareness questions, which may yield to an inaccurate measurement. This suggests that the need for an in-depth study investigating the best way to define and express the inference of personal information, particularly in the context of IoT. In addition, we presented every participant with the same set of possible and impossible inferences (answer options) for each scenario. Displaying different answer options based on the correctness of each participant's responses to the previous scenarios (i.e., computerized adaptive testing) might be a quicker and more effective way to measure his/her privacy awareness. Lastly, our approach will necessarily have scalability issues as long as researchers have to manually generate or review the inference rules according to the literature. Considering the rapid advancements in data mining and machine learning, it might not be feasible to keep these rules updated. One possibility, albeit not ideal, is to build a probabilistic information retrieval system that automatically extracts relationships between entities (e.g., sensor data and personal information) in text data (e.g., academic papers). To build such a system, however, we still need labeled training data (i.e., known relationships between sensor data and personal information), which does not exist to our knowledge. Thus, utilizing systems such as Snorkel[17], designed for programmatically generating labeled training datasets from raw data without much human intervention, is one possible strategy to realize the automated creation and maintenance of the inference rules.

[17] `https://hazyresearch.github.io/snorkel/`

5.8 Conclusion

In this chapter, we investigated how people's privacy awareness impacts the type and confidence of their privacy decisions in IoT. Through an online survey ($N = 488$), we collected people's privacy decisions as well as their levels of privacy awareness toward hypothetical IoT services. Through statistical analysis (random-effects model), we confirmed that people who are well aware of the potential privacy risks of using IoT services tend to make conservative and confident decisions. We also validated that confident privacy decisions form a better basis for building machine learning models (Random Forest) that accurately learn and predict people's preferred decisions toward unseen IoT services. To the best of our knowledge, this is the first attempt to empirically verify the implications of informed privacy decision-making in preserving user privacy, especially in IoT environments. Based on our findings, we also proposed strategies to better design and develop a privacy-aware system (PAS) which aims not only to increase IoT users' privacy awareness but also to further assist their decision-making through machine learning-based privacy recommendations.

Chapter 6

Understanding Causes and Effects of Confident Privacy Decision-Making in IoT

Our previous research indicated that people's privacy awareness impacts their privacy decision-making in IoT environments. The more aware they become of privacy risks in using IoT services, the more likely they make informed and hence confident privacy decisions. We also showed that the collection of confident privacy decisions would enable us to build more accurate machine learning models that predict people's future privacy decisions. In this chapter, we investigate the user perceptions and practical implications of informed privacy decision-making in IoT, so that we can support our previous findings more concretely. To this end, we conducted a new online survey based on the protocol we developed in the previous chapter, while differentiating sample population (university students; $N = 43$) and additionally collecting their qualitative feedback regarding privacy decisions they made. Through qualitative analysis on the collected survey responses, we confirmed the fact that the enhancement of privacy awareness could allow users to make a confident privacy decision toward IoT. We

also verified that privacy decisions made with more confidence would facilitate more accurate privacy predictions, by conducting machine learning experiments with training and test data, which were separately collected from two different user groups.

6.1 Introduction

During the past decade, user privacy has become an important issue in networked computing environments. For instance, smartphone applications are increasingly asking users' permission to access various types of data (e.g., location) collected by the user's device. Without careful consideration and decision-making regarding this request, a significant amount of sensor data can be collected, aggregated and processed by external entities, thereby causing serious privacy violations afterward (e.g., the undesired inference of personal information). This privacy-invasive practice is likely to further increase with the proliferation of sensor devices in the era of the Internet of Things (IoT). In IoT environments, users are surrounded by various computing devices capable of unnoticeably monitoring and analyzing users' every move [86, 70, 97]. For example, smart CCTVs in a shopping mall may collect customers' facial photos without their consent, or may not even give any notice to them when collecting the data. More importantly, IoT service providers may exploit the collected data so as to infer sensitive personal information such as age, race, or sexual orientation through computer vision and machine learning technologies, for the purpose of user profiling or targeted marketing. Therefore, users in IoT may feel less aware and in control of personal information being collected or inferred, compared to those in a traditional mobile computing environment. For these reasons, there are ongoing research efforts not only to inform IoT users of details of sensor data collection (and personal information inference) but also to allow them to disable the collection of specific sensor data (i.e., privacy-aware systems [71, 81, 32]).

However, it is still unclear how the enhanced privacy awareness impacts people's privacy

decisions toward IoT services. In this regard, we previously carried out an online survey study with Amazon MTurkers ($N = 488$) to figure out the role of privacy awareness in the process of privacy decision-making in IoT (see Chapter 5). Through statistical analysis on the collected survey data, we had revealed the fact that a higher level of privacy awareness is highly related to the likelihood of privacy decisions being made more confidently. By conducting machine learning (ML) experiments, we had also shown that confident privacy decision samples (i.e., training data) yield more accurate privacy decision prediction models. Even though our previous work sheds light on the importance of informed privacy decision-making in IoT, we still do not know how people actually perceive the informed privacy properties of IoT services and then make confident decisions accordingly. This is because the results of our previous work were primarily derived from quantitative analysis (no qualitative feedback collected). In addition, we performed ML experiments solely on the dataset collected from a single user population, thereby raising an issue of generalizability of our findings.

To address these limitations, we conducted a new online survey study based on the survey protocol used in our previous study, while both varying the sample population (university students; $N = 43$) and asking survey participants open-ended questions about their perceptions of privacy awareness and confident decision-making in IoT. After the participants completed the main body of the survey, they were asked to describe their privacy attitudes to using smart devices, opinions on the measurement of privacy awareness, perceived effects of privacy awareness on both their use (or non-use) decisions and decision confidence, and thoughts about confident privacy decision-making in IoT. Qualitative analysis on the collected responses showed that a higher degree of privacy awareness could lead people to make more confident privacy decisions, which is consistent with our previous finding. We also confirmed that the perceived balance between utility benefits and privacy risks (i.e., utility-privacy tradeoff) is the strongest factor that influences IoT users' final privacy decisions. We also validated that the level of user-stated confidence in privacy decision samples impacts the performance of ML models trained to predict future decisions. To that end, we first

split the whole dataset collected from the previous study (MTurkers) based on its privacy decision confidence and then trained multiple ML models on the training sets divided by decision confidence. We then evaluated the predictive performance of the trained ML models against a new dataset gathered in the current study (university students), varying the ML algorithm. Experimental results indeed indicated that the ML models trained on confident privacy decisions outperform the models based on unconfident decisions; this also supports the finding from our previous research.

To sum up, our work makes the following contributions in the field of modeling and prediction of privacy decision-making in IoT:

- We identified privacy awareness as one of the underlying causes of confident privacy decision-making in IoT, by collecting and analyzing people's opinions about the perceived privacy risk awareness and the resulting privacy decisions toward realistic IoT service scenarios.
- We verified the practical implications of confident privacy decision-making for privacy decision support, namely that confident decision samples would yield more accurate privacy decision prediction models, through the cross-population machine learning experiments.

6.2 Related Work

There exist some previous research efforts that aim to interpret and understand people's privacy decision-making processes in the context of IoT. Some of these works also describe ML-based prediction of privacy decisions of the user in IoT.

Emami-Naeini et al. performed an online survey study with the aim of understanding and

predicting people's privacy decisions toward hypothetical IoT services that can exist around the user [73]. Through statistical analysis on the collected dataset, the authors found that people's privacy decision-making not only varied from person to person but was also highly context-dependent. For instance, they showed that the physical location where sensor data collection occurred, user-perceived benefit, and the type of collectible sensor data, significantly impacted IoT users' privacy decisions. To verify the feasibility of predicting privacy decisions, the authors trained AdaBoost classifiers capable of inferring users' decisions of whether to accept or reject the given IoT service. The average accuracy of their predictions was as high as 86%.

Apthorpe et al. recently proposed a survey methodology for discovering users' privacy norms in a smart home environment [7], based on the contextual integrity (CI) privacy framework [75]. To be specific, the authors developed hypothetical IoT services possibly operating at a user's home using the combinations of several CI information flow parameters (i.e., senders, recipients, attributes, subjects, transmission principle), then constructed questions asking survey participants about the acceptability of the given information flows in each scenario. A pilot study with 1,731 participants showed that it was feasible to effectively identify users' privacy norms against diverse smart home application/service scenarios (e.g., IoT device owners are more accepting of smart home information flows). The authors insisted that their methodology was easily applicable to other types of IoT contexts (e.g., smart office).

Page et al. conducted an interview study to investigate differences in people's perceptions and their resulting adoption (or non-adoption) decisions for IoT technologies [78]. The authors recruited 39 interviewees (19 pairs each consisting of young people aged 18-26 and their parent) and asked their attitudes, understanding, and usage of commercially available wearable (e.g., fitness tracker) and environmental (e.g., in-home voice assistant) IoT devices. Through qualitative analysis, they revealed the fact that two conceptual models (user-centric and agentic technology perspectives) exist that explain people's adoption decisions toward

IoT more than generational differences between young and old users. The authors stated that most of the consumer-oriented IoT devices on the market are well-suited for users who have an agentic view, but they might be problematic (e.g., privacy violations) for those who approached IoT with an user-centric perspective. They also proposed some design recommendations for IoT devices so as to support both technological perspectives of the user.

All of these previous works provide us with useful implications of how people perceive and react to diverse IoT devices and services. However, there is a lack of consideration for privacy awareness that could impact the quality of privacy decision-making (namely, confident or unconfident decisions) in IoT environments. Also, we found no research examining the effects of user confidence in privacy decision-making instances on the predictive performance of the ML models that are trained to predict future decisions. The primary aim of this study is to fill these gaps.

6.3 Data Collection and Preprocessing

In this section, we describe how we measured people's privacy awareness and collected their decision-making behaviors toward IoT services. Based on the survey protocol we developed in our previous research (see Chapter 5.3.1), we additionally asked the participants open-ended questions to gather their qualitative feedback about informed privacy decision-making in IoT. We also explain how we performed privacy propensity modeling, which was also proposed in the previous study (see Chapter 5.3.2), to extract additional user-specific features for better predicting people's privacy decisions.

6.3.1 Online Survey Study

Here, we explain how we conducted a new online survey in more detail. As discussed, we modified the original survey protocol by additionally collecting survey participants' subjective opinions regarding the underlying causes of privacy decisions they made during the survey. We also recruited university students as study subjects to check for any differences of privacy decision-making between different user groups (MTurkers had been recruited in the previous study).

6.3.1.1 IoT Service Scenario

In order to collect users' reactions and opinions about their privacy decision-making in IoT, we need to present them with actual or hypothetical IoT services. Regarding this, we utilized 180 service scenarios created by our previous research (see Chapter 5.3.1.1). We had determined several contextual factors, such as service location, purpose, collectible sensor data, inferable personal information, and data privacy policies, that definitize hypothetical, but realistic IoT services. Specifically, each IoT service is equipped with some possible and impossible inferences of personal information based on the given sensor data. This information allowed us to objectively measure the degree of privacy awareness of the study subjects. We presented 15 randomly chosen IoT scenarios to each survey participant while he/she took the main body of our survey. We will elaborate on this procedure in the following section.

6.3.1.2 Survey Protocol

Our online survey was conducted in the following sequence: (1) verifying respondents' understanding of IoT and example services (IoT comprehension), (2) determining each user's

privacy segment (privacy segmentation), (3) modeling privacy decision-making behaviors (privacy decision modeling), (4) analyzing user perceptions of informed privacy decision-making (privacy perception analysis), (5) verifying the attentiveness in taking the survey (attention check), (6) measuring privacy self-efficacy, and (7) collecting demographic information of the survey population (study subject profiling). Figure 6.1 displays the overall procedure of the survey.

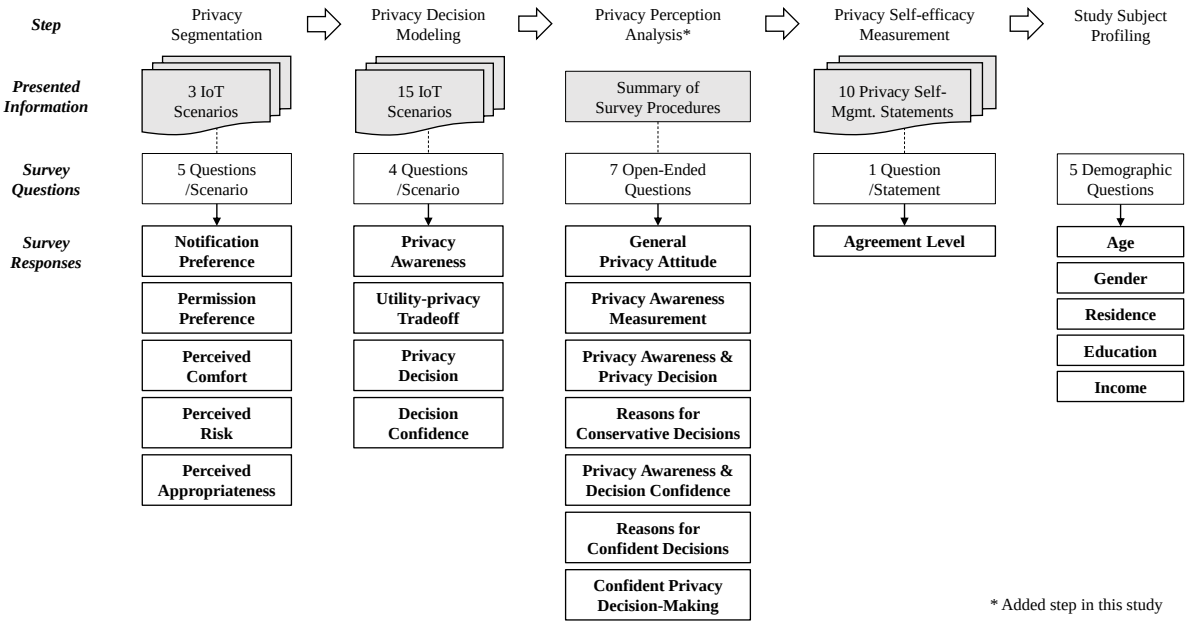

Figure 6.1: Survey Flowchart
Note: IoT comprehension (step 1) and attention check (step 5) are omitted for brevity.

Regarding privacy decision modeling (step 3), we gathered people's decisions of whether or not to use the IoT service, along with their pre/post-decisional factors, namely privacy awareness, utility-privacy tradeoff, and decision confidence. Each scenario is composed of a textual service description[1], an inference of personal information as a form of the *if-then* rule (e.g., *Photo* $\Rightarrow$ *Identity*), and its data privacy policy (data protection, retention, sharing). To gauge the privacy awareness of participants, we asked them to mark additionally possible inferences of personal information in the given scenario. We presented all the pre-defined

[1]E.g., At **your home** (*location*), your smart TV collects your **photo** to infer your **identity** (*inference*), thereby recommending TV shows based on your watching history, for your **convenience** (*purpose*).

possible and impossible inferences as answer options. Next, we measured participants' utility-privacy tradeoff using a 5-pt Likert scale (e.g., 5 indicates that benefits *are much less than* risks). Participants then make privacy decisions about whether or not to use the presented IoT scenario. Lastly, we inquired about their level of self-confidence regarding the decisions they made in the previous step on a 5-pt Likert scale (e.g., 5 indicates privacy decisions with *high confidence*).

