A Data-Driven Approach to Designing for Privacy in Household IoT

YANGYANG HE, PARITOSH BAHIRAT, and BART P. KNIJNENBURG, Clemson University, Clemson, South Carolina
ABHILASH MENON, IBM, Armonk, New York

In this article, we extend and improve upon a previously developed data-driven approach to design privacy-setting interfaces for users of household IoT devices. The essence of this approach is to gather users' feedback on household IoT scenarios *before* developing the interface, which allows us to create a navigational structure that preemptively maximizes users' efficiency in expressing their privacy preferences, and develop a series of 'privacy profiles' that allow users to express a complex set of privacy preferences with the single click of a button. We expand upon the existing approach by proposing a more sophisticated translation of statistical results into interface design, and by extensively discussing and analyzing the tradeoff between user-model parsimony and accuracy in developing privacy profiles and default settings.

CCS Concepts: • **Security and privacy** → **Usability in security and privacy**; • **Human-centered computing** → **User studies**;

Additional Key Words and Phrases: Designing for IoT, privacy

ACM Reference format:

Yangyang He, Paritosh Bahirat, Bart P. Knijnenburg, and Abhilash Menon. 2019. A Data-Driven Approach to Designing for Privacy in Household IoT. *ACM Trans. Interact. Intell. Syst.* 10, 1, Article 10 (September 2019), 47 pages.
<https://doi.org/10.1145/3241378>

1 INTRODUCTION

Our everyday life is being revolutionized by all kinds of smart electronic household devices, often known 'smart home' technology. Smart home technology is made possible by recent developments around the Internet of Things (IoT), and is still in its fledgling stages of development. A study by PWC [51], suggests that lower levels of Household IoT adoption are primarily due to high cost of ownership. Interestingly, the second-biggest reason of hesitation towards adoption is privacy and security concerns [51]. Arguably, such concerns may rise as costs decrease and adoption increases.

This work was supported by NSF award no. 1640664 and a gift from Samsung Research America.

The reviewing of this article was managed by special issue associate editors Mark Billingham, Margaret Burnett, and Aaron Quigley.

Authors' addresses: Y. He, P. Bahirat, and B. P. Knijnenburg, School of Computing, 821 McMillan Rd, Clemson, SC 29631; emails: {yyhe, pbahira}@g.clemson.edu, bartk@clemson.edu; A. Menon, IBM, 1 New Orchard Rd, Armonk, NY, 10504; email: abhilash.menon3@ibm.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

2160-6455/2019/09-ART10 \$15.00

<https://doi.org/10.1145/3241378>

Privacy is an inherent tradeoff in IoT, because IoT devices cannot provide their services without collecting data. Moreover, many IoT devices provide *personalized* services, which requires them to retain and process the data as well. Preserving users' privacy therefore means giving them control over this tradeoff, by allowing them to decide what information can be collected about them. Most household IoT devices have extensive privacy settings, but these are often hidden in the interfaces of the individual devices, making it difficult for users to set their preferences across their 'smart home' [24]. Having disparate settings on several devices is also at odds with the fact that these devices are usually interconnected, sharing both user data and opportunities for control across devices.

A promising solution to more suitably account for users' privacy preferences would be to allow them to control their privacy with global settings for the entire 'smart home'. But even with a centralized settings interface (which currently does not exist), users will likely still find it difficult to translate their preferences into the correct privacy settings [1, 15, 32]. Indeed, the modern 'smart home' has a vast number of household IoT devices—each with several features and interconnections to other devices—which makes choosing adequate privacy settings a very challenging task that is likely to result in information and choice overload [56].

What design process allows us to develop a usable privacy-setting interface for household IoT? The development of usable privacy interfaces commonly relies on user studies with existing systems. However, this method is not possible in our household IoT scenario, because there currently exists no centralized household IoT privacy-setting interface. We therefore propose to develop user interface designs for managing the privacy settings of Household IoT devices using a *data driven design* approach: rather than evaluating and incrementally improving an existing interface, we gather users' feedback on household IoT scenarios *before* developing the interface. This approach allows us to create a navigational structure that preemptively maximizes users' efficiency in expressing their privacy preferences. Moreover, it allows us to anticipate the most common privacy settings, and capture them into a series of 'privacy profiles' that allow users to express a complex set of privacy preferences with the single click of a button.

We first introduced this approach in our previous work [5], where we covered a wide variety of in-home as well as public IoT scenarios. In this article, we focus on household IoT in particular, and further refine our approach to allow us to create more carefully tailored user interfaces. Moving the context to a more narrow environment shifts the focus of the privacy decision from the entity collecting information (which was the dominant parameter in our previous work) to a more contextual evaluation of the content or nature of the information [33]. This results in more complex decisions, and thereby advances our previous approach.

The main contributions of our article are:

- Using an intricate mixed fractional factorial study design, we collect a dataset of 1,133 participants making 13,596 privacy decisions on 4,608 scenarios.
- We perform statistical analysis on this dataset to develop a layered IoT privacy-setting interface. As our analysis shows more complex decision patterns than our previous work, we present guidelines to translate our statistical results into a more sophisticated settings interface design.
- We perform machine learning analysis on our dataset to create a set of "smart profiles" for our IoT privacy-setting interface. Beyond our previous work, we conduct a deeper analysis regarding the tradeoff between parsimony and accuracy of our prediction models, leading to a better-informed selection of smart profiles.
- Aside from the privacy-setting interface and the smart profiles, we make specific design recommendations for household IoT devices that can help to minimize users' privacy concerns.

The remainder of this article is structured as follows: Section 2 summarizes key elements of related work. Section 3 provides relevant information of our experiment setup. Section 4 describes our statistical results which is followed by Section 5 where we discuss our preliminary prototype for privacy settings. Section 6 presents the machine learning analysis to create a set of “smart profiles”. Section 7 presents two separate interfaces which are modified version of the preliminary interface to carefully accommodate results from Machine Learning analysis. Finally, Section 9 concludes with a summary of contributions and pointers to future work.

2 RELATED WORK

Our goal is to develop intuitive interfaces for IoT privacy settings, using a data-driven approach. In this section we therefore discuss existing research on privacy-setting interfaces and on privacy prediction.

2.1 Personalization in IoT Systems

One of the key features of IoT environments is that they have a high potential for providing personalized services to their users [10, 13, 52]. For example, Russell et al. [42] use unobtrusive sensors and micro-controller to realize a human detection for further providing personalization in a scenario of a family making use of the IoT in their daily living. Henka et al. [16] propose an approach to personalize services in (household) IoT using the Global Public Inclusive Infrastructure’s [53] preference set to describe an individual’s needs and preferences, and then adapting a smart environment accordingly.

2.2 Privacy in Personalized Systems

Researchers have shown that privacy can play a limiting role in users’ adoption of personalized services [48]. For example, Awad and Krishnan [3] show that privacy concerns inhibit users’ use of personalized services, and Sutanto et al. [47] demonstrated that privacy concerns can prevent people from using a potentially beneficial personalized application. Kobsa et al. [27] demonstrate that the personalization provider is an important determinant of users’ privacy concerns.