Unlike the original survey design, we asked participants 7 open-ended questions regarding informed and confident privacy decision-making (privacy perception analysis; step 4), after they submitted responses to all the presented IoT scenarios in step 3. Since all of these questions are about each participant's reactions and experience to the previous step, we briefly reminded them what they had seen and answered. This was intended to figure out how the participants conceptualized privacy awareness in IoT and how it impacted the type and quality of privacy decisions they made. We summarized these questions as below (the full questionnaires can be found in Appendix C):

1. What is your opinion about potential privacy risks (e.g., *personal information inference*) in using smart devices?
2. How difficult or easy was it to answer *privacy awareness* questions?
3. How did your responses to *privacy awareness* questions impact your *privacy decisions*?
4. What were the reasons for making a *conservative* privacy decision?
5. How did your responses to *privacy awareness* questions impact your *privacy decision confidence*?
6. What were the reasons for making a *confident* privacy decision?
7. What is your opinion about *confident privacy decision-making* in IoT?

Privacy segmentation (step 2) and privacy self-efficacy measurement (step 6) will be explained in the following sections. Regarding the data type, all survey responses are categorical/ordinal except for the measured privacy awareness (numerical) and open-ended responses (text).

6.3.1.3 Survey Administration

We recruited 52 study participants on a university campus through e-mails in January 2019. Participants needed to be at least 18 years old, be proficient in English, and be affiliated with our institution (i.e., university students). Most of the recruited participants had a science and engineering background since we sent e-mails to students on an internal mailing list operated by the School of Information and Computer Sciences. The recruited participants were provided with a link to our online survey. All of them received a $5 Amazon e-gift card as compensation if they completed the survey. To assure the quality of survey responses, we discarded participants if they failed to pass IoT comprehension or the attention check question. We also eliminated some participants who finished the survey too rapidly (in less than 10 minutes, compared to the average response time of about 25 minutes). We thereby secured 43 qualified survey participants. Table 6.1 displays their demographics.

Age		Gender		US Residence (yr)		Education		Income	
18-24	34.9%	Male	44.2%	0-4	37.2%	Bachelor	39.5%	<$20,000	7.0%
25-34	**60.5%**	**Female**	**53.5%**	5-10	9.3%	**Master**	**53.5%**	**$20,000-$34,999**	**30.2%**
35-44	4.7%	n/a	2.3%	11-20	7.0%	Doctorate	4.7%	$35,000-$49,999	9.3%
				>20	**46.5%**	n/a	2.3%	$50,000-$74,999	16.3%
								$75,000-$99,999	2.3%
								>$100,000	7.0%
								n/a	27.9%

Table 6.1: Demographic Breakdown of Survey Population (UCI)

6.3.2 Privacy Propensity Modeling

Based on the collected survey responses, we extracted additional factors that represent people's personal privacy propensity, namely privacy segment, overall privacy awareness, and privacy self-efficacy.

6.3.2.1 Privacy Segment

By following an unsupervised clustering approach proposed by our previous works, we performed privacy segmentation as follows. We presented three representative IoT scenarios (different from the scenarios used in privacy decision modeling) and collected the reaction attitudes of our survey participants for each scenario (see Privacy Segmentation in Figure 6.1). Then we performed K-modes cluster analysis to segment users based on the commonality of their responses. Through the Elbow method, we uncovered that there exist at least three distinct privacy segments in our dataset ($K = 3$). Next, we ran the K-modes clustering algorithm on the dataset in order to determine each participant's cluster membership (privacy segment). Table 6.2 presents the cluster centroids learned from the user responses to the specific scenario. As it can be seen, each cluster is quite distinct, primarily in the perceived level of comfort, risk, and appropriateness: each cluster centroid has a unique combination of these reaction values. Therefore, we marked the clusters based on these three reactions (see Privacy Segment column in Table 6.2). We also statistically validated the distinctiveness of the clustering results by conducting Welch's t-tests. The tests confirm that the difference in the means of each reaction value between each pair of the resulting segments is statistically significant ($p < 0.025$, Bonferroni-corrected for two comparisons). Finally, we repeated this procedure for the remaining scenarios, and came up with similar results.

Ideally, the results of privacy segmentation should be independent of the scenario used for clustering. However, we noticed that some participants (about 19%) get clustered into three

Notification Preference	Permission Preference	Perceived Comfort	Perceived Risk	Perceived Appropriateness	Privacy Segment
Notify, always	**Allow**, always	Somewhat **comfortable**	Somewhat **risky**	Somewhat **appropriate**	***Somewhat Insensitive***
Notify, always	**Reject**, always	Somewhat **uncomfortable**	**Risky**	Somewhat **inappropriate**	***Somewhat Sensitive***
Notify, always	**Reject**, always	Very **uncomfortable**	Very **risky**	**Inappropriate**	***Very Sensitive***

Table 6.2: User Cluster Centroids for Sample IoT Scenario (UCI)

different privacy segments. Accordingly, we marked their segments as *undecided*. Otherwise, we adopted the majority voting approach to determine an individual's privacy segment. As a result, 12% of participants were classified as *somewhat privacy insensitive* ($N = 5$), 42% as *somewhat sensitive* ($N = 18$), and 28% as *very sensitive* ($N = 12$).

6.3.2.2 Overall Privacy Awareness

We also measured the privacy awareness of the participants by asking them to choose additional inferences of personal information that they believed to be true. For each of the 15 presented scenarios, we calculated a participant's privacy awareness as the ratio of correct answers. We used a modified mean of the measured scores after excluding the highest and lowest score as a measure of overall privacy awareness of the participant (population means 0.82, $SD = 0.14$).

6.3.2.3 Privacy Self-efficacy

To quantify people's self-confidence or belief of abilities in safeguarding their privacy by themselves, we performed confirmatory factor analysis (CFA) on the survey responses toward LaRose & Rifon's privacy self-efficacy scale. We first fitted a full CFA model on our dataset toward all the presented privacy management statements (i.e., indicators). The full CFA

model showed that 4 indicators should be dropped since they have unsatisfactory factor loadings < 0.7. Considering only the indicators with satisfactory factor loadings, we fitted a reduced CFA model with an acceptable convergent validity ($AVE = 0.59$). Based on the fitted reduced model, we computed the factor score for each participant. We then treated the computed factor score as the measurement of his/her privacy self-efficacy. The measured values of privacy self-efficacy range from -1.5 to 1.01 ($M = 0.00$, $SD = 0.53$).

ID[a]	Item (from LaRose & Rifon 2007)	FL[b]	R^2
PSE1	**It's easy to figure out which sites you can trust on the Internet.**	**0.70**	**0.49**
PSE2	**I am confident I know how to protect my credit card information online.**	**0.73**	**0.53**
PSE3	I know how to identify sites with secure servers.	0.34	0.12
PSE4	I know how to evaluate online privacy policies.	0.60	0.36
PSE5	**It's easy to set up dummy email account to shield my identity.**	**0.72**	**0.52**
PSE6	**I know how to change the security settings of my browser to increase privacy.**	**0.81**	**0.66**
PSE7	**I know how to use a virus-scanning program.**	**0.71**	**0.51**
PSE8	I am able to protect myself against the release of personal information.	0.60	0.35
PSE9	I know how to block unwanted E-mails.	0.50	0.25
PSE10	**Overall, I am confident that I can protect my privacy online.**	**0.91**	**0.82**

[a]Indicator
[b]Factor Loading

Table 6.3: Confirmatory Factor Analysis for Privacy Self-efficacy (UCI)

6.3.3 Dataset Summary

Through the abovementioned survey procedures, we collected people's diverse attitudes, reactions, and subjective opinions toward the presented IoT service scenarios. We utilized them to derive the causes and effects of confident privacy decision-making in IoT via both qualitative analysis and machine learning experiments.

To perform qualitative analysis, we isolated all responses to open-ended questions in its original form. As 43 survey participants individually responded to 7 questions, we came up with 301 qualitative responses. For machine learning, we prepared a dataset with 19 attributes that characterize both the IoT-enabled spaces and users who interact with them (see Table 6.4). There exist 7 attributes with contextual information about IoT services,

4 attributes for users' privacy decisions and pre/post-decisional factors (privacy decision-making behavior), 3 attributes for indicating their privacy propensity, and 5 attributes for demographic information. 15 attributes are categorical/ordinal and 4 numerical. Since each participant responded to 15 scenarios, we came up with 645 user-scenario instances. There were 284 accept decisions (*use* the service) and 361 rejects. The dataset will be fed into ML models to investigate the effects of the decision confidence on the prediction of future decisions. Like the previous work, we excluded *utility-privacy tradeoff* from the input features for building ML models since it only constitutes a semi-final privacy decision. We also ruled *decision confidence* out of input features since we utilized it as a criterion for generating separate training sets with different levels of confidence, thereby constructing multiple ML models accordingly. Regarding the numerical attributes, we performed min-max normalization in order to rescale the range of original values to $[0, 1]$, thereby achieving better fitted predictive models.

6.4 User Perceptions of Confident Privacy Decision Making

In this section, we describe how we analyzed the collected qualitative responses in more detail. We outline the overall analysis process, summarize findings for each open-ended question, and derive insights for understanding the causes and effects of confident privacy decision-making in IoT.

6.4.1 Analysis Procedure

We first divided each open-ended question into 2 to 3 sub-questions. For each sub-question, we identified response categories by grouping similar answers. We then calculated the fre-

Category	Attribute	Type	Range	Note
Contextual Information	Location	Categorical	3-class	Feature
	Purpose	Categorical	4-class	Feature
	Core Inference	Categorical	12-class	Feature
	Data Protection	Categorical	2-class	Feature
	Data Retention	Categorical	3-class	Feature
	Data Sharing	Categorical	2-class	Feature
	# Possible Inferences	Numerical	[0,1]	Feature
Privacy Decision-making Behavior	Privacy Awareness	Numerical	[0,1]	Feature
	Utility-privacy Tradeoff	Ordinal	5-class	*Excluded*
	Privacy Decision	Binary	2-class	**Label**
	Decision Confidence	Ordinal	5-class	***Excluded***
Personal Privacy Propensity	Overall Privacy Awareness	Numerical	[0,1]	Feature
	Privacy Segment	Categorical	4-class	Feature
	Privacy Self-efficacy	Numerical	[0,1]	Feature
Demographic Information	Age	Categorical	3-class	Feature
	Gender	Categorical	3-class	Feature
	Residence	Categorical	4-class	Feature
	Education	Categorical	4-class	Feature
	Income	Categorical	7-class	Feature

Table 6.4: Dataset Summary (UCI)

quency distribution of answers under each response category. In doing this, we did not alter the answers given by the survey participants.

6.4.2 Analysis Results

6.4.2.1 General Privacy Attitude

In the first question, we intended to grasp the participants' general privacy attitudes toward the use of their smart devices (e.g., face unlock phone). Since smart devices are typically designed to collect and analyze sensor data (e.g., facial photo) to provide the user with an intelligent function, there exist potential privacy risks accordingly. We aimed to understand how much our study subjects care about their privacy and how it impacts their actual usage

of the device.

Rephrased Question #1	Answer	Ratio
Have you ever cared about the undesired disclosure of personal information before using a smart device (e.g., face unlock phone)?	**Yes**	**67.4%**
	No	27.9%
	I do not know	4.7%
If you answered *Yes* above, how did your privacy concerns impact your usage of the smart device?	**I do not use**	**55.17%**
	It depends	31.03%
	I still use	3.45%
	I do not know	10.34%

Table 6.5: Response Summary – General Privacy Attitude
Note: "I do not know" includes no answer (n/a); this applies to other tables in Chapter 6.4.2.

As can be seen in Table 6.5, about 67% of the participants consider potential privacy threats that can be caused by a smart device before start using it. More than half of these participants also mentioned that they are not willing to use the device if the perceived privacy risks exist. For instance, a participant said:

> "Yes, I am aware of this although it is hard to track with smarter devices. Now, for example, I tend to check in Settings that apps and programs have only the desired access to files and pictures. I also tend to erase history in Google by hand because it tends to save it. Also, when I'm talking about something important, I'll put away my phone because I know that Google Assistant and other things are listening and I can't control them 100%." [P11]

Very few (3.45%) participants mentioned that they would use the device even if they realized the risks. A participant said:

> "I have thought about it, especially considering that a lot of that information is shared with unspecified third parties. It has not drastically affected how I interact with my devices, as it's ultimately unlikely to have a direct impact on how I perceive my life." [P29]

On the other hand, approximately one-third of the participants who have privacy concerns responded that they would make use (or non-use) decision depending on other factors (e.g., benefits). For example, a participant said:

> "I rarely utilize intelligent functionalities of smart devices because I do not think they are safeguarding my data or my privacy, I'm not convinced they are accurate, and I find them to be socially intrusive. However, there are times when I'm willing to sacrifice one form of safety for the convenience of another – for example, asking my phone to look something up or write a text while I'm driving." [P19]

Through this analysis, we noticed that our participants are privacy-sensitive in general, but it is also possible that they would accept the privacy risks for the sake of their needs.

6.4.2.2 Privacy Awareness Measurement

In question #2, we checked whether our strategy for measuring the degree of privacy awareness was appropriate or not. Since we had confirmed that the measured privacy awareness meaningfully impacts people's privacy decision-making, we wanted to make sure that privacy awareness questions were neither too difficult nor too easy to answer, thereby securing accurate and reliable privacy awareness measurement.

Table 6.6 indicates that about 70% of the participants thought the privacy awareness questions were easy[2]. Given the fact that our participants are undergraduate/graduate students majoring in computer science, human-computer interaction, and statistics, they were probably familiar with various technologies used to infer personal information from sensor data. This might be the reason that they perceived the questions were easy. In fact, about 43%

[2]Actual mean score was 0.82 with a standard deviation of 0.14.

Rephrased Question #2	Answer	Ratio
How difficult or easy was it to answer privacy awareness questions?	Difficult	18.6%
	Medium	9.3%
	Easy	**69.77%**
	I do not know	2.33%
Why was it *easy*?	**Prior knowledge/experience**	**43.33%**
	Contextual information (e.g., data sharing)	26.67%
	I assumed all inferences are possible	26.67%
	I do not know	3.33%
Why was it *difficult*?	**Lack of prior knowledge/experience**	**75%**
	No details about sensor data (e.g., data size)	25%

Table 6.6: Response Summary – Privacy Awareness Measurement

of the participants stated that they determined the possible inferences based on their prior knowledge or experience. A participant said:

> "I thought it was reasonably easy although I wasn't sure about some possible inferences. I've seen demo machines at the school I went to that showed how machine vision can infer your emotional state and read books about how big data can make accurate associations/predictions if it learns on a big enough data set." [P36]

Some of them (about 27%) also mentioned that the given information (e.g., data privacy policies) was helpful to find out answers. For instance, a participant said:

> "It was easy to answer the questions. This is due to the fact that all the details pertaining to the data were mentioned beforehand like time period the data will be shared, whether it is prone to unauthorized usage and its sharing policy." [P23]

Conversely, about 19% of the participants found the questions difficult to answer because they had no prior knowledge or insufficiently detailed information about sensor data collection (e.g., how big the collected data would be). Participants said:

"It was difficult because I am not familiar with that type of data analysis or how advanced (for example) computer vision algorithms have become. I guessed a lot. I think I was pretty accurate but I'm not sure, and I am sure my knowledge is out of date." [P27]

"Fairly difficult. The answer depends on what existing other information is correlated, parameters of the sensor, and whether you believe the stated use purpose." [P15]

We now recognized that the presented privacy awareness questions were generally easy for university students who had a science and engineering background. Therefore, we believe that the measured privacy awareness score adequately quantifies the study subjects' level of understanding of potential privacy risks in IoT environments.

6.4.2.3 Privacy Awareness and Privacy Decision

Through statistical analysis, we had shown in our previous study that people who were well aware of potential privacy risks in a given IoT scenario tended to make more conservative (non-use) privacy decisions (see Chapter 5.4.1.3). Here, we aimed to verify this claim by explicitly asking our participants their opinions regarding the relationships between privacy awareness and privacy decision.

Rephrased Question #3	Answer	Ratio
Did your answers to privacy awareness questions impact your privacy decisions?	**Yes**	**69.77%**
	No	18.6%
	I do not know	11.63%
If you answered *Yes* above, how did it impact your privacy decisions?	**More awareness, more conservative decision**	**70%**
	It made me think twice before making decisions	26.66%
	I do not know	3.33%

Table 6.7: Response Summary – Privacy Awareness and Privacy Decision

As expected, privacy awareness impacts privacy decision-making in a way that we confirmed in our previous study (see Table 6.7). About 70% of the participants responded that their answers to privacy awareness questions (i.e., privacy awareness measurement) influenced their privacy decisions. A majority of these participants stated that they were less likely to use the presented IoT service if they became more aware of potential privacy threats. Participants said:

> "Yes, most definitely. The more sensitive the inferences made in the area of privacy awareness, the less likely I would choose to use it." [P01]

> "Of course! I read a lot of privacy and safety articles. If I weren't knowledgeable in the area, it would be easier to give my data without meaning to or sub-estimating the real risks. Also, a lot of the devices lack proper privacy protocols and security protocols, so it makes me more weary. I have also been in risky places where simple things like name, address and phone numbers can be used to damage someone, so I guess it influences my privacy decisions." [P11]

Some of these participants also mentioned that privacy awareness measurement allowed them to think twice, even if they finally made a permissive (use) decision toward IoT services.

About 19% of the participants insisted that the level of privacy awareness did not significantly impact their decision-making. They mentioned that there existed more important factors (e.g., data privacy, purpose) for making a privacy decision. For example, participants said:

> "What the IoT was tracking was less relevant to me than who was tracking the data, why, and the level of security around the resulting data." [P19]

> "I think what really impacted my choices is whether I had a choice. All of the

governmental, airport questions ... do I really have a choice to participate or not?" [P25]

In summary, privacy risk awareness is one of the critical factors impacting people's privacy decision-making. To explain this phenomenon, we showed that the degree of privacy awareness is generally proportional to the probability of conservative privacy decisions, through the analysis of open-ended responses to this question. This finding is well aligned with our previous quantitative research.

6.4.2.4 Reasons for Conservative Decision-Making

There could also exist other privacy decision factors other than privacy awareness as addressed by P19 & P25. In our previous study, we had shown that service location (e.g., workplace), the presented inference of personal information (e.g., *Photo* ⇒ *Identity*), data privacy (e.g., data shared with third parties), and utility-privacy tradeoff (e.g., benefits *are much less than* risks) significantly influenced people's decisions on whether to allow or reject the IoT service scenario. Here we asked our survey participants for their views on the causal factors for conservative privacy decision-making.

Answer for Question #4	Ratio
I thought that privacy risks are greater than utility benefits	**41.86%**
I just felt uncomfortable	25.58%
I had concerns regarding data privacy (retention, protection, sharing)	23.26%
IoT can eventually infer all my personal information	4.65%
It was hard to evaluate utility-privacy tradeoff, so I just decided not to use	2.33%
IoT service was owned or operated by the entity which I do not trust	2.33%

Table 6.8: Response Summary – Reasons for Conservative Decision-Making

Table 6.8 displays a list of response categories ordered by the number of answers within each category. As can be seen, the perceived balance between privacy risks and utility benefits was the most important factor for conservative privacy decisions. Since the survey participants

evaluated utility-privacy tradeoff right before making a final decision, it can actually be considered as a *semi-final* decision. We also believe that privacy awareness measurement helped remind the participants of potential privacy risks, thereby allowing them to better judge the tradeoff. For instance, a participant said:

> "Trade-offs not being worthwhile. Most of these functions can easily be implemented with much less intrusive tools – for instance, RFID could be used instead of video to verify your identity and it is less risky in my opinion." [P28]

About 26% of the participants pointed out the discomfort of being monitored by IoT as a cause of conservative privacy decision-making. Generally speaking, these participants felt uncomfortable having personal data being recorded, analyzed, and used by external entities, regardless of the benefits they could gain. A participant said:

> "I don't want to *feel* watched. I know that I'm being tracked on my phone, my Apple Watch, etc, but as soon as you *feel* watched, and you are aware that you're being monitored or watched, it becomes way more uncomfortable. I also don't ever want to feel like my employer is watching/tracking me because that's just a boundary I'm not comfortable crossing." [P37]

Another 24% of the participants stated that data privacy policies primarily impacted their conservative decision-making behaviors. They worried about the situation in which the service provider does not protect the collected data (e.g., without encrypting data) or share it with external entities unknown to themselves. A participant said:

> "I was mostly concerned with third party sharing and lack of protection. I don't want my data to be used against me by hacker/attackers and I don't want to be

> aggressively marketed to by companies that have all of my information and can manufacture desire or catch me when I'm in a vulnerable mood." [P36]

A small portion (about 5%) of the participants were very skeptical about privacy protection mechanisms in IoT environments, then decided not to use the service. They believe that it is hard to avoid the collection or inference of sensitive personal information since end users generally have limited controls over the entire data processing protocol in an IoT system. To explain, participants said:

> "When there is a seemly complete lack of privacy (almost all information can be inferred)." [P22]

> "Even if you trust a particular device, do you trust its maker? The middleware maker? The data broker? No; trust requires trusting the whole chain, and in IoT that is not possible." [P25]

To sum up, most people tend to make a conservative privacy decision when they found that the risks are greater than the benefits of using the IoT service. The (privacy) risks are open to multiple interpretations: an unreasonable practice of retaining, protecting, or sharing the collected data, too many possible inferences of personal information, and low trustworthiness of the entity who operates the service. We believe that those who did not specify the reason for conservative decision-making ("I just felt uncomfortable"; about 26%) would have identified these privacy risks or similar negative consequences.

6.4.2.5 Privacy Awareness and Decision Confidence

According to our previous work, the degree of people's privacy awareness has been shown to have a positive correlation with the confidence level of their resulting privacy decisions. In

this question, we aimed to double check this finding.

Rephrased Question #5	Answer	Ratio
Did your answers to privacy awareness questions impact your privacy decision confidence?	**Yes**	**67.44%**
	No	18.61%
	I do not know	13.95%
If you answered *Yes* above, how did it impact your privacy decision confidence?	**More awareness, more confident decision**	**62.06%**
	Less awareness, less confident decision	**10.34%**
	I do not know	27.59%

Table 6.9: Response Summary – Privacy Awareness and Decision Confidence

As can be seen in Table 6.9, privacy awareness impacts privacy decision confidence in a way that aligned with our previous study; if people are made more aware of privacy risks, then they may be more confident in their privacy decisions (or vice versa). Around 67% of the participants stated that their responses to privacy awareness questions impacted their privacy decision confidence, and a majority of them agreed with our previous findings. For example, participants said:

> "I felt that my responses to the privacy awareness questions impacted my privacy decision confidence. Knowing all of the potential information that can be inferred from the data allowed me to weigh the value of the smart feature with my personal values on whether or not I would confidently use that feature or service." [P30]

> "Yes. If I know/suspect that many other things may be inferred from the data, and I don't want that to be a possibility, it is straight forward to decide on whether or not to opt in for the service." [P41]

There are participants (about 19%) who held opposite opinions to this. In general, they thought that personal information inference would be fine if the IoT service handled the sensor data and inferred information in a privacy-preserving manner. A participant said:

"No not by much, the facts about the data had more weight than the privacy awareness questions. As long as the data stored securely and deleted promptly and shared with no one, it didn't matter what info got collected." [P23]

We now confirmed that privacy awareness is one of the key factors of letting people make a confident privacy decision toward IoT services. Specifically, we showed that the enhancement of privacy awareness would induce more confident decision-making behaviors.

6.4.2.6 Factors Impacting Decision Confidence

In question #6, we asked survey participants about factors that impacted their confident privacy decision-making.

Answer for Question #6	Ratio
Utility-privacy tradeoff was clear to me	**28.33%**
I was informed about data privacy (retention, protection, sharing)	26.67%
I thought about what the additionally possible inferences are	16.67%
I do not know	11.67%
IoT service was owned or operated by a specific entity (e.g., government)	6.67%
IoT service was operated at the specific location (e.g., airport)	6.67%
I have prior knowledge/experience regarding IoT	3.33%

Table 6.10: Response Summary – Reasons for Confident Decision-Making

Similar to question #4 (see Chapter 6.4.2.4), the participants pointed out both utility-privacy tradeoff and data privacy policies as important factors. As can be seen in Table 6.10, people could make a confident privacy decision when they reasonably balanced the tradeoff (about 28%) or acknowledged details about data privacy (about 27%). To be specific, participants said:

"I would be more confident if I knew that the scenario would greatly improve my lifestyle and day-to-day functions." [P14]

> "If I am aware that data collected will be deleted soon, not shared and protected against unauthorized access, I was more confident." [P26]

About 17% of the participants made a confident decision if it seemed that the IoT service inferred too many personal information without a reasonable justification. A participant said:

> "I was confident when the scenario seemed to be inferring too much about individual." [P09]

There also exist some contextual factors such as who is providing the IoT service and where the service is actually operating. It is also true that these factors eventually impacted the user-perceived balance between privacy risks and utility benefits. Participants said:

> "I was confident about decisions that involved the government or my employer, and also about the scenarios I would actually find useful (health, car monitoring)" [P02]

> "The less personal data out there the better. So, technology in home is not necessary for me even if it is convenient. So, I had strong confidence about my decisions in those scenarios. I had less confidence in the airport scenarios/safety scenarios because they involve more individuals than just myself. So, I don't really know if the benefits outweigh the risks in those scenarios." [P07]

Lastly, if participants had technological knowledge regarding IoT, they could make a privacy decision more confidently. A participant said:

> "I already use IoT devices, and know people who do research in IoT, so I am generally confident in using them. I'm also aware of the machine learning methods that go into a lot of IoT applications which makes me more comfortable with them." [P40]

Through this analysis, we realized that people were more likely to be confident in their privacy decisions if they were informed about diverse privacy implications (that includes the perception of possible inferences of personal information) or already had prior knowledge or experience with IoT.

6.4.2.7 Confident Privacy Decision-Making

Finally, we questioned survey participants on their opinions and attitudes about confident privacy decision-making in IoT. We first asked if it was important (or not) for safeguarding their privacy in IoT environments. We also asked them to explain why.

Rephrased Question #7	Answer	Ratio
Do you think confident privacy decision-making is important?	**Yes**	**65.12%**
	It depends	9.3%
	No	16.28%
	I do not know	9.3%
If you answered *Yes* above, why is it important?	**I make non-use decisions after careful consideration**	**82.14%**
	Accountable decision-making is important anyway	17.86%
If you answered *No* above, why is it unimportant?	IoT service providers can do anything they want	50%
	My mind and situation will change	50%

Table 6.11: Response Summary – Confident Privacy Decision-Making

Table 6.11 presents the summarized responses. About 65% of the participants thought that confident decision-making is important in terms of privacy protection in IoT. Most of them insisted that confidence came from how they truly understood privacy implications in using IoT, and usually it led to their conservative decision-making (i.e., lowered privacy risk). A participant said:

> "Yes. A conscious choice between achieved benefit and privacy implications is important to me. It's binary to me: if I don't think the service is (very) beneficial, I won't use it because there is a chance that the collected data may be misused (e.g., increased health insurance price due to unhealthy habits)." [P41]

Some of them considered confident decision-making important regardless of the type of decision-making (i.e., conservative vs. permissive decision). A participant said:

> "Yes, making a confident decision means you are very well aware of the risks and are willing to embrace/not embrace the risks." [P14]

On the other hand, about 16% of the participants did not think their confidence mattered. It seemed that they did not trust IoT service providers in general. Also, the participants mentioned that IoT technologies and business practices are rapidly changing and evolving, thereby making their decisions meaningless. To explain, participants said:

> "In this day and age privacy is a myth so it doesn't matter what my decision is. Corporations and people want data and buy and sell data to make money, and to improve IoT services, to make more money." [P40]

> "I don't think my confidence has anything to do with protecting my privacy, I think that IoT service providers should be explicit about what data they are collecting and how it CAN be used, not necessarily what they claim to use it for. I could be confident today about how a service uses my data, but the technology or the company's decisions could change tomorrow, making my confidence irrelevant." [P27]

Considering all these responses, we claim that confident privacy decision-making is important for lowering the possibility of privacy violations that can be caused by IoT services.

6.4.3 Gained Insights

Here, we summarize implications gained from the collected qualitative feedback.

First, our survey participants (university students who major in Information and Computer Sciences) are sensitive to their privacy regarding the usage of smart devices. Accordingly they gave us diverse opinions about the relationships between privacy awareness and decision-making behaviors. In addition, they were comfortable for answering privacy awareness questions because they were familiar with IoT technologies that used to infer personal information from raw sensor data.

Regarding the impact of privacy awareness on the type and quality of privacy decision-making, we found that people who were more aware of possible inferences of personal information tended to make more conservative and confident privacy decisions. Participants also stated that both the perceived balance between privacy risks and utility benefits and data privacy policies were the factors meaningfully impacted their privacy decision-making. These results are consistent with our previous findings based on a quantitative method with a different sample population (Amazon MTurkers).

Lastly, most of the participants mentioned that confident privacy decision-making was important not only for minimizing the negative outcomes of using IoT services but also for making an accountable privacy decision by themselves.

6.5 Impacts of Decision Confidence on Privacy Decision Prediction

In this section, we perform machine learning (ML) experiments to validate our previous finding, namely that privacy decision samples made with more user confidence would enable

more accurate predictions of future privacy decisions (i.e., privacy decision support). An accurate and reliable privacy decision support is important since a single false prediction can cause serious privacy violations that harm users' IoT experience.

6.5.1 Motivation

In our previous study, we had experimented with *confidence-wise* prediction of privacy decisions in IoT (see Chapter 5.5.2). We constructed and evaluated three ML models using a privacy behavior dataset, which was divided by user-stated decision confidence. In doing this, we adopted 10-fold cross validation (CV) as a model evaluation method since we could not find publicly available datasets that we can use as test data. In other words, we trained and tested the ML models solely based on our own dataset collected from a single user population. However, it would be more desirable to employ the pre-trained ML models for predicting privacy decisions of new users who have different characteristics (i.e., cross-population prediction). No doubt it is necessary to make privacy predictions with reasonable accuracy in this scenario. More importantly, we need to check whether decision confidence works in the same way as we found in our previous study, under the condition that training and test data derived from two distinct user groups. In this study, we therefore trained ML models by utilizing the previously collected dataset (training data from 488 Amazon MTurkers) and measured their predictive performance on the newly collected dataset (test data from 43 university students).