Moreover, research has shown users’ willingness to provide personal information to personalized services depends on both the risks and benefits of disclosure [17, 19, 39], and researchers therefore claim that both the benefits and the risks meet a certain threshold [49], or that they should be in balance [7].

2.3 Privacy in IoT

The argument that using user-generated data for personalization can result in privacy concerns has also been made in IoT environments [59]. One of the first examples in this regard was the work by, Sheng et al. [45], who showed that users of “u-commerce” services (IoT-driven mobile shopping) felt less inclined to use personalized (rather than non-personalized) u-commerce services, unless the benefits were overwhelming (i.e., providing help in an emergency).

In response, researchers have proposed frameworks with guidelines for evaluating the security and privacy of consumer IoT applications, devices, and platforms [31, 38]. Most of these guidelines are focused on minimizing data acquisition, storage, and collection sources. Along these guidelines, several researchers have proposed architectures that restrict unwanted access to users’ data by IoT devices. For example, Davies et al. [8] propose “privacy mediators” to the data distribution pipeline that would be responsible for data redaction and enforcement of privacy policies even before the data is released from the user’s direct control. Likewise, Jayraman et al.’s [20] privacy preserving architecture aggregates requested data to preserve user privacy.

Other research has considered IoT privacy from the end-user perspective [12], both when it comes to research (e.g., Ur et al. [50] investigated how privacy perceptions differ among teens and their parents in smart security systems installed in homes) and design (e.g., Williams et al. [56] highlight the importance of designing interfaces to manage privacy such that they are usable to the end users of IoT devices, and Feth et al. [12] investigated the creation of understandable and usable controls). The current article follows this approach by outlining a novel methodology for the development of usable and efficient privacy-setting interfaces and applying it to household IoT privacy management.

2.4 Existing Privacy Control Schemes

Smartphones give users control over their privacy settings in the form of prompts that ask whether the user allows or denies a certain app access to a certain type of information. Such prompts are problematic for IoT, because IoT devices are supposed to operate in the background. Moreover, as the penetration of IoT devices in our homes continues to increase, prompts would become a constant noise which users will soon start to ignore, like software EULAs [14] or privacy policies [21].

Pejovic and Musolesi [37] presented the design and implementation of an efficient online learner that can serve as a basis for recognizing opportune moments for interruption. The design of the library is based on an in-depth study of human interruptibility. Comparatively, our work tries to find the most suitable privacy-setting profile for each user based on their privacy preference on different household IoT scenarios.

2.5 Privacy-Setting Interfaces

Beyond prompts, one can regulate privacy with global settings. The most basic privacy-setting interface is the traditional “access control matrix”, which allows users to indicate which entity gets to access what type of information [44]. This approach can be further simplified by grouping recipients into relevant semantic categories, such as Google+’s *circles* [55]. Taking a step further, Raber et al. [40] proposed *Privacy Wedges* to manipulate privacy settings. Privacy Wedges allow users to make privacy decisions using a combination of semantic categorization (the various wedges) and inter-personal distance (the position of a person on the wedge). Users can decide who gets to see various posts or personal information by “coloring” parts of each wedge.

Privacy wedges have been tested on limited numbers of friends, and in the case of household IoT they are likely to be insufficient, due to the complexity of the decision space. To wit, IoT privacy decisions involve a large selection of devices, each with various sensors that collect data for a range of different purposes. This makes it complicated to design an interface that covers every possible setting [56]. A wedge-based interface will arguably not be able to succinctly represent such complexity, and therefore either be impossible, or still lead to a significant amount of information and choice overload.

We propose a data-driven approach to solve this problem: statistical analysis informs the construction of a layered settings interface, while machine-learning-based privacy prediction helps us find smart privacy profiles.

2.6 Privacy Prediction

Several researchers have proposed privacy prediction as a solution to the privacy settings complexity problem—an approach known as “user-tailored privacy” (UTP) [23]. Systems that implement UTP first predict users’ privacy preferences and behaviors based on their known characteristics. They then use these predictions to provide automatic default settings or suggestions in line with users’ disclosure profiles, to educate users’ about privacy features they are unaware of, to tailor the privacy-setting user interfaces to make it easier for users to engage with their preferred

privacy management tools, or to selectively restrict the types of personalization in which a system is allowed engage.

Most existing work in line with this approach has focused on providing automatic default settings. For example, Sadeh et al. [43] used a k -nearest neighbor algorithm and a random forest algorithm to predict users' privacy preferences in a location-sharing system, based on the type of recipient and the time and location of the request. They demonstrated that users had difficulties setting their privacy preferences, and that the applied machine learning techniques can help users to choose more accurate disclosure preferences. Similarly, Pallapa et al. [36] present a system which can determine the required privacy level in new situations based on the history of interaction between users. Their system can efficiently deal with the rise of privacy concerns and help users in a pervasive system full of dynamic interactions.

Dong et al. [9] use a binary classification algorithms to give users personalized advice regarding their privacy decision-making practices on online social networks. They found that J48 decision trees provided the best results. Li et al. [29] similarly use J48 to demonstrate that taking the user's cultural background into account when making privacy predictions improves the prediction accuracy. Our data stems from a culturally homogeneous population (U.S. Mechanical Turk workers), so cultural variables are outside the scope of our study. We do however follow these previous works in using J48 decision trees in our prediction approach.

We further extend this approach using *clustering* to find several smart default policies ("profiles"). This is in line with Fang and LeFevre [11], who present an active learning algorithm that comes up with privacy profiles for users in real time. Since our approach is based on an existing dataset, our algorithm does not classify users in real time, but instead creates a static set of profiles 'offline', from which users can subsequently choose. This avoids cold start problems, and does not rely on the availability of continuous real-time behaviors. This is beneficial for household IoT privacy settings, because users often specify their settings in these systems in a "single shot", leaving the settings interface alone afterwards.

Ravichandran et al. [41] employ an approach similar to ours, using k -means clustering on users' contextualized location sharing decisions to come up with several default policies. They showed that a small number of policies could accurately reflect a large part of the location-sharing preferences. We extend their approach to find the best profiles based on various novel clustering approaches, and take the additional step of designing user interfaces that incorporate the best solutions.

2.7 Data-Driven Design

In our previous work [5], we leveraged data collected by Lee and Kobsa [28], which asked 200 participants about their intention to allow or reject the IoT features presented in 14 randomized scenarios. They varied the scenarios in a mixed fractional factorial design along the following dimensions: 'Who', 'What', 'Where', 'Reason', 'Persistence'.

We conducted a statistical analysis on this dataset to determine the relative influence of these parameters on users' privacy-related decisions. The outcome informed the design of a 'layered interface', which presents privacy settings with the most prominent influence first, relegating less prominent aspects to subsequently lower layers. Users can use this interface for making manual privacy settings.

We also conducted a machine learning analysis to predict participants' reactions to the scenarios. We used the outcomes of this analysis to develop a "smart" default setting, which preempts the need for many users to manually change their settings [46]. However, since people differ extensively in their privacy preferences [35], it is not possible to achieve an optimal default that is the same for everyone. Instead, different people may require vastly different settings [25, 35, 57].

By partitioning the participants in a number of clusters, we were able to construct a number of ‘privacy profiles’, which represented a selection of default settings for the user to choose from. These profiles automate (part of) the privacy-setting task.

As noted in the introduction, our current paper builds upon this existing work by applying it to a newly collected dataset focused on household IoT privacy decisions, and by refining both the statistical and machine learning procedures underlying this approach. The resulting procedure can be considered a blueprint for researchers interested in applying data-driven design to their (privacy-)settings interfaces.

3 EXPERIMENTAL SETUP

To develop our data-driven design approach, we collected a new dataset of users’ privacy behaviors in various IoT contexts. In this section, we first discuss the factorial procedure by which we developed 4,608 highly specific IoT scenarios, as well as the questions we asked participants to evaluate these scenarios. We then describe the participant selection and experimental procedures used to collect over 13,500 responses from 1,133 participants.

3.1 Contextual Scenarios

The scenarios evaluated in our study are based on a full factorial combination of five different Parameters: Who, What, Purpose, Storage, and Action. A total of $8(\text{who}) * 12(\text{what}) * 4(\text{purpose}) * 4(\text{storage}) * 3(\text{action}) = 4,608$ scenarios were tested this way.

The scenarios asked participants to imagine that they were owners and active users of the presented IoT devices, trying to decide whether to turn on or off certain functionalities and/or data-sharing practices. To avoid endowment effects, the scenarios themselves made no indication as to whether the functionality was currently turned on or off (such endowment effects were instead introduced by manipulating the framing of the Decision question; see Section 3.2). An example scenario is: “*Your smart TV (Who) uses a camera (What) to give you timely alerts (Purpose). The data is stored locally (Storage) and used to optimize the service (Action).*” This scenario may, for example, represent a situation where the smarthome system has detected (via camera) a delivery of package and then alerts the user (via the smart TV) about its arrival. In this particular scenario, we note that the video data is stored locally to optimize service; this could mean that the smarthome system uses the video stream to (locally) train a package detection algorithm. Similarly, another example of scenario is: “*Your Smart Assistant uses a microphone to detect your location in house. The data is stored on a remote server and shared with third parties to recommend you other services.*” Similarly, this scenario could represent a situation where the smarthome has detected (via microphone) it’s user’s location in the house and this information is shared to smart assistant. In the scenario, the data is stored on remote server and shared with third parties so that it can recommend additional services (like weather or local transportation) via third parties to the user.

The levels of all five parameters used in our experiment are shown in Table 1. The parameters were highlighted in the scenario for easy identification, and upon hovering the mouse cursor over them each parameter would show a succinct description of the parameter. Figure 34 in the Appendix shows a screenshot of a scenario as shown to participants in the study. A thirteenth scenario regarding the interrelated control of various IoT devices (e.g., “*You can use your smart TV to control your smart refrigerator*”) was also asked, but our current analysis focuses on the information-sharing scenarios only.

The parameters used in the current study deviate from the ones in the Lee and Kobsa [28] dataset. In our previous work, we observed that the *Where* parameter in this dataset did not have a significant effect on user decision making [5], hence we removed it from the scenarios in the

Table 1. Parameters Used to Construct the Information-Sharing Scenarios

Parameter	Levels	Code
Who: <i>Your Smart...</i>	1. Home Security System 2. Refrigerator 3. HVAC System 4. Washing Machine 5. Lighting System 6. Assistant 7. TV 8. Alarm Clock	SS RE HV WM SL SA TV SC
What: <i>...uses information collected by your...</i>	1. Home Security System 2. Refrigerator 3. HVAC System 4. Washing Machine 5. Lighting System 6. Assistant 7. TV 8. Alarm 9. uses a location sensor 10. uses a camera 11. uses a microphone 12. connects to your smart phone/watch	CSE CRE CHV CWA CLI CAS CTV CAL CLO CCA CMP CSW
Purpose : <i>...to...</i>	1. detect whether you are home 2. detect your location in house 3. automate its operations 4. give you timely alerts	PH LH AO TA
Storage: <i>The data is stored...</i>	1. locally 2. on remote server 3. on a remote server and shared with third parties	L R T
Action: <i>...and used to...</i>	1. optimize the service 2. give insight into your behavior 3. recommend you other services 4. [None]	O I R N

The “codes” are used as abbreviations in graphs and figures throughout the paper and the Appendix.

current study. Likewise, in public, IoT encounters are often ephemeral, so persistent tracking is rather uncommon. Hence, we removed *Persistence* of tracking from the scenarios as well, since this parameter is much more relevant in public IoT than in household IoT. The original *Reason* parameter is similar to the current *Purpose* and *Action* parameters, although the reasons/purpose for tracking are obviously different in public IoT than in household IoT.

Moreover, we learned from the qualitative feedback in our previous study that the secondary use of information was a prominent concern among users of IoT systems. Hence, we consider *Purpose* as the primary purpose of tracking, separate from *Action*, a secondary purpose that requires *Storage*—a parameter we added because it is possible for household IoT systems to operate (and thus store data) locally, and because the sharing of data with third parties is not as common in household IoT as in public IoT.

3.2 Scenario Evaluation Questions

The first question participants were asked about each scenario was whether they would enable or disable the particular feature mentioned in scenario (Decision). Subsequently, they were asked about their attitudes regarding the scenario in terms of their perceived Risk, Appropriateness, Comfort, Expectedness and Usefulness regarding the presented scenario (e.g., “*How appropriate do you think this scenario is?*”). These questions were answered on a 7-point scale (e.g., “*very inappropriate*” to “*very appropriate*”). In every 4th scenario, the Risk and Usefulness questions were followed by an open question asking the participants to describe the potential Risk and Usefulness of the scenario. We asked these question mainly to encourage participants to carefully evaluate the scenarios. A screenshot of the questions asked about each scenario is depicted in Figure 34 in the Appendix.

The framing and default of the Decision question were manipulated between-subjects at three levels each: positive framing (“Would you enable this feature?”, options: Yes/No), negative framing (“Would you disable this feature?”, options: Yes/No) or neutral framing (“What would you do with this feature?”, options: Enable/Disable); combined with a positive default (enabled by default), negative default (disabled by default), or no default (forced choice).

3.3 Participants and Procedures

To collect our dataset, 1133 adult U.S.-based participants (53.53% Female, 45.75% Male, 8 participants did not disclose) were recruited through Amazon Mechanical Turk. Participation was restricted to Mechanical Turk workers with a high reputation (at least 50 completed tasks, average accuracy of >95%). Participants were paid \$2.00 upon successful completion of the study. The participants were warned about not getting paid in case they failed attention checks (see below).

The study participants represented a wide range of ages, with 9 participants less than 20 years old, 130 aged 20–25, 273 aged 25–30, 418 aged 30–40, 175 aged 40–50, 80 aged 50–60, and 43 participants over 60 years old (5 participants did not disclose their age). This significant increase in participants over the Lee and Kobsa [28] dataset is commensurate with our expectation of more complex privacy decision behaviors in household IoT compared to public IoT.