6.5.2 Experimental Setup

We performed feature engineering on both training and test instances to create pairs of input features along with ground-truth labels (i.e., binary privacy decision) as we described in Chapter 6.3.2 and 6.3.3. We then divided the preprocessed training/test sets by the three

levels of decision confidence (see Table 6.12). Since there were very few *unconfident* and *very unconfident* decisions in both datasets, we merged them with *neutral* decisions into a *neutral and unconfident* category. As a result, we wound up with 152 (23.6%) *neutral and unconfident*, 319 (49.5%) *confident*, and 174 (27%) *very confident* decisions in our test data[3].

Confidence Level	# Instances	# Accept Decisions	# Reject Decisions	Accept Ratio
Neutral and Unconfident	***152***	***71***	***81***	***46.71%***
Very Unconfident (1)	4	0	4	0.00%
Unconfident (2)	32	14	18	43.75%
Neutral (3)	116	57	59	49.14%
Confident (4)	**319**	**172**	**147**	**53.92%**
Very Confident (5)	**174**	**41**	**133**	**23.56%**
Total	645	284	361	44.03%

Table 6.12: Data Size and Class Distribution (UCI)

Like the previous work, we used area under the ROC curve (AUC) as a performance measure to overcome the class imbalance problem. To be specific, test data is composed of about 24% of *very confident* accept (use) and 76% of *very confident* reject (non-use) decisions. Therefore we may have misleading results if we computed classification accuracy, which is the number of correct predictions divided by the total number of predictions made. AUC is unaffected by this issue. We used `scikit-learn` machine learning library for all of our experiments.

6.5.3 Experiment Results

6.5.3.1 Decision Confidence and Predictive Performance

To begin with, we trained ML models for all possible combinations of training and test data, each grouped by the three levels of decision confidence. We mainly intended to investigate how privacy decision confidence impacts the ML models' predictive performance. As men-

[3]Previously collected training data contains relatively less *neutral and unconfident* and more *very confident* decisions compared to test data; there were 1,402 (19.2%) *neutral and unconfident*, 3,254 (44.5%) *confident*, and 2,664 (36.4%) *very confident* decisions.

tioned above, we used training and test data sampled from two different user populations, namely MTurkers and students, respectively. In this experiment, we constructed Random Forests (RF) models to compare the prediction results with our previous experiments.

Train MTurkers\Test Students	NU[a] (152)	Confident (319)	Very Confident (174)	*All (645)*
Neutral and Unconfident (1402)	**57.27%**	66.28%	73.63%	*66.72%*
Confident (3254)	52.75%	**71.21%**	79.05%	*69%*
Very Confident (2664)	51.77%	69.2%	**87.24%**	*71.97%*
All (7320)	*48.3%*	*69%*	*81.44%*	*69.24%*

[a]Neutral and Unconfident

Table 6.13: Predictive Performance – Confidence Matrix (Random Forests)

Table 6.13 displays the predictive performance of the RF models trained on MTurkers' privacy decision behaviors separated by its underlying decision confidence. Numbers in parentheses indicate the sample size of the corresponding privacy decisions. As can be seen, the performance of the trained RF models against both *very confident* and *all* test instances gradually improves as the level of decision confidence of training instances increases (see the last two columns in Table 6.13). Interestingly, we observed the opposite trend toward *neutral and unconfident* test instances; the predictive performance tends to degrade as training instances' decision confidence increases. In general, we recognized that we could obtain the optimal results if we used training and test data with the same level of decision confidence (see the diagonal of Table 6.13).

However, privacy decision prediction mechanism should perform well on all unseen decision-making instances (decision confidence is not available when making predictions). In this regard, we believe that the results against all test instances were the most accurate metric for evaluating the trained ML model. We achieved the best AUC of 71.97% by the model trained with MTurkers' *very confident* decision behaviors and an overall AUC of 69.24%. However, both of them are worse than the result we showed in our previous study (mean AUC of 76.41% through 10-fold CV). We suspect that two sample populations have different characteristics (i.e., professional survey takers vs. STEM students), leading to inconsistent

privacy behavioral patterns. It might be the reason why the RF models trained on a specific dataset showed unimpressive predictive performance on a totally different dataset [11].

6.5.3.2 Decision Confidence and ML Algorithms

In order to assure the abovementioned finding was not limited to a specific algorithm investigated (Random Forests), we repeated the experiment using seven different ML algorithms widely used in recommender system research (including privacy decision support, e.g., [36, 87, 16]). These include Extremely Randomized Trees (ERT), Gradient Boosting (GB), Support Vector Machines (SVM) with RBF kernel, Logistic Regression (LR), k-Nearest Neighbors (kNN), Decision Tree (DT), and Naïve Bayes (NB). Like the experiment described in Chapter 6.5.3.1, we used training data extracted from MTurkers' privacy decision-making behaviors (grouped by decision confidence) for building different ML models. Regarding test data, however, we used all privacy decisions of university students to gauge the trained ML models' predictive performance in a real-world situation.

Confidence\Algorithm	RF	ERT	GB	SVM	LR	kNN	DT	NB
Neutral and Unconfident	66.72%	62.91%	68.43%	68.26%	65.2%	60.72%	54.94%	64.98%
Confident	69%	66.4%	**71.63%**	70.56%	64.47%	61.58%	**60.56%**	64.95%
Very Confident	**71.97%**	**68.79%**	68.5%	70.94%	**67.29%**	**65.99%**	60.53%	**65.89%**
All	69.24%	66.07%	70.67%	**71.91%**	66.35%	65.62%	59.57%	65.58%

Table 6.14: Predictive Performance – Algorithm Benchmark
Note: We used the default hyper-parameters for all ML algorithms. We configured the same seed of the pseudo-random number generator used when shuffling training data to achieve the reproducibility of our experiments.

Table 6.14 summarizes the experiment results. As can be seen, training instances of *confident* or *very confident* privacy decisions yielded the highest prediction accuracy for most of the ML algorithms we tested, except for SVM. We noticed that the SVM model trained on all available training instances slightly outperformed the model based on *very confident* privacy decisions only. We suspect that SVM could perform well even without the consideration of decision confidence because it inherently has a good generalization ability [99]. From

all training instances, SVM probably learned the latent patterns of privacy decision-making with different levels of decision confidence and then produced binary classification results accordingly. Regarding GB and DT, we noticed that the models based on *confident* decisions slightly outperformed *very confident* models. As there exists more *confident* decisions (3,254 instances) than *very confident* decisions (2,664 instances) in our training data, the resulting ML models might be able to recognize more diverse patterns of privacy decision-making behaviors. Other than that, we confirmed that the ML models performed the best when they were trained on *very confident* decisions across all types of algorithms.

Through the experimentation mentioned above, we showed the fact that the level of decision confidence of training instances is related to the predictive performance of ML models. To be specific, it is possible to build more accurate predictive models if we utilized privacy decisions made with confidence as training data. Unlike our previous work, we validated this claim using training and test data separately collected from two different groups of IoT users. At the same time, however, we obtained less accurate overall predictive performance (about 7% lower AUC) than our previous result based on a single user population. We consider that training and test users showed somewhat different privacy behaviors, thereby making the learned behavioral patterns inconsistent between two user groups.

6.6 Discussion

We have investigated the causes and effects of confident privacy decision-making in IoT, through qualitative analysis and machine learning experiments on the collected survey responses. In this section, we present some issues that need to be considered and addressed.

First, our study participants were skewed toward students who aged 25-34 (60.5%). Also, most of them had a science and engineering background. This may result in a sampling bias

that makes our findings less general. For example, about 70% of the participants mentioned that privacy awareness questions were easy to answer. As discussed, the collected responses to these questions (i.e., privacy awareness measurement) had an impact on not only the type but confidence level of their privacy decisions toward IoT services. However, it would be the opposite for other populations (e.g., old people who are not familiar with IoT-related technologies). In this regard, it is necessary to validate our findings with more representative samples.

We also identified the necessity of obtaining the generalizability of the ML model for the realization of privacy decision support for a broader range of users. ML models trained on privacy decision samples gathered from a specific user group (MTurkers) did not perform well on unseen test instances given by a different group of users (students). To tackle this difficulty, we may need to investigate ways to secure a greater amount of training data to better accommodate diverse privacy decision-making behaviors of IoT users. One possible approach is to augment the preexisting training data (i.e., data augmentation) by generating additional synthetic decision-making instances through over-sampling methods such as SMOTE [21]. However, we should be careful about over-fitting problems during the model training procedure since we will not feed entirely new data instances into the model.

Lastly, we analyzed people's privacy decisions about hypothetical IoT scenarios and not actual behavior toward working IoT services. Although we tried to make participants believe they were in a real situation, we still do not know how they would actually behave in their everyday lives. Previous research repeatedly indicated that user-stated privacy decisions are often inconsistent with actual behaviors (i.e., privacy paradox [3, 45, 30]). Thus, we need to let study subjects freely interact with an operational IoT environment such as TIPPERS [81], possibly including multiple IoT services, and collect their actual privacy behaviors accordingly.

6.7 Conclusion

In this chapter, we conducted an online survey study to understand the user perceptions and practical implications of confident privacy decision-making in IoT environments. We asked and analyzed people's opinions about the rationale for their privacy decisions toward the presented IoT service scenarios. The analysis results indicated that people who are well aware of potential privacy risks (i.e., high privacy awareness) tend to have more confidence in their decisions of whether to use the service or not. We also performed machine learning experiments to confirm the fact that privacy decision instances made with high confidence would enable us to construct more accurate privacy prediction models. In doing this, we not only used training and test data collected from two distinct user groups but also tested different types of machine learning algorithms, to prove the generalizability of our finding. As future work, we plan to focus on the collection and analysis of privacy decision-making behaviors of a wide variety of users who interact with working IoT systems.

Chapter 7

Privacy-Aware System for Privacy-Preserving IoT Environments

IoT services collect and analyze sensor data to provide users with intelligent functionality tailored to their needs. However, users are often unaware of privacy risks relating to sensor data collection and the inferences possible from this data. Even if aware of the data collection and possible inferences, users lack ways to manage the collection, processing, and transmission of the data. To address this problem, we designed and implemented a novel web-based privacy awareness system called `IoT Service Store` (`ISS`) that allows users to easily browse nearby IoT services, understand the privacy implications of these IoT services, and control the collection and usage of sensor data. To better inform users about the potential privacy risks in using IoT services, `ISS` displays detailed information on what personal information might be inferred from the sensor data being collected. `ISS` also allows each user to give a rating or to view other users' ratings regarding the perceived utility-privacy tradeoff for each IoT service. `ISS` is designed to communicate with IoT services to modify those services' data collection and usage practices, according to a user's privacy preferences. Using the preferred privacy settings in the proposed system, users will be more *confident* in

their decisions of whether to subscribe to IoT services and less concerned with privacy risks in using the services.

7.1 Introduction

The ubiquity and density of the Internet of Things (IoT) are rapidly increasing. In the near future, we will be surrounded by numerous sensor devices that unobtrusively and collaboratively extract myriad types of sensor data from the user's environment. Much of this data will be information involving people's presence, behaviors, or states. For instance, an automated student attendance system might infer the presence of a particular user based on the Wi-Fi MAC address of the user's device and/or video collected by cameras associated with the system. Rapid advancements in big data processing and analysis will additionally allow IoT service providers to infer more diverse user-related information from this sensor data. The inferred information may include sensitive personal information of users, such as their attendance patterns, emotions, health condition, or even sexual orientation. Our proposed system is predicated on the notion that these and other types of personal information should only be collected and processed with a user's meaningful consent.

A more fundamental issue is that users generally do not know about the existence of nearby sensors (except possibly for sensors in the user's residence), let alone the nature of services operating these sensors. Also, it is hard to make sensor devices transparent to users (e.g., by adopting a conventional notice-and-consent model) due to the lack of natural communication channels between the users and services. Therefore, users need to have a unified way to discover nearby IoT services and understand the privacy properties of these services. In this way, they can become aware of the privacy implications of using IoT services of interest [24]. To be specific, users not only need to be informed about the whole process of sensor data collection, but also to understand what types of personal information might be inferred from

the collected sensor data. With this information, users will be better positioned to define their own privacy settings (i.e., preferences) more *confidently* and therefore possibly give IoT services a way to respect these preferences [40, 25].

With these aims in mind, we designed and implemented a web-based system called `IoT Service Store` (`ISS`) for privacy-aware IoT service discovery and interaction. `ISS` is a web server that manages various privacy-related information of multiple IoT services. Here is how it works: First, each IoT service registered to `ISS` broadcasts its unique identifier (Uniform Resource Locator; URL) through Bluetooth beacon(s). The beacon-generated URLs are automatically detectable by nearby users' Bluetooth-enabled Android or iOS smartphones. When a user clicks on a specific URL, she will be redirected to the designated web page, hosted by `ISS`. This web page contains a visualization of the IoT service's privacy-related information, including its data collection policies, inferable personal information, and users' collaborative evaluations of the service's utility-privacy tradeoff. Most notably, we developed a novel information architecture of both sensor data and personal information, as well as their relationships (e.g., sensor data `A` implies personal information `B`), in order to better inform users of the privacy properties of the service. Also, we adopted a five-star rating system for letting users collaboratively evaluate the service in terms of the balance between utility benefits and privacy risks.

`ISS` is designed to allow IoT services to comply with user-defined privacy settings. For instance, a user may utilize the automated student attendance system while allowing his/her Wi-Fi MAC address to be gathered but at the same time disallowing the collection of face photos. We plan to integrate `ISS` with several IoT services running on an operational IoT framework called TIPPERS [71, 81]. Through TIPPERS's open data APIs, we can programmatically make *opt-out* requests for specific sensor data to be collectible by the IoT service. Using the proposed system, users can subscribe to specific IoT services that offer and honor their preferred privacy settings, thereby helping minimize user-perceived inappropriate in-

formation flows.

In summary, our work makes the following contributions to the field of usable privacy and pervasive computing:

- We designed and implemented a novel web-based IoT service management system called `IoT Service Store` for privacy-aware service discovery and interaction.
- We developed an information architecture describing the IoT service's privacy properties, specifically the types/attributes of raw sensor data and the personal information that may be inferred. Our architecture clearly separates the data from the inferences, motivated by the ever-increasing capabilities of machine learning.
- We designed interfaces to allow users to inspect the privacy properties of the IoT service, to collaboratively evaluate its potential privacy risks, and to accordingly adjust data collection practices.

7.2 Related Work

One of the first privacy awareness systems for ubiquitous computing environments was proposed by Langheinrich [56]. The author proposed a system called pawS. This system aimed not only to allow data collectors to announce data usage policies, but also to provide data subjects (i.e., users) with the technical means of managing how their personal information is stored and processed by the service. The author assumes that all entities in an environment have their own privacy proxies, continuously running services that handle privacy-related interactions between the entities. Each user has his/her own personal privacy proxy which contains privacy preferences with respect to the multiple service privacy proxies that codify data collection and usage processes. Service providers describe their data collection policies

using a machine-readable XML format such as a P3P privacy policy. Using such a policy, each service provider can describe, for example, who is collecting data, what data is being collected, and for what purpose, in each case. Correspondingly, end users can express their own privacy preferences via a machine-readable preference language such as APPEL[1], which consists of a set of (updatable) rules. On pawS, all data collection and usage are therefore performed in accordance with the user's privacy preferences.