Each participant was first shown a video with a brief introduction to various smart home devices, which also mentioned various ways in which the different appliances would cooperate and communicate within a home. After the video, participants were asked to answer three attention-check questions depicted in Figure 30 in the Appendix. If they got any of these questions wrong, they would be asked to read the transcript of the video and re-answer the questions. The transcript is depicted in Figure 31 in the Appendix.

After the introduction video, each participant was presented with 12 information-sharing scenarios (and a 13th control scenario, not considered in this article). These scenarios were selected from the available 4608 scenarios using fractional factorial design¹ that balances the within- and between-subjects assignment of each parameter’s main effect, and creates a uniform exposure for each participant to the various parameters (i.e., to avoid “runs” of near-similar scenarios). Participants were asked to carefully read the scenario and then answer all questions about it. Two of the 13 scenarios had an additional attention-check question (e.g., “Please answer this question with Completely Agree”, see Figure 33 in the Appendix), and there was an additional attention check question asking participants about the remaining time to finish the study (which was displayed right there on the same page, see Figure 32 in the Appendix). Participants rushing through the experiment and/or repeatedly failing the attention-check questions were removed from the dataset.

¹The scenario assignment scheme is available at <https://www.usabart.nl/scenarios.csv>.

Fig. 1. Different tests conducted for mediation analysis.

4 INSPECTING USERS' DECISIONS

In this section we explain the different regression analyses performed on the dataset to understand how different scenario parameters affected the decisions made by participants. We begin by explaining the effects of scenario parameters on participants' decision to enable or disable the feature mentioned in the scenario. Similar to Ajzen and Fishbein [2] and Bahirat et al. [5], we also present the results of the mediation analysis, which are on the lines of attitude-behavior models. As shown in Figure 1, we test whether participants' attitudes mediate the effects of the scenario parameters on their decisions. This mediation analysis involves the following tests:

- Test 1:* The effect of parameters (Who, What, Purpose, Storage, Action) on attitudes (Risk, Comfort, Appropriateness, Expectedness and Usefulness).
- Test 2:* The effect of attitudes on decision.
- Test 3:* The effect of both parameters *and* attitudes on decision.

If tests 1 and 2 are significant and test 3 reveals a drastic reduction in the conditional direct effect of parameters, then we can say that the effects of scenario parameters on participant's decision are mediated by their attitudes [5].

Finally, we present a post-hoc analysis of differences between individual levels of the parameters on attitudes and decision.

4.1 Effect of Scenario Parameters on Decision

To understand the effect of the scenario parameters on participants' allow/reject decision, we developed a *generalized linear mixed effects regression (glmer)* with a random intercept (to account for repeated measures on the same participant) and a logit link function (to account for the fact that the outcome variable is binary). We used a forward stepwise approach, where we added the strongest remaining parameter into the model at each step and then comparing it using ANOVA tests against the previous model. If new parameter makes a significant improvement to the previous model, it has a significant overall effect on the outcome variable. Once all significant main effects are added to the model, two-way interaction effects are tested one by one.

Table 2 shows the effects of the parameters on the allow/reject decision. All parameters had a significant effect. Particularly, **Storage** had the strongest effect on participants' decisions, followed by **What**, **Who**, and **Purpose** (all similar), and finally **Action**.

Moreover, we find many significant interaction effects, but some of them are not substantial compared to the main effects.² Substantial two-way interaction effects were observed between **Who**, **What**, and **Purpose**. It should be noted that the interactions are added separately, not accumulatively. This reduces overfitting and multicollinearity.

²Very small but still significant interaction effects are a common occurrence in the analysis of large datasets.

Table 2. Effect of Scenario Parameters on Decision

Model	χ^2	<i>df</i>	<i>p</i> -value
<i>decision</i> ~ (1 <i>sid</i>)			
+storage	1487.76	2	<.0001
+purpose	206.97	11	<.0001
+what	202.48	3	<.0001
+who	195.91	7	<.0001
+action	77.20	3	<.0001
<i>Interactions</i>			
+what:who	138.03	77	<.0001
+who:purpose	87.92	21	<.0001
+what:purpose	68.30	33	.0002

Table 3. Effect of Scenario Parameters on Appropriateness

Model	<i>df</i>	<i>Chi.Sq.</i>	<i>p</i> -value
<i>appropriateness</i> ~ (1 <i>sid</i>)	3		
+storage	5	2346.19	<.0001
+what	16	398.63	<.0001
+purpose	19	359.98	<.0001
+who	26	179.09	<.0001
+action	29	91.05	<.0001
<i>Interactions</i>			
+what:who	106	167.01	<.0001
+who:purpose	50	113.73	<.0001
+what:purpose	62	55.67	.0081

4.2 Effect of Scenario Parameters on Attitudes

Test 1 of the mediation model is a test of the effect of the scenario parameters on participants' attitudes. For this, we developed a separate *linear mixed effects regression model (lmer)* with a random intercept (to account for repeated measures on the same participant) for each dependent variable (Risk, Comfort, Appropriateness, Expectedness, and Usefulness), using the scenario parameters as independent variables. As in the previous section, we took a forward stepwise approach.

Tables 3–7 show the effects of the parameters on the different attitudes. All parameters had a significant effect on all attitudes. Substantial two-way interaction effects were again observed between **Who**, **What**, and **Purpose**. Again, the interactions are added separately, not accumulatively.

4.3 Effect of Attitudes on Decision

Test 2 of the mediation model is a test of the effect of participants' attitudes on their allow/reject decision. We perform this test by creating a *glmer* model with a random intercept and a logit link function. Using a forward stepwise approach, we find that all attitudes except **Expectedness** have a significant effect on decision (see the top part of Table 8). Specific effects are as follows:

- Each 1-point increase in **Comfort** (measured on a 7-point scale) results in a 2.30-fold increase in the odds that the participant will allow the scenario ($p < 0.001$).

Table 4. Effect of Scenario Parameters on Comfort

Model	<i>df</i>	<i>Chi.Sq.</i>	<i>p</i> -value
<i>comfort</i> ~ (1 <i>sid</i>)	3		
+storage	5	2822.57	<.0001
+what	16	391.10	<.0001
+purpose	19	381.69	<.0001
+action	22	113.68	<.0001
+who	29	90.57	<.0001
<i>Interactions</i>			
+what:who	106	132.86	<.0001
+who:purpose	50	89.20	<.0001
+what:purpose	62	58.24	.0043

Table 5. Effect of Scenario Parameters on Risk

Model	<i>df</i>	<i>Chi.Sq.</i>	<i>p</i> -value
<i>risk</i> ~ (1 <i>sid</i>)	3		
+storage	5	47240.72	<.0001
+purpose	16	421.08	<.0001
+action	19	355.65	<.0001
+who	26	81.35	<.0001
+what	29	70.64	<.0001
<i>Interactions</i>			
+what:who	106	77.14	0.0017
+who:purpose	50	19.91	<.0001
+what:purpose	62	37.19	0.0352

- Each 1-point increase in **Usefulness** results in a 2.09-fold increase in the odds that the participant will allow the scenario ($p < 0.001$).
- Each 1-point increase in **Appropriateness** results in a 44% increase in the odds that the participant will allow the scenario ($p < 0.001$).
- Each 1-point increase in **Risk** results in a 14% decrease in the odds that the participant will allow the scenario ($p < 0.001$).
- **Expectedness** had no significant influence on the participant's decision ($p = 0.201$).

The strongly significant relationship between attitudes and behavior is interesting in light of the “privacy paradox” [34], an attitude-behavior gap that has been studied extensively by privacy researchers. Arguably, the privacy paradox is an artifact of the fact that *general* privacy concerns (which are commonly high) do not match *specific* behaviors (which subsequently ignore these general concerns). Since in our study attitudes and behaviors are measured at the same contextual level, their relationship is much stronger than in other studies. This may explain why we do not find an attitude-behavior gap.

4.4 Mediation Analysis

With tests 1 and 2 of our mediation analysis confirmed, we conduct test 3 by adding the scenario parameters to the *g*lmer of participants' decisions on their attitudes. The bottom half of Table 8

Table 6. Effect of Scenario Parameters on Usefulness

Model	<i>df</i>	<i>Chi.Sq.</i>	<i>p</i> -value
<i>usefulness ~ (1 sid)</i>	3		
+what	5	939.91	<.0001
+storage	12	457.36	<.0001
+purpose	23	401.18	<.0001
+action	26	328.88	<.0001
+who	29	117.57	<.0001
<i>Interactions</i>			
+what:who	106	214.48	<.0001
+who:purpose	50	184.48	<.0001
+what:purpose	62	85.39	<.0001

Table 7. Effect of Scenario Parameters on Expectedness

Model	<i>df</i>	<i>Chi.Sq.</i>	<i>p</i> -value
<i>expectedness ~ (1 sid)</i>	3		
+storage	5	841.24	<.0001
+who	16	425.92	<.0001
+what	19	422.31	<.0001
+purpose	22	231.98	<.0001
+action	29	29.45	<.0001
<i>Interactions</i>			
+what:who	106	262.80	<.0001
+who:purpose	50	138.73	<.0001
+what:purpose	62	84.89	<.0001

Table 8. Effect of Attitudes on Decision; Conditional Effects of Parameters Are Added Subsequently

Model	χ^2	<i>df</i>	<i>p</i> -value
<i>decision ~ (1 sid)</i>			
+Comfort	7934.72	1	<.0001
+Usefulness	1249.51	1	<.0001
+Appropriateness	149.15	1	<.0001
+Risk	10.90	1	.0009
+Expectedness	1.62	1	.201
<i>Adding Scenario Parameters</i>			
+action	0.332	3	0.953
+what	13.871	11	0.2401
+purpose	3.60	3	0.3069
+storage	14.57	2	0.0006
+who	24.53	7	0.0009

Fig. 2. Final mediation model.

shows these conditional effects of the significant scenario parameters on participants' allow/reject decision, controlling for their attitudes. **Action**, **What**, and **Purpose** are no longer significant in this model, suggesting that these effects are *fully mediated* by participants' attitudes. **Storage** and **Who** are still significant, but their conditional effects are smaller than their marginal effects on decision (without controlling for attitude, see Table 2). Their χ^2 s are reduced drastically by 98% and 87%, respectively. Overall, there was a substantial mediation effect. Figure 2 shows the final model mediation model.

4.5 Post-Hoc Results

To understand the effects of different values of each parameter on participants' various attitudes and their allow/reject decision, we conducted post-hoc tests using Tukey's method to adjust p -values to account for family-wise error. This subsection highlights the key insights from these tests. For an overview of the differences between parameter values, the reader is invited to visually inspect them by referring to Figures 24–28 in the Appendix.

Storage. Participants perceive more risk (d range = [0.568, 1.707], all $ps < .001$), are less comfortable (d range = [0.538, 1.741], all $ps < .001$) and find it inappropriate (d range = [0.436, 1.550], all $ps < .001$) when their information is shared to 'third parties' or 'stored on a remote server' as compared to when it is stored 'locally'. Participants also found it less useful to share their information with third parties as compared to storing it locally or on a remote server (d range = [0.28, 1.02], $p < .001$). Interestingly, participants expected it less that the information is stored locally rather than stored on remote server or shared to third parties (d range = [0.212, 0.894], all $ps < .001$). Finally, the odds of enabling a feature when information is stored locally were 1.96 times higher than when information is stored on a remote server ($p < .001$) and 8.36 times higher than when information is shared with third parties ($p < .001$).

Action. Participants were less comfortable (d range = [0.158, 0.348], all $ps < .001$) and found it more risky (d range = [0.145, 0.262], all $ps < .001$) when their information is used to give them

recommendations instead of optimizing services or giving them insight into their behavior. Sharing information was also found less useful ($d = 0.293, p < .001$) and less appropriate ($d = 0.256, p < .001$) for recommendation purposes as opposed to when the scenario did not specify any purpose. Participants also expected it less ($d = 0.123, p = .0021$) when their information was being shared for recommendation purposes as opposed to when the scenario did not specify any purpose. Finally, the odds of enabling a feature for recommendation purposes were 1.53 times lower as opposed to when the scenario did not specify any purpose ($p < .001$). Additionally, the odds of enabling a feature for optimization purposes were 1.65 times higher than for recommendation purposes ($p < .001$) and 1.26 times higher than for giving behavioral insights ($p < .001$).

Purpose. Participants found it inappropriate (d range = [0.343, 0.411], all $ps < .001$) when information is collected for the purpose of detecting their presence in the house as compared to the purposes of automating operations or giving timely alerts, and it was even more inappropriate to collect information for the purpose of detecting their location in the house (d range = [0.163, 0.574], all $ps < .001$). Participants also found it more risky when information is used for location detection as compared to presence detection ($d = 0.598, p < .001$), but they found it less risky to share information for the purpose of giving timely alerts or for automating operations (d range = [0.550, 0.601], p range = [0.002, 0.004]). Participants also found it more useful when information is collected for the purpose of providing alerts ($d = 0.558, p < .001$) or for automating operations ($d = 0.603, p < .001$) compared to the purpose of detecting their location in the house. Finally, the odds of enabling a feature were 1.29 times higher for detecting their presence in house than for detecting their location ($p = 0.0002$). Moreover, the odds of enabling a feature for the purpose of giving timely alerts and automating operations were 1.59 ($p < .001$) and 1.65 ($p < .001$) times higher respectively.