Even though privacy awareness systems like pawS provide standardized ways to safeguard user privacy, service providers are still required to change their service infrastructure and/or reveal internal data handling practices, which could be a significant barrier to adoption. To handle this issue, Kolter et al. prototyped a user-centric privacy architecture that enables provider-independent privacy awareness in using Internet services [53]. The core part of this architecture is an online privacy community that lets multiple users post and share privacy-related information regarding the service. Just like in Wikipedia, users can edit diverse information about a specific web service (e.g., Amazon.com), including required amounts of personal data, practices of data sharing with third parties, adherence to the stated privacy policies, and the subject evaluation of privacy risks. Users can also share their personal privacy preferences with others. Inexperienced users may import the pre-defined preferences of a trusted privacy expert and utilize the imported preferences as their baseline choices.

Recently, researchers are realizing privacy awareness in a real world IoT environment. Mehrotra et al. are developing a privacy-aware IoT framework named TIPPERS and deploying it in Donald Bren Hall (DBH) at the University of California, Irvine (UCI) [71, 81]. In order to transform DBH into a smart environment, TIPPERS captures raw data from various sensors installed in DBH and makes the collected data publicly accessible through open data APIs. Third-party developers are then able to create various IoT services (e.g., indoor location awareness apps) on the TIPPERS framework. Regarding user privacy, service providers (or

[1] `https://www.w3.org/TR/P3P-preferences/`

building administrators) need to advertise building policies which give detailed information about data collection procedures regarding the service. Then, end users set their privacy preferences and ask TIPPERS to enforce these preferences while operating the service. Both building policies and privacy preferences are defined using a custom JSON-schema for supporting data request and access from devices outside TIPPERS (e.g., user's smartphone). The research team is currently developing a remote storage called IoT Resource Registries (IRRs) for administering building policies. They are also developing a smartphone app called IoT Assistants (IoTA) that notifies users about available building policies, thereby configuring privacy preferences, whether via interactions with the users or automatically.

The common goal of this body of work is making personal data collection as transparent as possible to users, thereby helping them make an informed privacy decision. However, it is still unclear how well users understand the implications of some data collection. For instance, not all users understand that the Wi-Fi MAC address is an identifier that may be used to track the location of the corresponding user device or its owner. Service providers can explicitly describe the implications of various sorts of sensor data collection; however, they may not cover all inferences of personal information based on the sensor data. For instance, it's obvious that images from a camera might reveal users' identity, but it's less well-known that these images may also reveal their sexual orientation [100]. The IoT service may not rely on these other inferences and in fact the service provider may be oblivious to them. Nevertheless, for the sake of user privacy, the user should know about these inferences. Thus, our aim is the development not only of an information architecture for describing the privacy properties of IoT services, but also of user interfaces (including the underlying system) that efficiently convey possible inferences of data collected. To the best of our knowledge, our platform is the first to address user understanding of inferences and to define an architecture that distinguishes inferences from raw sensor data collection. In addition, this body of work has not considered collaborative privacy management strategy (e.g., crowdsourced evaluations of privacy risks) in improving users' privacy awareness in

IoT. Since Kolter et al. showed its feasibility in web environments [53], we are also applying this approach in our work.

7.3 Web-Based Privacy Awareness System for IoT

In this section, we discuss the design and implementation of our system in detail. We first describe an overall system architecture including software/hardware specifications, and then explain the functional details of the system.

7.3.1 System Architecture

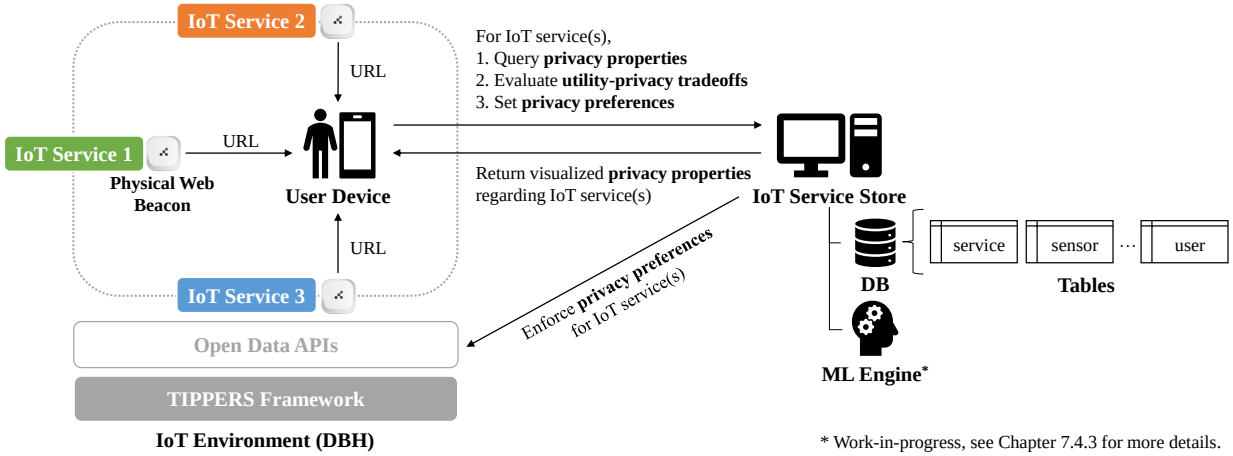

Figure 7.1: `IoT Service Store` — System Architecture

Our system is designed as a client-server model (Figure 7.1). Users can query diverse privacy-related information (e.g., privacy properties) of IoT services of interest through web browsers installed on their smartphones. The web server (`IoT Service Store`; `ISS`) then retrieves requested information from the database (DB) and returns the results back to the users. `ISS` is regarded as a trustworthy entity. Users can also send requests for other types of operations (e.g., giving a rating or setting up privacy preferences for the target service) to the server.

With the intent of making the system as easily accessible as possible, we followed standard web protocols in the implementation of functionalities for communication and interactions between the entities under the system. In addition, all information sent between the user and `ISS` is secured through use of the HTTPS protocol.

In order to let IoT service providers uniformly inform users about service descriptions along with privacy implications, we adopted the Physical Web[2] as an underlying communication mechanism. Physical Web is an open source project that aims to transform all physical objects (e.g., parking meters) into smart agents by allowing the objects to interact with the web via Bluetooth Low Energy beacon profile called Eddystone. Eddystone beacons are capable of broadcasting object-specific identifiers like URLs, which are automatically searchable by nearby users' Bluetooth-enabled Android or iOS smartphones. Using these URLs, users are then able to browse web pages containing relevant information about physical objects (e.g., parking rates), and also to perform additional actions through their smartphones (e.g., payment of parking fees). We chose Physical Web since it enables users to easily find and interact with resources in their physical environments, without first downloading an additional app. In the current implementation, `ISS` assigns a unique URL to each of the registered IoT services. Service providers need to deploy Eddystone beacon(s) broadcasting the assigned URLs to advertise their services.

`ISS` is a standalone web server running on a virtual private cloud. We installed a standard web service stack composed of Linux, Apache, MySQL, and PHP (LAMP) on an Amazon Elastic Compute Cloud (EC2) instance and deployed a server program on this virtual machine. We assume that each IoT service provider registers its service(s) to `ISS`, while providing the detailed data collection practices in an honest manner. Based on this information (stored in a MySQL database), `ISS` systematically composes a specific web page with user interfaces displaying privacy-related information of each IoT service. After that, `ISS` assigns

[2] `http://google.github.io/physical-web/`

a unique URL for the service. We designed and implemented all user interfaces for `ISS` using HTML5 and CSS. We also added some JavaScript functions (e.g., jQuery for changing the content of a web page without reloading) to make the user interfaces more interactive and responsive for the users. We will further elaborate on ways to formulate the web page in the following sections. As discussed above, `ISS` is also designed to communicate with the TIPPERS IoT framework so as to ask IoT services to follow user-defined privacy settings (see the lower left of Figure 7.1).

7.3.2 Workflow

We now explain the functional workflow of the proposed system (see Figure 7.2).

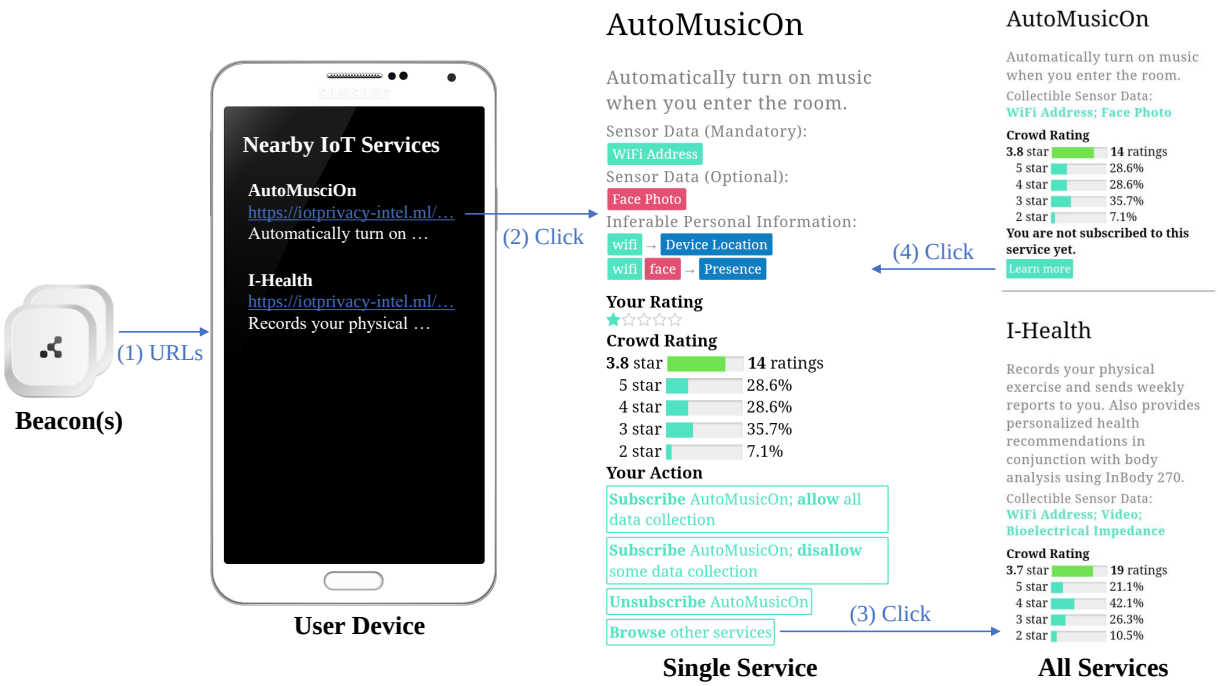

Figure 7.2: `IoT Service Store` — Functional Workflow

Step 1 As discussed, IoT service providers broadcast their service URLs through Physical Web-compatible beacons. These URLs, mapped with web pages summarizing privacy-related information of the IoT services, are automatically detectable by nearby users'

Bluetooth-enabled smartphones. Android users might become aware of available services (URLs) via OS-level notifications (i.e., Google Nearby); iOS users, in contrast, will get notifications from Chrome browsers installed on their devices.

Step 2 Users click on URLs of the IoT services of interest and then check various privacy-related information on these web pages being hosted by `ISS`. Each web page first visualizes all possible combinations of sensor data as well as personal information which can be inferred from each combination of data (i.e., privacy properties). Next, it provides user interfaces for both receiving an individual user's evaluation of utility-privacy tradeoff and showing all other users' evaluations with statistics. The last part of the web page is a list of all available actions for the user (e.g., disallowing some data collection).

Step 3 Users are also able to browse other IoT services running at their current location by clicking a `Browse other services` button. When this happens, users will be redirected to a different page enumerating nearby IoT services with short descriptions. In this page, IoT services are sorted in descending order by the users' ratings about the utility-privacy tradeoff.

Step 4 If users find any other interesting services, they can check the details by clicking a `Learn more` button.

7.3.3 Privacy Properties of IoT Services

To make users more aware of privacy risks related to IoT services they are using, we defined an information architecture capable of expressing privacy properties of the services. As discussed earlier, users have a limited understanding of the implications about the collection of various kinds of sensor data. The primary goal of the proposed system is reinforcing users' privacy awareness by letting them understand each IoT service's privacy properties,

namely relationships between raw sensor data (e.g., Wi-Fi MAC address) and higher-level descriptions of personal information (e.g., user identity) which can be inferred from the sensor data.

7.3.3.1 Sensor Data and Personal Information

To begin with, we need to build a taxonomy of sensor data collectible in IoT environments. Since we plan to integrate the proposed system with the TIPPERS IoT framework, currently running on a six-story building (DBH) at UCI, we first defined the types of sensor data according to available sensor devices installed in DBH (see Sensor Data in Table 7.1). As explained, IoT service providers need to register their services with detailed data collection policies. In order to systematically express the policies for collecting sensor data, we also defined the following attributes: `service_id` (identifier of service collecting and processing data), `mandatory` (indication that data is mandatory or optional for using the service), `source` (sensor device collecting data), `storage` (location where data is stored), `retention` (time duration for which data is stored), `protection` (security mechanism for protecting data), and `sharing` (the existence of third-parties allowed to access data). As all this information is stored in a MySQL database, service providers (or system administrators) can update their data collection practices via simple web user interfaces.

Next, we developed a taxonomy of (inferable) personal information in the context of IoT. Since we were unable to find a classification scheme for personal information in IoT, we created a broad-brush classification scheme using the P3P specification V1.1[3] as a baseline. To do this, we considered the `CATEGORIES` element that describes 16 different types of personal information available in web environments. This element was originally designed to help Internet users define generalized preferences and rules (i.e., P3P privacy policy) for the exchange of their personal data through the web. We augmented this element to consider

[3] `https://www.w3.org/TR/P3P11/`

Sensor Data	(Inferable) Personal Information	
(1) Image/Video	(1) *Identity*	(9) *Preference*
(2) Audio	(2) *Purchase*	(10) *Presence (Location)*
(3) Wi-Fi	(3) *Financial*	(11) ***Physical Activity***
(4) Temperature	(4) *Device*	(12) ***Physical State***
(5) HVAC	(5) *Behavioral*	(13) ***Emotion***
(6) Electricity	(6) *Demographic*	(14) ***Personality***
(7) Light	(7) *Political*	(15) ***Cognitive Activity***
(8) Motion	(8) *Health*	(16) ***Social Relationship***

Table 7.1: Sensor Data and Personal Information in IoT

personal data available in physical environments, as opposed to web environments. For instance, IoT devices are argued to accurately recognize users' emotional states by analyzing video footage captured by a security camera [47]. We added the following types of personal information (bold-faced items in Table 7.1): *Physical Activity* (e.g., walking or running), *Physical State* (e.g., sitting position), *Emotion* (e.g., Ekman's six basic emotions), *Personality* (e.g., The Big Five personality traits), *Cognitive Activity* (e.g., intention), and *Social Relationship* (e.g., workplace dynamics). We also excluded some personal information which is not directly related to IoT (e.g., user-generated online content). As a result, we wound up with 16 types of personal information. Note that each type of sensor data or inferable personal information may have subcategories (e.g., *Image>Face Photo*) to better express its meaning.

7.3.3.2 Inference of Personal Information

Providing users with knowledge of the possible inferences of sensitive personal information is important for increasing their privacy awareness in IoT environments. The most straightforward way would be for IoT service providers to generally describe the inferred personal information from the collected sensor data. This declaration may, however, reveal confidential business strategies [53]. Also, each service provider may not know other possible inferences that provide no utility to the service it offers. For instance, a service provider who

is utilizing facial recognition software for the purpose of user authentication (*Identity*) might be unaware of the fact that a similar technique can be used to infer users' sexual orientation (*Preference>Sexuality*) [100] from previously collected image data. For these reasons, we believe that the construction of a knowledge base about the inference of personal information is necessary for reinforcing awareness of the privacy risks in IoT. As a starting point, we therefore defined *if-then* rules for specifying the privacy properties of the IoT services, composed of the combination of available sensor data (antecedent) and possible inferences of personal information (consequent). In doing this, we referred to literature on mobile sensing and data mining related to personal information of the user [66, 63, 6]. We will also discuss strategies for extending and managing this knowledge base in a later section.

We generated 36 different rules for five hypothetical IoT services registered to the proposed system. As an example, `AutoMusicOn` is a service that aims to automatically play music upon the user's entrance into a specific room. To check whether a registered user enters the room, `AutoMusicOn` is collecting Wi-Fi MAC addresses of mobile devices and/or face photos of people inside the room. In this case, the Wi-Fi MAC address will need to be collected because it can imply both the current location of nearby devices (see *Device>Location* in Figure 7.3) and the identity of their owners (*Identity*), thereby possibly inferring user location. Optionally, the service may also collect face photos in order to verify a user's identity via facial recognition and finally confirm the presence of this user in the room (*Presence*). With the collected facial images, however, the service provider (or third-party) might infer the sexual orientation of the user (*Preference>Sexuality*) in the future. We manage these rules in a MySQL database under the system for commonly applying them to all the registered services. Therefore, each IoT service's web page summarizes its privacy properties as depicted in Figure 7.3. As can be seen, (mandatory/optional) sensor data and inferable personal information are distinguished through color coding. In addition, all items are visualized as *clickable* buttons; users are then able to click the button to view the popup window giving additional explanations about the item.

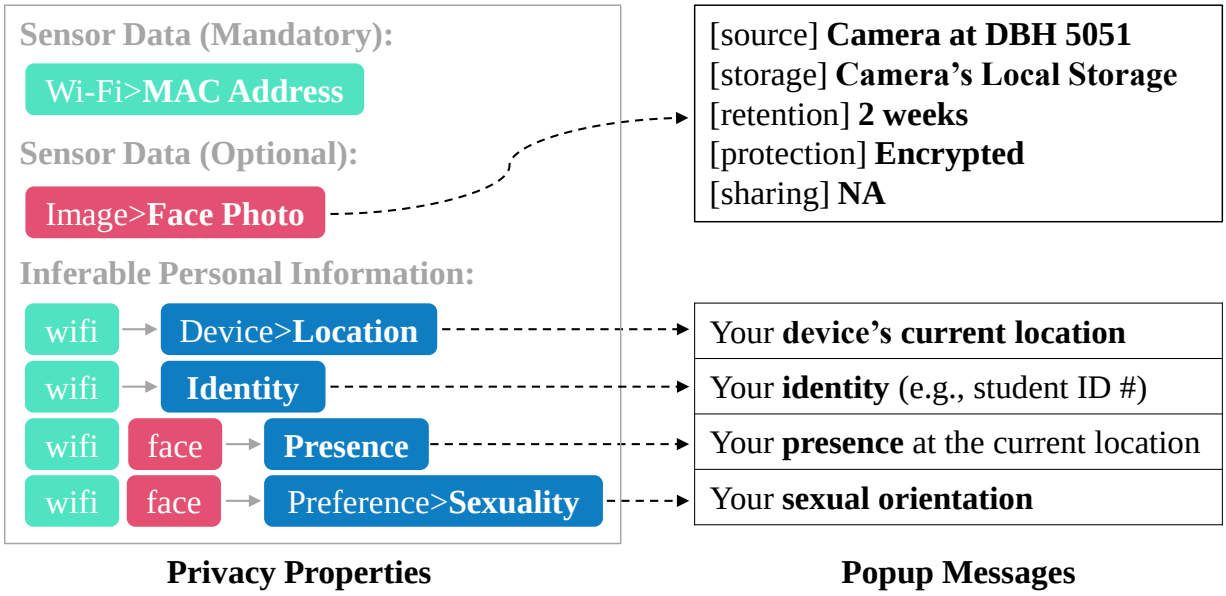

Figure 7.3: `IoT Service Store` — User Interfaces

7.3.4 User-Driven Assessment and Control of Privacy Risks

Aside from providing the privacy properties, `ISS` lets users evaluate and share their opinions about nearby IoT services. We adopted a conventional five-star rating system to allow users to collaboratively assess the subjective balance between the benefits and privacy risks of using a specific IoT service. After checking privacy properties of the service, users can leave their ratings (ranging from 1 to 5 stars), just as they do in general e-commerce systems like Amazon.com (see `Your Rating` in Figure 7.2). Since the system distinguishes users with server-generated identifiers (i.e., random number assigned to `$_SESSION` variable in PHP), users are able to change previously submitted ratings whenever they want. For instance, a user who gave 1 star to `AutoMusicOn` solely based on its stated privacy properties might update his/her rating after actually using the service. Except for the session identifier, the system does not collect or use any personal data (e.g., phone number) to recognize the user at this moment. To present multiple users' evaluations about the service at a glance, `ISS` also calculates its average rating and displays it with the distributions of star ratings (see `Crowd Rating` in Figure 7.2).

Lastly, each service web page provides users with the following control options: (1) subscribe service, allow all data collection, (2) subscribe service, disallow some data collection, (3) unsubscribe service, and (4) browse other services. Users can make choices by clicking buttons in the web page (see `Your Action` in Figure 7.2). Regarding option (2), `ISS` is designed to send a message requesting the *opt-out* of specific data collection (i.e., user's privacy preferences) in using the service. It can be simply done through JSON-based REST APIs and a policy enforcement engine, both of which are being provided by the TIPPERS framework [81]. Using the abovementioned session identifiers, the system keeps track of all actions performed by the users, thereby allowing them to change their preferences if necessary.

7.4 Discussion and Future Work

In this section, we explore future opportunities for research that will potentially help further improve the effectiveness of the proposed system, in making people aware of the privacy implications regarding the IoT service and available options for them to avoid potential privacy breaches.

7.4.1 Scalable Knowledge Base for Privacy Properties

We consulted existing literature to define a set of rules specifying possible inferences that can be drawn from sensor data available in IoT, and then used these rules to present the privacy properties of the IoT services. However, we recognize that in practice, claimed inferences are complex, highly context-dependent, and open to misinterpretation. This approach will necessarily have scalability issues as long as humans are a component of the system. One possibility, albeit not ideal, is to build a probabilistic information retrieval system that auto-

matically extracts relationships between entities (e.g., sensor data and personal information) in text data (e.g., machine learning literature). To build such a system, however, we still need labeled training data (i.e., known relationships between sensor data and personal information), which does not exist to our knowledge. One approach is to utilize systems such as Snorkel, designed for programmatically generating labeled training datasets from raw data without much human intervention [84].

7.4.2 Privacy Decision-Making in Operational IoT Environments

Through our previous survey studies, we showed that people's privacy decisions are significantly affected by the awareness of inferable personal information in IoT environments (see Chapter 3 and 5.4). Regarding the collection of image data, for instance, users are more worried when they realize the implications of image-based human age estimation. In contrast, they are very open to providing information about their devices (e.g., phone identifier) if they perceive that this information is not related to their sensitive personal information. However, these findings are based on people's stated privacy decision-making towards hypothetical service scenarios, not actual behavior of using working IoT systems. It is therefore necessary to analyze users' privacy behavior captured in real world situations, both for validating our previous findings and for extracting additional insights about the privacy awareness in IoT. In this vein, we plan a collaboration with the TIPPERS research team in order to incorporate their internally (or externally) developed service apps into the proposed privacy-aware system. Thereafter, we will deploy the integrated system to DBH and conduct field experiments with real users (i.e., building inhabitants) to collect and analyze their privacy-related behavioral data generated while using the system.

7.4.3 Ubiquitous Privacy Decision Support

Even though the proposed system presents users with information that will help them understand some privacy implications of using various IoT services, users still need to configure their privacy settings by themselves. As discussed before, however, some users may have difficulties in doing so due to limits in their available time, motivation, and cognitive decision-making abilities [2, 92]. Therefore, the system may need to assist users with making better privacy choices, perhaps by predicting future decisions based on their historical decision-making behavior and recommending privacy settings accordingly. We are currently considering using our machine learning (ML) methodologies (see Chapter 4 and 5.5) to realize this functionality. By using the abovementioned privacy behavioral data (i.e., users' interaction logs of using the system and their submitted privacy settings) as training data, we can train ML model(s) predicting privacy decisions of the users. We expect that the collected training data contains a sufficient amount of *confident* privacy decision-making instances since users will be given privacy properties of IoT services from the proposed privacy-aware system. It would yield a reasonable predictive performance (see Chapter 5.5.2 and 6.5.3). We then embed the trained ML models(s) into the proposed system (see `ML Engine` in Figure 7.1) for providing users with machine-generated privacy recommendations.

7.5 Conclusion

In this chapter, we designed and implemented a novel web-based system called `IoT Service Store` (`ISS`). Our goal is to allow users to comprehend the privacy implications of nearby IoT services through gaining a better understanding of the data collected and possible inferences that may be drawn from this data. `ISS` also allows users the ability to control the collection of their data. For concreteness, we adopted the Physical Web as an underlying communication channel between the users and IoT services, thereby realizing an easy dis-

covery of IoT services operating near the user. In order to efficiently express and convey the detailed privacy properties of each available IoT service to the user, we developed an information architecture for describing the relationships between collected sensor data and inferable personal information. We designed and implemented privacy-aware user interfaces, as a presentation layer of `ISS`, not only visualizing the privacy properties of an IoT service, but also allowing users to collaboratively assess its potential privacy risks and configure privacy settings according to their privacy expectations. Future work will mainly focus on the following topics: the automated extension of a knowledge base used for presenting IoT services' privacy properties, collection and analysis of people's perceptions of privacy in operational IoT environments, and a ML-based privacy decision support system that alleviates users' cognitive burden of configuring privacy settings for diverse IoT services.

Chapter 8

Conclusion

In this dissertation, we first aimed to extract diverse contextual factors comprising IoT service scenarios and understand how these factors impact people's privacy perceptions and decisions toward IoT. To this end, we conducted both online and situated survey studies with two different user populations (Amazon MTurkers; $N = 200$ and university students; $N = 172$, respectively) and analyzed the collected survey responses using a K-modes clustering algorithm. We identified four clusters of scenarios ($K = 4$), with clearly distinctive associated user reactions. By comparing the different clusters, we identified contextual parameters (namely, parameter *who* and *what*) that are strongly associated with higher or lower acceptance of personal information collection/inference caused by IoT environments. We also hypothesized that there exists an important user-specific factor (namely, *privacy awareness*) that could impact the type and quality of privacy decisions that each individual makes. To validate this claim, we performed a new online survey (Amazon MTurkers; $N = 488$) and analyzed the collected responses using a mixed-effect logistic regression model. Through careful model selection and fitting procedures, we confirmed the fact that the degree of user's privacy awareness is positively correlated with the likelihood of making conservative and *confident* privacy decisions. Machine learning experiments with a Random Forests algo-

rithm also verified that confident decision samples (i.e., training data) yield more accurate privacy decision prediction models.

In order to better implement a privacy decision support system, we also proposed a novel machine learning (ML) mechanism for predicting people's privacy decisions in IoT environments. We aimed to predict binary privacy decisions of each user, namely whether to allow or reject a given personal information monitoring scenario in IoT. To begin with, we adopted linear model and deep neural networks (LMDNN) as the ML algorithm for our study. Using a privacy behavioral dataset collected from our situated survey study (university students; $N = 172$), we confirmed that LMDNN provides better predictive performance than conventional ML algorithms that have been widely used in the literature. Next, we utilized both contextual and user-specific information as input features for training LMDNN models. We adopted a wide range of contextual factors comprising diverse IoT scenarios. We then generated *privacy segment* information of our study participants by clustering their responses toward a single selected IoT scenario. Lastly, we proposed a new model training strategy called *one-size-fits-segment* modeling, and compared its performance with two commonly used strategies: individual and one-size-fits-all modeling. Experimental results indicated that *one-size-fits-segment* outperforms other modeling strategies.

Lastly, we designed and implemented a novel web-based privacy-aware system called `IoT Service Store` (`ISS`). The ultimate aim of `ISS` is to allow users to comprehend the privacy implications of nearby IoT services through gaining a better understanding of the data collected and possible inferences that may be drawn from this data (i.e., enhanced *privacy awareness*). For concreteness, we adopted the Physical Web as an underlying communication channel between the users and IoT services, thereby realizing an easy discovery of IoT services operating near the user. In order to efficiently express and convey the detailed privacy properties of each available IoT service to the user, we developed an information architecture for describing the relationships between collected sensor data and inferable personal infor-

mation. We designed and implemented privacy-aware user interfaces, as a presentation layer of `ISS`, not only visualizing the privacy properties of an IoT service, but also allowing users to collaboratively assess its potential privacy risks and configure privacy settings according to their privacy expectations. Thereby, users are more likely to make informed and hence confident privacy decisions. We believe that confident privacy decision-making is important since it yields more reliable and consistent behavioral patterns that can be efficiently learned by ML models for privacy decision support in IoT.

Future work will mainly focus on collecting privacy-related human behavioral data from more representative samples of users interacting with operational IoT systems and validating our findings against this new dataset.

Bibliography

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

[2] A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

[3] A. Acquisti and J. Grossklags. Privacy Attitudes and Privacy Behavior. In *Economics of Information Security*, pages 165–178. Springer, 2004.

[4] D. G. Altman. *Practical statistics for medical research*. CRC press, 1990.

[5] I. Altman. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Brooks/Cole Pub. Co., 1975.

[6] N. Apthorpe, D. Reisman, S. Sundaresan, A. Narayanan, and N. Feamster. Spying on the Smart Home: Privacy Attacks and Defenses on Encrypted IoT Traffic. *arXiv preprint arXiv:1708.05044*, 2017.

[7] 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*, 2(2):59, 2018.

[8] L. Atzori, A. Iera, and G. Morabito. The Internet of Things: A survey. *Computer Networks*, 54(15):2787–2805, 2010.

[9] P. Bahirat, Y. He, A. Menon, and B. Knijnenburg. A Data-Driven Approach to Developing IoT Privacy-Setting Interfaces. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces*, IUI '18, pages 165–176, New York, NY, USA, 2018. ACM.

[10] D. Bates, M. Mächler, B. Bolker, and S. Walker. Fitting Linear Mixed-Effects Models using lme4. *arXiv preprint arXiv:1406.5823*, 2014.

[11] J. Beel, C. Breitinger, S. Langer, A. Lommatzsch, and B. Gipp. Towards reproducibility in recommender-systems research. *User Modeling and User-Adapted Interaction*, 26(1):69–101, 2016.

[12] Y. Bengio, O. Delalleau, and C. Simard. Decision trees do not generalize to new variations. *Computational Intelligence*, 26(4):449–467, 2010.