Who. Participants expected it more that their smart security systems will access information as compared to other devices such as their smart HVAC, TV, alarm, and washing machine (d range = [0.267, 0.618], all $ps < .001$). Users perceived data access by their security systems as more useful compared other devices like their smart refrigerator, washing machine, TV and HVAC (d range = [0.386, 0.627], all $ps < .001$). Participants were more comfortable ($d = 0.196, p = .002$) and found it less risky ($d = 0.263, p < .001$) for their security systems to access collected data as compared to their smart lighting systems. Also, participants were more comfortable (d range = [0.173, 0.356], all $ps < .05$) and found it less risky (d range = [0.256, 0.338], all $ps < .05$) for their lighting systems to access collected data compared to their smart assistant, TV and alarm clock. Finally, the odds of users enabling access to their smart security system were higher than to their smart refrigerator and washing machine by 1.8 times ($p < .001$), their smart TV by 1.7 times ($p < .001$) and their smart alarm clock by 1.6 times ($p < .001$). We found similar results for participants' smart assistant which had odds higher than their smart TV (1.76 times higher, $p < .001$), their smart alarm clock (1.68 times higher, $p < .001$), their smart washing machine (1.90 times higher, $p < .001$) and their smart refrigerator (1.85 times higher, $p < .001$).

What. This parameter had twelve different values and there were numerous combinations that were significant when we checked the post-hoc effects. We limit our discussion to the differences between the 'Smart Assistant' and the other values of this parameter, because these specific differences are consistently significant. The reader is invited to inspect Figure 27 in the Appendix for other differences. Participants found it more appropriate (d range = [0.213, 0.756], all $ps < .001$) and useful (d range = [0.365, 0.683], all $ps < .01$) when information collected by their smart assistant was being accessed as compared to other devices like cameras or microphones. The participants also found it less risky (d range = [0.385, 0.759], all $ps < .05$) and were more comfortable (d range = [0.430, 0.821], all $ps < .01$) to grant access to information collected by their smart

assistant than their camera or microphone. Participants also expected more to share information collected by their smart assistant as compared to other devices such as cameras ($d = 0.62, p < .01$), microphones ($d = 0.39, p < .01$), or their smart alarm clock ($d = 0.21, p = .027$). The odds of giving access to information collected by their smart assistant were higher than for cameras by 2.7 times ($p < .001$), microphones by 1.8 times ($p < .001$), their Smart TV by 1.15 times ($p < .001$), and their smart washing machine by 1.8 times ($p < .001$).

4.6 Defaults and Framing

As outlined in Section 3.2, the framing and default of the Decision question in our study were manipulated between-subjects at three levels each: positive, negative, or neutral framing; combined with a positive, negative, or no default. We analyze the effects of defaults and framing in a separate article [4]. In short, the analysis shows that defaults and framing have direct effects on disclosure: Participants in the negative default condition are less likely to enable the functionality, while participants in the positive default condition are more likely to enable the scenario (a traditional default effect). Likewise, participants in the negative framing condition are more likely to enable the functionality (a loss aversion effect).

Moreover, there are interaction effects between defaults/framing and attitudes on disclosure: the effects of attitudes are generally weaker in the positive and negative default conditions than in the no default condition, and they are also weaker in the negative framing condition.

Importantly, there are no interaction effects between defaults/framing and parameters on attitude or disclosure. Hence, the main findings in this section regarding the structure and relative importance of the effects of parameters remain the same, regardless of the effects of defaults and framing. For a more thorough discussion of the effects of defaults and framing, we refer the reader to [4].

4.7 Discussion

We split this section in two parts: First, we discuss the consequences of our analyses—and especially our post hoc test results—for the development and adoption of household IoT devices. Second, we discuss how our results can inform the design of household IoT privacy-setting interfaces.

4.7.1 Consequences for the Development and Adoption of Household IoT Devices. In the introduction we mentioned that privacy risk is an increasingly important barrier to the adoption of household IoT devices. Interestingly, though, in our study, Comfort, Usefulness, and Appropriateness had a stronger effect on users' allow/reject decisions than Risk. This suggests that IoT devices with a trust-inspiring design, a strong value proposition, and a clear explanation of the appropriateness of their data collection practices can overcome initial perceptions of privacy risk.

The tradeoff between Comfort, Usefulness, and Appropriateness embodies an interesting tradeoff: Usefulness is associated with the utility of a feature, whereas Appropriateness is a contextual evaluation (Is this acceptable, given the *situation*?) and Comfort is a self-relevant evaluation (Is this acceptable for *me*?).

The interaction between **What**, **Who**, and **Purpose** also suggests that users make context-relevant evaluations: scenarios are not accepted based on the sum of their components; rather, certain combinations of devices and purposes are more acceptable than others. While this is outside the scope of the current paper, future work could look into this context-dependency to find specific synergistic combinations.

The **Storage** parameter had the most significant influence on the participants' decision and all attitudes, but most prominently on Risk. ($\chi^2 = 47240.72, p < .001$). This indicates that users' risk perceptions are mostly dependent on the way household IoT systems store and share their data.

Household IoT device manufacturers who want to reduce risk perceptions may want to opt for storing all data locally instead of on a remote server (something users are actually more likely to expect).

Finally, the **Action** parameter had the least significant influence. Arguably, once users allow information to be collected, they care less about how exactly it is being used (or possibly, they do not expect to be able to control how it is being used).

4.7.2 Designing for IoT by Prioritizing Parameters. The results of our analyses uncover an intuitive reality about our household IoT scenarios, namely, they consist of two somewhat separate parts: On the one hand, there is a device (**Who**) that accesses information collected by another device (**What**), for a purpose certain **Purpose**. At the same time, this collected information may be stored somewhere (**Storage**) and some **Action** may be performed on it.

For the first aspect, we observed substantial interaction effects between all the three parameters, indicating that users want to make intricate decisions about what information is going where and for what purposes. Specifically, Unlike Bahirat et al. [5], we cannot use an interface with a separate ‘layer’ for each parameter; the interaction effects suggest that when users decide on one parameter, they inherently take another parameter into account. Therefore the settings interface for device/sensor management should show all three parameters at the same time to allow users to make these decisions.

For the second aspect, data storage had a very strong impact, while the action had the weakest impact. Additionally, there were no interactions between these two parameters, nor did they interact with any of the other parameters. This suggests that data storage and use can be separated in our privacy-setting interface.

5 PRIVACY-SETTING PROTOTYPE DESIGN

Designers of household IoT privacy-setting interfaces face a difficult challenge. Since there currently exists no centralized system for setting one’s household IoT privacy settings, designers must rely on existing data (cf. [28]) or self-collected survey results (cf. this paper) to inform the design of these interfaces. Moreover, these privacy-setting interfaces will likely be complex, as they require users to navigate settings for the collection of various types of data for multiple purposes across various devices.

Our dataset presents a simplified version of possible scenarios one might encounter in routine usage of smart home technology. Still it is a daunting task to design an interface, even for these simplified scenarios: We want to enable users to navigate their information collection and sharing preferences across 12 different sources (*What*), 7 different devices trying to access this information *Who* for 4 different *Purposes*. Additionally, this information is being stored/shared in three ways (*Storage*) and being used for four different longer-term *Actions*.

In this section, we present our prototype, which is based on the observations made from our statistical analysis. Section 7 extends this prototype to cover findings from our machine learning analysis to create default privacy profiles, but before we do so, we first want to design an intuitive interface that gives users manual control over their privacy settings. This interface should be able to present a vast amount of settings information in a concise and understandable manner, and allow some users to set their settings with little effort at a coarse level while still allowing others to spend the effort to micro-manage their privacy settings in more detail.