[13] M. Benisch, P. G. Kelley, N. Sadeh, and L. F. Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, 15(7):679–694, 2011.

[14] M. Bergmann. Testing Privacy Awareness. In *IFIP Summer School on the Future of Identity in the Information Society*, pages 237–253. Springer, 2008.

[15] G. Bigwood, F. B. Abdesslem, and T. Henderson. Predicting location-sharing privacy preferences in social network applications. *Proc. of AwareCast*, 12:1–12, 2012.

[16] 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*, 25:125–142, 2016.

[17] I. Bilogrevic, K. Huguenin, B. Agir, M. Jadliwala, and J.-P. Hubaux. Adaptive information-sharing for privacy-aware mobile social networks. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 657–666. ACM, 2013.

[18] T. A. Brown. *Confirmatory Factor Analysis for Applied Research*. Guilford Publications, 2014.

[19] G. Burel, H. Saif, and H. Alani. Semantic Wide and Deep Learning for Detecting Crisis-Information Categories on Social Media. In *International Semantic Web Conference*, pages 138–155. Springer, 2017.

[20] A. Cetto, M. Netter, G. Pernul, C. Richthammer, M. Riesner, C. Roth, and J. Sänger. Friend Inspector: A Serious Game to Enhance Privacy Awareness in Social Networks. In *Proceedings of the 2nd International Workshop on Intelligent Digital Games for Empowerment and Inclusion (IDGEI)*, pages 1–8, 2014.

[21] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of artificial intelligence research*, 16:321–357, 2002.

[22] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, et al. Wide & deep learning for recommender systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, pages 7–10. ACM, 2016.

[23] 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 *Proceedings of the 13th international conference on Ubiquitous computing*, pages 41–44. ACM, 2011.

[24] R. Chow. IoT privacy: can we regain control? In *Proceedings of the 3rd ACM Workshop on Information Hiding and Multimedia Security*, pages 3–3. ACM, 2015.

[25] R. Chow. The Last Mile for IoT Privacy. *IEEE Security & Privacy*, 15(6):73–76, 2017.

[26] R. Chow, S. Egelman, R. Kannavara, H. Lee, S. Misra, and E. Wang. HCI in Business: A collaboration with academia in IoT privacy. In *International Conference on HCI in Business*, pages 679–687. Springer, 2015.

[27] D. Christin, A. Reinhardt, S. S. Kanhere, and M. Hollick. A survey on privacy in mobile participatory sensing applications. *Journal of systems and software*, 84(11):1928–1946, 2011.

[28] J. Cohen. Statistical power analysis for the behavioural sciences, 1988.

[29] F. T. Commission et al. Internet of Things: Privacy & security in a connected world. *Washington, DC: Federal Trade Commission*, 2015.

[30] K. Connelly, A. Khalil, and Y. Liu. Do I do what I say?: Observed versus stated privacy preferences. In *IFIP Conference on Human-Computer Interaction*, pages 620–623. Springer, 2007.

[31] G. Danezis. Inferring privacy policies for social networking services. In *Proceedings of the 2nd ACM workshop on Security and artificial intelligence*, pages 5–10. ACM, 2009.

[32] A. Das, M. Degeling, D. Smullen, and N. Sadeh. Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice. *IEEE Pervasive Computing*, 17(3):35–46, 2018.

[33] A. Das, M. Degeling, X. Wang, J. Wang, N. Sadeh, and M. Satyanarayanan. Assisting Users in a World Full of Cameras: A Privacy-aware Infrastructure for Computer Vision Applications. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 1387–1396. IEEE, 2017.

[34] A. Deuker. Addressing the Privacy Paradox by Expanded Privacy Awareness – The Example of Context-Aware Services. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 275–283. Springer, 2009.

[35] C. Dwork. Differential Privacy. *Encyclopedia of Cryptography and Security*, pages 338–340, 2011.

[36] L. Fang and K. LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World wide web*, pages 351–360. ACM, 2010.

[37] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami. Internet of Things (IoT): A vision, architectural elements, and future directions. *Future generation computer systems*, 29(7):1645–1660, 2013.

[38] B. Hill. Confidence and decision. *Games and Economic Behavior*, 82:675–692, 2013.

[39] C. Hine. Privacy in the Marketplace. *The Information Society*, 14(4):253–262, 1998.

[40] J. Hong. The Privacy Landscape of Pervasive Computing. *IEEE Pervasive Computing*, 16(3):40–48, 2017.

[41] T. Hothorn, K. Hornik, and A. Zeileis. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical statistics*, 15(3):651–674, 2006.

[42] C. Hu, J. Zhang, and Q. Wen. An identity-based personal location system with protected privacy in IoT. In *Broadband Network and Multimedia Technology (IC-BNMT), 2011 4th IEEE International Conference on*, pages 192–195. IEEE, 2011.

[43] Z. Huang. A Fast Clustering Algorithm to Cluster Very Large Categorical Data Sets in Data Mining. *Data Mining and Knowledge Discovery*, 3(8):34–39, 1997.

[44] Z. Huang. Extensions to the k-Means Algorithm for Clustering Large Data Sets with Categorical Values. *Data Mining and Knowledge Discovery*, 2(3):283–304, 1998.

[45] C. Jensen, C. Potts, and C. Jensen. Privacy practices of Internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63(1-2):203–227, 2005.

[46] J. B. Kadane and N. A. Lazar. Methods and Criteria for Model Selection. *Journal of the American statistical Association*, 99(465):279–290, 2004.

[47] S. E. Kahou, X. Bouthillier, P. Lamblin, C. Gulcehre, V. Michalski, K. Konda, S. Jean, P. Froumenty, Y. Dauphin, N. Boulanger-Lewandowski, et al. Emonets: Multimodal deep learning approaches for emotion recognition in video. *Journal on Multimodal User Interfaces*, 10(2):99–111, 2016.

[48] R. Kang, S. Brown, L. Dabbish, and S. Kiesler. Privacy Attitudes of Mechanical Turk Workers and the U.S. Public. In *10th Symposium On Usable Privacy and Security*, SOUPS '14, pages 37–49. USENIX Association, 2014.

[49] T. Kang and L. Kagal. Enabling Privacy-Awareness in Social Networks. In *2010 AAAI Spring Symposium Series*, 2010.

[50] S. Karayev, M. Trentacoste, H. Han, A. Agarwala, T. Darrell, A. Hertzmann, and H. Winnemoeller. Recognizing image style. *arXiv preprint arXiv:1311.3715*, 2013.

[51] B. P. Knijnenburg, A. Kobsa, and H. Jin. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies*, 71(12):1144–1162, 2013.

[52] A. Kobsa, H. Cho, and B. Knijnenburg. The effect of personalization provider characteristics on privacy attitudes and behaviors: An Elaboration Likelihood Model approach. *Journal of the Association for Information Science and Technology*, 67(11):2587–2606, 2016.

[53] J. Kolter, T. Kernchen, and G. Pernul. Collaborative privacy management. *computers & security*, 29(5):580–591, 2010.

[54] B. Könings, F. Schaub, and M. Weber. Who, How, and Why? Enhancing Privacy Awareness in Ubiquitous Computing. In *2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pages 364–367. IEEE, 2013.

[55] P. Kumaraguru and L. F. Cranor. Privacy indexes: a survey of Westin's studies. Technical report, Carnegie Mellon University, Pittsburgh, PA, 2005.

[56] M. Langheinrich. A Privacy Awareness System for Ubiquitous Computing Environments. In *Proceedings of the 4th International Conference on Ubiquitous Computing*, UbiComp '02, pages 237–245. Springer, 2002.

[57] N. Lankton, D. McKnight, and J. Tripp. Privacy Management Strategies: An Exploratory Cluster Analysis. In *Proceedings of the 22nd Americas Conference on Information Systems (AMCIS 2016)*, pages 1–10, 2016.

[58] R. LaRose and N. J. Rifon. Promoting i-Safety: Effects of Privacy Warnings and Privacy Seals on Risk Assessment and Online Privacy Behavior. *Journal of Consumer Affairs*, 41(1):127–149, 2007.

[59] S. Lederer, J. Mankoff, and A. K. Dey. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *CHI'03 extended abstracts on Human factors in computing systems*, pages 724–725. ACM, 2003.

[60] H. Lee, C. Upright, S. Eliuk, and A. Kobsa. Personalized object recognition for augmenting human memory. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, pages 1054–1061. ACM, 2016.

[61] Y. Li, A. Kobsa, B. P. Knijnenburg, C. Nguyen, et al. Cross-Cultural Privacy Prediction. *Proceedings on Privacy Enhancing Technologies*, 2017(2):113–132, 2017.

[62] D. J. Liebling and S. Preibusch. Privacy considerations for a pervasive eye tracking world. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, pages 1169–1177. ACM, 2014.

[63] C. S. Liew, T. Y. Wah, J. Shuja, B. Daghighi, et al. Mining personal data using smartphones and wearable devices: A survey. *Sensors*, 15(2):4430–4469, 2015.

[64] J. Lin, B. Liu, N. Sadeh, and J. I. Hong. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Proceedings of the 10th Symposium on Usable Privacy and Security (SOUPS 2014)*, pages 199–212, 2014.

[65] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, A. Acquisti, and Y. Agarwal. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *12th Symposium on Usable Privacy and Security*, SOUPS '16, pages 27–41. USENIX Association, 2016.

[66] C. Liu, S. Chakraborty, and P. Mittal. DEEProtect: Enabling Inference-based Access Control on Mobile Sensing Applications. *arXiv preprint arXiv:1702.06159*, 2017.

[67] M. Madden. Privacy, Security, and Digital Inequality. *New York: Data & Society*, 2017.

[68] T. S. Madhulatha. An overview on clustering methods. *arXiv preprint arXiv:1205.1117*, 2012.

[69] D. Malandrino, A. Petta, V. Scarano, L. Serra, R. Spinelli, and B. Krishnamurthy. Privacy Awareness about Information Leakage: Who knows what about me? In *Proceedings of the 12th ACM Workshop on Privacy in the Electronic Society*, pages 279–284. ACM, 2013.

[70] C. M. Medaglia and A. Serbanati. An Overview of Privacy and Security Issues in the Internet of Things. In *The Internet of Things*, pages 389–395. Springer, 2010.

[71] S. Mehrotra, A. Kobsa, N. Venkatasubramanian, and S. R. Rajagopalan. TIPPERS: A privacy cognizant IoT environment. In *Pervasive Computing and Communication Workshops (PerCom Workshops), 2016 IEEE International Conference on*, pages 1–6. IEEE, 2016.

[72] D. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac. Internet of Things: Vision, applications and research challenges. *Ad hoc networks*, 10(7):1497–1516, 2012.

[73] P. E. Naeini, S. Bhagavatula, H. Habib, M. Degeling, L. Bauer, L. Cranor, and N. Sadeh. Privacy Expectations and Preferences in an IoT World. In *Proceedings of the 13th Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 399–412, 2017.

[74] H. Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.

[75] H. Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.

[76] P. A. Norberg, D. R. Horne, and D. A. Horne. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.

[77] I. I. of Things et al. IoT Scenarios. `https://iot.ieee.org/iot-scenarios.html`. Accessed: 2016-04-28.

[78] X. Page, P. Bahirat, M. I. Safi, B. Knijnenburg, and P. Wisniewski. The Internet of What?: Understanding Differences in Perceptions and Adoption for the Internet of Things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(4):183, 2018.

[79] L. Palen and P. Dourish. Unpacking "Privacy" for a Networked World. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 129–136, New York, NY, USA, 2003. ACM.

[80] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010.

[81] P. Pappachan, M. Degeling, R. Yus, A. Das, S. Bhagavatula, W. Melicher, P. E. Naeini, S. Zhang, L. Bauer, A. Kobsa, et al. Towards Privacy-Aware Smart Buildings: Capturing, Communicating, and Enforcing Privacy Policies and Preferences. In *Distributed Computing Systems Workshops (ICDCSW), 2017 IEEE 37th International Conference on*, pages 193–198. IEEE, 2017.

[82] C. Perera, R. Ranjan, L. Wang, S. U. Khan, and A. Y. Zomaya. Big data privacy in the Internet of Things era. *IT Professional*, 17(3):32–39, 2015.

[83] S. Pötzsch. Privacy Awareness: A Means to Solve the Privacy Paradox? In *IFIP Summer School on the Future of Identity in the Information Society*, pages 226–236. Springer, 2008.

[84] A. J. Ratner, S. H. Bach, H. R. Ehrenberg, and C. Ré. Snorkel: Fast training set generation for information extraction. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1683–1686. ACM, 2017.

[85] 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*, 13(6):401–412, 2009.

[86] A. C. Sarma and J. Girão. Identities in the Future Internet of Things. *Wireless personal communications*, 49(3):353–363, 2009.

[87] M. Shehab, G. Cheek, H. Touati, A. C. Squicciarini, and P.-C. Cheng. User centric policy management in online social networks. In *Policies for Distributed Systems and Networks (POLICY), 2010 IEEE International Symposium on*, pages 9–13. IEEE, 2010.

[88] M. Shehab and H. Touati. Semi-supervised policy recommendation for online social networks. In *Advances in Social Networks Analysis and Mining (ASONAM), 2012 IEEE/ACM International Conference on*, pages 360–367. IEEE, 2012.

[89] S. Shi, M. Zhang, H. Lu, Y. Liu, and S. Ma. Wide & Deep Learning in Job Recommendation: An Empirical Study. In *Asia Information Retrieval Symposium*, pages 112–124. Springer, 2017.

[90] S. Sicari, A. Rizzardi, L. A. Grieco, and A. Coen-Porisini. Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76:146–164, 2015.

[91] A. Sinha, Y. Li, and L. Bauer. What you want is not what you get: Predicting sharing policies for text-based content on Facebook. In *Proceedings of the 2013 ACM Workshop on Artificial Intelligence and Security*, pages 13–24. ACM, 2013.

[92] D. J. Solove. Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.*, 126:1880–1903, 2012.

[93] E. Spyromitros-Xioufis, G. Petkos, S. Papadopoulos, R. Heyman, and Y. Kompatsiaris. Perceived Versus Actual Predictability of Personal Information in Social Networks. In *International Conference on Internet Science*, pages 133–147. Springer, 2016.

[94] J. A. Stankovic. Research directions for the Internet of Things. *IEEE Internet of Things Journal*, 1(1):3–9, 2014.

[95] T. M. Therneau, E. J. Atkinson, et al. An introduction to recursive partitioning using the RPART routines. Technical report, Mayo Foundation, 1997.

[96] G. Tutz and W. Hennevogl. Random effects in ordinal regression models. *Computational Statistics & Data Analysis*, 22(5):537–557, 1996.

[97] J. Virkki and L. Chen. Personal perspectives: Individual privacy in the IoT. *Advances in Internet of Things*, 3(02):21, 2013.

[98] 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)*, 14(3s):64:1–64:24, 2018.

[99] X. Wang and Q. He. Enhancing Generalization Capability of SVM Classifiers with Feature Weight Adjustment. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, pages 1037–1043. Springer, 2004.

[100] Y. Wang and M. Kosinski. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psychology*, 114(2):246–257, 2018.

[101] R. H. Weber. Internet of Things – New security and privacy challenges. *Computer Law & Security Review*, 26(1):23–30, 2010.

[102] T. Winkler and B. Rinner. User-centric privacy awareness in video surveillance. *Multimedia Systems*, 18(2):99–121, 2012.