Our statistical analysis (see Section 4) reveals what the most significant parameters are in our dataset, as well as which parameters interact with each other. The results show that the *Storage* parameter had by far the strongest effect on participants’ decision to enable/disable the smart home feature described in the scenario. After *Storage*, *Who*, *What*, and *Purpose* had similar-sized

effects. Moreover, we found fairly strong significant two-way interaction effects between these parameters. Finally, the *Action* parameter had a weak but still significant effect.

Based on these results, we decided to split our settings interface into two separate sections: ‘Device/Sensor Management’ and ‘Data Storage & Use’. The landing page of our design (Screen 1 in Figure 3) gives users access to these two sections. The former section is based on *Who*, *What*, and *Purpose* and allows users to “Manage device access to data collected in your home” (Screen 2-3). The latter section is based on *Storage* and *Action*, and allows users to “Manage the storage and long-term use of data collected in your home” (Screen 4). Both sections are explained in more detail below.

Device/Sensor Management. This screen (Figure 3, screen 2) allows users to control the *Purposes* for which each device (*Who*) is allowed to access data collected by itself, other devices, and the smart home sensors installed around the house (*What*). This screen has a collapsible list of data-collecting devices and sensors (*What*). For each device/sensor, the user can choose what devices can access the collected data (*Who*; in rows), and what it may use that data for (*Purpose*; in columns).

In the example of Figure 3, the user does not give the ‘Refrigerator’ access to information collected by the ‘Smart Assistant’ for any of the four purposes, while they give the ‘Smart TV’ access to this data for the purpose of giving ‘timely alerts’. In this example the ‘Smart Assistant’ is allowed to use its own data to ‘automate operations’ and to ‘know your location in your home’.

Showing *Who*, *What*, and *Purpose* at the same time allows users to enable/disable specific combinations of settings—the significant interaction effects between these parameters suggest that this is a necessity. The icons for the *Purpose* requirement allow this settings grid to fit on a smartphone or in-home control panel. We expect that users will quickly learn the meaning of these icons, but they can always click on ‘I want to know more’ to learn their meaning (see Figure 3, screen 3).

Data Storage & Use. This screen (Figure 3, screen 4) allows users to control how their data is stored and shared (*Storage*), as well as how stored data is used (*Action*). These settings are independent from each other and from the Device/Sensor Management settings.

For ‘Storage & Sharing’, users can choose to turn storage off altogether, store data locally, store data both locally and on a remote server, or store data locally and on a remote server *and* allow the app to share the data with third parties. Note that the options for *Storage* are presented as ordered, mutually exclusive settings. Our scenarios did not present them as such (i.e., participants were free to reject local storage but allow remote storage). However, the *Storage* parameter showed a very clear separation of levels (see Figure 24 in the Appendix), so this presentation is justified. For ‘Data Use’, the users can choose to enable/disable the use of the collected data for various secondary purposes: behavioral insights, recommendations, service optimization, and/or other purposes.

In the subsequent sections, we describe the results from our machine learning analysis and further explain how these results impact the designs presented in this section. For this purpose, Section 7 revisits the interface designs presented here.

6 PREDICTING USERS’ BEHAVIORS

In this section, we predict participants’ *enable/disable* decision using machine-learning methods. Our goal is to find suitable default settings for our IoT privacy-setting interface. Consequently, we do not attempt to find the best possible solution; instead we make a conscious tradeoff between parsimony and prediction accuracy. Accuracy is important to ensure that users’ privacy preferences are accurately captured and/or need only few manual adjustments. Parsimony, on the other hand, prevents overfitting and promotes fairness: we noticed that more complex models tended to increase overall accuracy by predicting a few users’ preferences more accurately, with no effect on other users. Parsimony also makes the associated default setting easier to understand for the user.

Fig. 3. Screen 1 (top left) is the landing page of our manual settings interface, screen 2 (top right) is the Device/Sensor Management page, screen 3 (bottom left) shows the explanation when you click on “I want to learn more”, and screen 4 (bottom right) is the Data Storage & Use page.

Table 9. Comparison of Clustering Approaches (Highest Parsimony and Highest Accuracy)

Approach	Initial clusters	Final # of profiles	Complexity (avg. tree size/profile)	Accuracy
Naive (enable all)	1	1	1	46.74%
Naive (disable all)	1	1	1	53.26%
One Rule (Figure 4)	1	1	3	61.39%
Overall (Figure 7)	1	1	8	63.32%
	1	1	264	63.76%
Attitude-based clustering (Figure 9)	2	2	2	69.44%
	2	2	121.5	72.66%
	3	3	2.67	72.19%
	3	3	26.67	73.47%
	5	4	3	72.61%
	5	4	26	73.56%
Agglomerative clustering (Figure 12)	1133	4	2	79.4%
	1133	5	2.4	80.35%
	1133	6	3.17	80.60%
Fit-based clustering (Figure 16)	2	2	2	74.43%
	2	2	151.5	76.72%
	3	3	7	79.80%
	3	3	65.33	80.81%
	4	4	9.25	81.88%
	4	4	58.25	82.41%
	5	5	4.2	82.92%
	5	5	51.4	83.35%

Our prediction target is the participants’ decision to *enable* or *disable* the data collection described in each scenario. The scenario parameters serve as input attributes. These are nominal variables, making decision tree algorithms such as ID3 and J48 a suitable prediction approach. Unlike ID3, J48 uses gain ratio as the root node selection metric, which is not biased towards input attributes with many values. Moreover, by using J48 decision trees, the amount of pruning for the model can be easily manipulated to investigate the tradeoff between the accuracy and parsimony. We therefore use J48 throughout our analysis.

Using Java and Weka’s Java library [58] for modeling and evaluation, we implement progressively sophisticated methods for predicting participants’ decisions. After discussing naive (enable/disable all) solutions and One Rule Prediction, we first present a cross-validated tree learning solution that results in a single “smart default” setting that is the same for everyone. Subsequently, we discuss three different procedures that create a number of “smart profiles” by clustering the participants and creating a separate cross-validated tree for each cluster. For each procedure, we try various numbers of clusters and pruning parameters. The solutions with the most parsimonious trees and the highest accuracies of each approach are reported in Table 9; more detailed results of the parsimony/accuracy tradeoff are presented in Figures 7, 9, 12, and 16 throughout the article, and combined in Figure 20.

6.1 Naive Prediction Model

We start with the naive or “information-less” predictions. Compared to our previous work [5], our current dataset shows that it is even less amenable to a ‘simple’ default setting: it contains

Fig. 4. A “smart default” setting based on the “One Rule” algorithm (4 nodes, accuracy: 61.39%). Parameter value abbreviations correspond to the “code” column in Table 1.