[103] J. Xie, B. Knijnenburg, and H. Jin. Location Sharing Preference: Analysis and Personalized Recommendation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*, IUI '14, pages 189–198, New York, NY, USA, 2014. ACM.

[104] L. Yang, K. Ting, and M. B. Srivastava. Inferring occupancy from opportunistically available sensor data. In *Pervasive Computing and Communications (PerCom), 2014 IEEE International Conference on*, pages 60–68. IEEE, 2014.

[105] Y. Zhao, J. Ye, and T. Henderson. Privacy-aware Location Privacy Preference Recommendations. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, pages 120–129, Brussels, Belgium, Belgium, 2014. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[106] J. H. Ziegeldorf, O. G. Morchon, and K. Wehrle. Privacy in the Internet of Things: threats and challenges. *Security and Communication Networks*, 7(12):2728–2742, 2014.

Appendices

A Base IoT Scenarios for Privacy Decision Modeling

In this appendix, we list the textual descriptions and factor values of all base scenarios.

S01. At **your home**, your smartphone-connected wearable device collects your **vital signs** (e.g., blood pressure) to let you easily track, monitor, and share data with your doctors, for your **health**.

- Location: **Private**
- Purpose: **Health**
- Core Inference:
  - ***Vital*** $\Rightarrow$ ***Nothing***
- Non-core Inference(s):
  - ***Vital*** $\Rightarrow$ ***Health Information*** (e.g., risk of diseases)
- Impossible Inference(s):
  - ***Vital*** $\Rightarrow$ ***Sexual Orientation*** (e.g., homosexuality)
  - ***Vital*** $\Rightarrow$ ***Social Relationship***

S02. At **your home**, your smart electricity meter collects your **electricity usage** to infer your **energy consumption patterns**, thereby suggesting energy saving methods, for your **savings**.

- Location: **Private**
- Purpose: **Saving**
- Core Inference:
  - ***Electricity Usage*** ⇒ ***Energy Consumption Pattern***
- Non-core Inference(s):
  - ***Electricity Usage*** ⇒ ***Number of Household Members***
  - ***Electricity Usage*** ⇒ ***Presence***
  - ***Electricity Usage*** ⇒ ***Occupancy Pattern***
  - ***Electricity Usage*** ⇒ ***Sleep Pattern***
- Impossible Inference(s):
  - ***Electricity Usage*** ⇒ ***Emotion***
  - ***Electricity Usage*** ⇒ ***Movie Preference***

S03. At **your home**, your smart TV collects your **photo** to infer your **identity**, thereby recommending TV shows based on your watching history, for your **convenience**.

- Location: **Private**
- Purpose: **Convenience**
- Core Inference:
  - ***Photo*** ⇒ ***Identity***
- Non-core Inference(s):
  - ***Photo*** ⇒ ***Demographic Information*** (e.g., age, race, gender)

* ***Photo*** ⇒ ***Emotion***
* ***Photo*** ⇒ ***Personality Type***
* ***Photo*** ⇒ ***Sexual Orientation*** (e.g., homosexuality)

– Impossible Inference(s):

* ***Not specified***

S04. At **your home**, your smart TV collects your **voice** to let you control your TV with **voice commands** (e.g., raise the volume), for your **convenience**.

– Location: **Private**

– Purpose: **Convenience**

– Core Inference:

* ***Voice*** ⇒ ***Device Control Intention***

– Non-core Inference(s):

* ***Voice*** ⇒ ***Identity***
* ***Voice*** ⇒ ***Demographic Information*** (e.g., age, race, gender)
* ***Voice*** ⇒ ***Emotion***
* ***Voice*** ⇒ ***Personality Type***

– Impossible Inference(s):

* ***Voice*** ⇒ ***Physical Activity***

S05. At **your home**, your smart thermostats detect people's **motions** to infer their **presence** in the house, thereby automatically adjusting temperature settings, for your **savings**.

– Location: **Private**

– Purpose: **Saving**

- Core Inference:
  - ***Motion*** ⇒ ***Presence***
- Non-core Inference(s):
  - ***Motion*** ⇒ ***Moving Pattern***
  - ***Motion*** ⇒ ***Occupancy Pattern***
  - ***Motion*** ⇒ ***Sleep Pattern***
- Impossible Inference(s):
  - ***Motion*** ⇒ ***Identity***
  - ***Motion*** ⇒ ***Demographic Information*** (e.g., age, race, gender)

S06. At **your home**, your smart voice assistant collects your **voice** to infer your **identity**, thereby verifying that you are a registered driver to remotely control your connected car through voice commands (e.g., OK Google, start up my car), for your **safety**.

- Location: **Private**
- Purpose: **Safety**
- Core Inference:
  - ***Voice*** ⇒ ***Identity***
- Non-core Inference(s):
  - ***Voice*** ⇒ ***Demographic Information*** (e.g., age, race, gender)
  - ***Voice*** ⇒ ***Emotion***
  - ***Voice*** ⇒ ***Personality Type***
- Impossible Inference(s):
  - ***Voice*** ⇒ ***Dietary Habit***
  - ***Voice*** ⇒ ***Sexual Orientation*** (e.g., homosexuality)

S07. While you are driving **your connected car**, your vehicle collects its **on-board diagnostics (OBD) data** (e.g., RPM, speed, pedal position) to let you easily track, monitor, and share data with vehicle maintenance facilities, for your **safety**.

- Location: **Private**
- Purpose: **Safety**
- Core Inference:
  - ***OBD*** $\Rightarrow$ ***Nothing***
- Non-core Inference(s):
  - ***OBD*** $\Rightarrow$ ***Driving Habit***
- Impossible Inference(s):
  - ***OBD*** $\Rightarrow$ ***Social Relationship***
  - ***OBD*** $\Rightarrow$ ***Political View***
  - ***OBD*** $\Rightarrow$ ***Religion***

S08. While you are driving **your connected car**, your vehicle collects your **photo** to infer your **emotion**, thereby giving real-time alerts if you are emotionally unstable while driving, for your **safety**.

- Location: **Private**
- Purpose: **Safety**
- Core Inference:
  - ***Photo*** $\Rightarrow$ ***Emotion***
- Non-core Inference(s):
  - ***Photo*** $\Rightarrow$ ***Identity***
  - ***Photo*** $\Rightarrow$ ***Demographic Information*** (e.g., age, race, gender)

    - * ***Photo*** ⇒ ***Personality Type***
    - * ***Photo*** ⇒ ***Sexual Orientation*** (e.g., homosexuality)
  - Impossible Inference(s):
    - * ***Photo*** ⇒ ***Financial Condition***

S09. At **your workplace**, your employer's face recognition device collects your **photo** to infer your **identity**, thereby allowing you to freely enter the building without a badge, for your **convenience**.

  - Location: **Workplace**
  - Purpose: **Convenience**
  - Core Inference:
    - * ***Photo*** ⇒ ***Identity***
  - Non-core Inference(s):
    - * ***Photo*** ⇒ ***Presence***
    - * ***Photo*** ⇒ ***Attendance Pattern***
    - * ***Photo*** ⇒ ***Emotion***
    - * ***Photo*** ⇒ ***Personality Type***
    - * ***Photo*** ⇒ ***Sexual Orientation*** (e.g., homosexuality)
  - Impossible Inference(s):
    - * ***Photo*** ⇒ ***Cognitive Activity***
    - * ***Photo*** ⇒ ***Energy Consumption Pattern***

S10. At **your workplace**, your employer's network devices collect a **unique ID of your mobile device** (e.g., Wi-Fi MAC address of your phone) to infer your **location**, thereby helping you find a specific meeting room, for your **convenience**.

  - Location: **Workplace**

- Purpose: **Convenience**
- Core Inference:
  - ***Device ID* ⇒ *User Location***
- Non-core Inference(s):
  - ***Device ID* ⇒ *Identity***
  - ***Device ID* ⇒ *Moving Pattern***
- Impossible Inference(s):
  - ***Device ID* ⇒ *Emotion***
  - ***Device ID* ⇒ *Physical Activity***

S11. At **your workplace**, your employer's smart CCTVs collect your **video** to infer your **physical activity** (e.g., types of physical movement) inside the building, thereby providing you personalized health recommendations (e.g., you need to walk at least 30 minutes per day), for your **health**.

- Location: **Workplace**
- Purpose: **Health**
- Core Inference:
  - ***Video* ⇒ *Physical Activity***
- Non-core Inference(s):
  - ***Video* ⇒ *Presence***
  - ***Video* ⇒ *Attendance Pattern***
  - ***Video* ⇒ *Social Relationship***
  - ***Video* ⇒ *Moving Pattern***
  - ***Video* ⇒ *Health Information*** (e.g., risk of diseases)
- Impossible Inference(s):

    - ***Video*** $\Rightarrow$ ***Cognitive Activity***
    - ***Video*** $\Rightarrow$ ***Political View***

S12. At **your workplace**, your employer's smart CCTV collects your **video** to infer your **presence** at your office, thereby automatically managing your attendance (i.e., auto-timesheet), for your **convenience**.

- Location: **Workplace**
- Purpose: **Convenience**
- Core Inference:
    - ***Video*** $\Rightarrow$ ***Presence***
- Non-core Inference(s):
    - ***Video*** $\Rightarrow$ ***Attendance Pattern***
    - ***Video*** $\Rightarrow$ ***Social Relationship***
    - ***Video*** $\Rightarrow$ ***Physical Activity***
    - ***Video*** $\Rightarrow$ ***Smoking Habit***
- Impossible Inference(s):
    - ***Not specified***

S13. At a **public airport**, US government's face recognition device collects your **photo** to infer your **identity**, thereby expediting identity verification process with your machine-readable passport, for your **convenience**.

- Location: **Public**
- Purpose: **Convenience**
- Core Inference:
    - ***Photo*** $\Rightarrow$ ***Identity***

- Non-core Inference(s):
  - ***Photo*** ⇒ ***Presence***
  - ***Photo*** ⇒ ***Attendance Pattern***
  - ***Photo*** ⇒ ***Emotion***
  - ***Photo*** ⇒ ***Personality Type***
  - ***Photo*** ⇒ ***Sexual Orientation*** (e.g., homosexuality)
- Impossible Inference(s):
  - ***Not specified***

S14. At a **public airport**, US government's network devices collect a **unique ID of your mobile device** (e.g., Wi-Fi MAC address of your phone) to infer your **location**, thereby helping you find a boarding gate, for your **convenience**.

- Location: **Public**
- Purpose: **Convenience**
- Core Inference:
  - ***Device ID*** ⇒ ***User Location***
- Non-core Inference(s):
  - ***Device ID*** ⇒ ***Identity***
  - ***Device ID*** ⇒ ***Moving Pattern***
- Impossible Inference(s):
  - ***Device ID*** ⇒ ***Personality Type***
  - ***Device ID*** ⇒ ***Health Information*** (e.g., risk of diseases)

S15. At a **public airport**, US government's smart CCTVs collect your and others' **video** to infer their **identity**, thereby identifying wanted criminals or suspicious people around you, for your **safety**.

- Location: **Public**
- Purpose: **Safety**
- Core Inference:
  - ***Video* ⇒ *Identity***
- Non-core Inference(s):
  - ***Video* ⇒ *Presence***
  - ***Video* ⇒ *Attendance Pattern***
  - ***Video* ⇒ *Social Relationship***
  - ***Video* ⇒ *Physical Activity***
  - ***Video* ⇒ *Moving Pattern***
- Impossible Inference(s):
  - ***Video* ⇒ *Financial Condition***

B IoT Scenarios for Privacy Segmentation

In this appendix, we list the three IoT scenarios used for privacy segmentation.

S1. A device of **Starbucks** (*who*) takes a **photo** of you to determine your **mood** (*what*). This happens **once** (*persistence*), while you are at **Starbucks** (*where*), for **commercial** (*reason*) purposes, namely to suggest the most suitable product to you based on your mood.

S2. A device of **your employer** (*who*) collects your **phone ID** (e.g., Wi-Fi MAC address) to verify your **identity** (*what*). This happens **continuously** (*persistence*), while you are at **your workplace** (*where*), for your **convenience** (*reason*), namely to automate mundane tasks (e.g., making a payment at the cafeteria).

S3. A device of **the government** (*who*) takes a **video** (*what*) of you. This happens **continuously** (*persistence*), while you are in **a public park** (*where*), for your **safety** (*reason*), namely to identify wanted criminals around you.

We also present a questionnaire for collecting user reactions to the abovementioned scenarios.

Q1. Would you want to be **notified** (*notification*) about this monitoring?

1. Yes, always
2. Yes, but ask me again the next time
3. No, but ask me again the next time
4. No, never

Q2. Would you want to **allow** (*permission*) this monitoring?

1. Yes, always
2. Yes, but ask me again the next time
3. No, but ask me again the next time
4. No, never

Q3. How **comfortable** (*comfort*) is this monitoring?

1. Very uncomfortable
2. Uncomfortable
3. Somewhat uncomfortable
4. Neutral
5. Somewhat comfortable
6. Comfortable

7. Very comfortable

Q4. How **risky** (*risk*) is this monitoring?

1. Very risky
2. Risky
3. Somewhat risky
4. Neutral
5. Somewhat safe
6. Safe
7. Very safe

Q5. How **appropriate** (*appropriateness*) is this monitoring?

1. Very inappropriate
2. Inappropriate
3. Somewhat inappropriate
4. Neutral
5. Somewhat appropriate
6. Appropriate
7. Very appropriate

C Open-ended Questions regarding Confident Privacy Decision-Making

In this appendix, we list open-ended questions used in our online survey. Survey participants responded to these questions right after they finished the main part of the survey (i.e., privacy

decision modeling).

We now ask you a few open-ended questions regarding privacy in IoT.

Q1. In general, have you ever cared about the possibility of the **undesired disclosure of your personal information** before using intelligent functionalities of your smart device? For instance, you can activate your smartphone through face unlock, but your face photo might also be used to infer your age. If you have cared about this, please explain how it impacted your usage of that function or device.

*For each IoT scenario, we asked you about the possible inferences from collected sensor data (e.g., Photo ⇒ Emotion). This was intended to measure your degree of awareness of potential privacy risks posed by using an IoT service (that is, your **privacy awareness**).*

Q2. Keeping in mind all the IoT scenarios you encountered, how difficult or easy was it to answer **privacy awareness** questions? Please explain why you felt it was difficult or easy.

*After you answered a privacy awareness question for a specific IoT scenario, you were asked to make a decision on whether or not to use the IoT service (that is, your **privacy decision**).*

Q3. Keeping in mind all the IoT scenarios you encountered, do you think your responses to **privacy awareness** questions impacted your **privacy decisions**? If so, please explain how it impacted them.

Q4. Considering all the IoT scenarios that you decided NOT to use, what were the reasons for not using IoT services (that is, make a **conservative** decision)?

*You were also asked to indicate your confidence about your privacy decision for each IoT scenario (that is, your **privacy decision confidence**).*

Q5. Keeping in mind all the IoT scenarios you encountered, do you think your responses to **privacy awareness** questions impacted your **privacy decision confidence**? If so, please explain how it impacted your confidence.

Q6. Keeping in mind all the IoT scenarios for which you made **confident** decisions, what made you **confident** using or not using the IoT service?

Q7. Do you think making a **confident privacy decision** is important in protecting your privacy while using the IoT service? If so, please explain why and how it matters.