Table 10. Confusion Matrix for the One Rule Prediction

Observed	Prediction		Total
	Enable	Disable	
Enable	5,085 (TP)	1,270 (FN)	6,355
Disable	3,262 (FP)	3,979 (TN)	7,241
Total	7,192	6,404	13,596

6,335 *enable* cases and 7,241 *disable* cases, which means that predicting *enable* for every setting gives us a 46.74% prediction accuracy, while making a *disable* prediction for every setting gives us an accuracy of 53.26%. In other words, if we disable all information collection by default, only 53.26% users will on average be satisfied with this default settings. Moreover, such a default setting disallows any ‘smart home’ functionality by default—arguably not a solution the producers of smart appliances can get behind.

6.2 One Rule Prediction

Next, we use a “One Rule” (OneR) algorithm to predict users’ decision using the simplest prediction model possible. OneR is a very simple but often surprisingly effective learning algorithm [18]. It creates a frequency table for each predictor against the target, and then find the best predictor with the smallest total error based on the frequencies.

As shown in Figure 4, the OneR model predicts users’ decision solely based on the **Storage** parameter with an accuracy of 61.39%. Based on this model, if we enable all information-sharing *except* with third parties, we will on average satisfy 61.39% of users’ preferences—a 15.3% improvement³ over the naive “disable all” default. Note, though, that this default setting is overly permissive, with 3,262 false positive predictions (see Table 10).

6.3 Overall Prediction

Moving beyond a single parameter, we create a “smart default” setting by predicting the *enable/disable* decision with all scenario parameters using the J48 decision tree algorithm. The resulting tree has an accuracy of 63.76%. As shown in Figure 5, this model predicts users’ decision on **Storage** first. It predicts *disable* for every scenarios with collected data stored on a remote server

³61.39 / 53.26 = 1.153.

Fig. 5. A “smart default” setting with 264 nodes (accuracy: 63.76%). Parameter value abbreviations correspond to the “code” column in Table 1.

Table 11. Confusion Matrix for the Overall Prediction

Observed	Prediction		Total
	Enable	Disable	
Enable	4,753 (TP)	2,488 (FN)	7,241
Disable	2,439 (FP)	3,916 (TN)	6,355
Total	7,192	6,404	13,596

and shared with third party. For scenarios that store collected data on remote server without sharing, the default settings will depend on the ‘purpose’ of information sharing. There is a further drill down based on ‘who’ and ‘what’. For scenarios that store collected data locally, the default settings will depend on the ‘what’. There is a further drill down based on ‘who’, ‘what’, and ‘action’. With this default setting, users would on average be satisfied with 63.76% of these settings—a 19.7% improvement over the naive “disable all” default.

On the downside, this “smart default” setting is quite complex—the “smart default” in our previous work [5] contained only 49 nodes, whereas the “smart default” for our current dataset has 264 nodes. Compared to *One Rule* algorithm, which only has four nodes in its decision tree and is thus much easier to explain, the accuracy improvement of Smart Default is only 3.8%. This highlights the tradeoff between parsimony and prediction accuracy that we have to make when developing “smart default” settings. On the upside, though, the prediction of the J48 decision tree algorithm is more balanced, with a roughly equal number of false positives and false negatives (see Table 11).

To better understand the parsimony/accuracy tradeoff, we vary the degree of model pruning to investigate the effect of increasing the parsimony (i.e., more trimming) on the accuracy of the resulting “smart default” setting. The parameter used to alter the amount of post-pruning performed on the J48 decision trees is called Confidence Factor (*CF*) in Weka, and lowering the Confidence Factor will incur more pruning. We tested the J48 classifier with a Confidence Factor ranging from 0.01 to 0.25 (the default setting in Weka) with an increments of 0.01.

Figure 6 displays the accuracy and the size of the decision tree as a function of the Confidence Factor. The X-axis represents the Confidence Factor; the left Y-axis and the orange line represent the accuracy of the smart default setting; the right Y-axis and the dotted blue line represent the size

Fig. 6. Accuracy and parsimony (tree size) of the smart default change as a function of Confidence Factor.

of the decision tree for that setting. The highest accuracy, 63.75%, is achieved with the 264-node decision tree produced by $CF = 0.25$. The lowest accuracy, 62.9%, is achieved with the 44-node decision tree produced by $CF = 0.19$. When $CF \leq 0.16$, the decision tree contains only 8 nodes. The 8-node profile with the highest accuracy is produced by $CF = 0.10$ with an accuracy of 63.32%.

Figure 7 summarizes accuracy as a function of parsimony. The X-axis represents the number number of nodes in the decision tree (more = lower parsimony); the Y-axis represents the accuracy of the decision tree. The figure shows the most accurate J48 solution for any given tree size, and includes the One Rule and Naive predictions for comparison. Reducing the tree from 264 to 8 nodes incurs a negligible 0.67% reduction in accuracy. This decision tree is shown in Figure 8, and is still 3.1% better than the One Rule prediction model and 18.9% better than the naive “disable all” default. This more parsimonious “smart default” setting can easily be explained to users as follows:

- All sharing with third parties will be disabled by default.
- Remote storage is allowed for automation and alerts, but not for detecting your presence or location in the house.
- Local storage is allowed for all purposes.

While the “smart default” setting makes a considerable improvement over a naive default, there is still a lot of room for improvement—even our best prediction model only correctly models on average 63.76% of the user’s desired settings. This should come at no surprise, as one of the most consistent findings in the field of privacy is that people differ substantially in their privacy preferences [25]. As a result, our “one-size-fits-all” default setting—smart as it may be—is not very accurate. Recent work in the field of privacy suggest to *tailor* the privacy settings to the user to accommodate for these interpersonal differences [23]. Our previous work therefore moved beyond “smart default” settings by clustering participants with similar privacy preferences and creating a

Fig. 7. Parsimony/accuracy comparison for Naive, One Rule, and Overall Prediction.

Fig. 8. A “smart default” setting with only 8 nodes (accuracy: 63.32%). Parameter value abbreviations correspond to the “code” column in Table 1.

set of “smart profiles” covering each of the clusters [5]. The idea is that the accuracy of the tree for each cluster will likely exceed the accuracy of our overall prediction model.

In the remainder of this section, we apply existing and new clustering methods with the aim of creating separate “smart profiles” for each cluster. As our goal is to develop simple, understandable profiles, we keep the parsimony/accuracy tradeoff in mind during this process.

6.4 Attitude-Based Clustering

As shown in Figure 1, our statistical results indicate that the effects of scenario parameters on users’ decisions are mediated by their attitudes (Risk, Comfort, Appropriateness, Expectedness and Usefulness). Therefore, our first attempt to develop “smart profiles” is to cluster participants

Fig. 9. Parsimony/accuracy comparison for attitude-based clustering.

with similar attitudes towards the 12 scenarios they evaluated. We averaged the values per attitude across each participant's 12 answers, and ran a *k-means* clustering algorithm to divide them into 2, 3, 4, 5, and 6 clusters. We then added the participants' cluster assignments back to our original dataset, and ran the J48 decision tree algorithm on the dataset with this additional *Cluster* attribute for each number of clusters, varying the Confidence Factor from 0.01 to 0.25 with increments of 0.01. The results are summarized in Figure 9, which displays the most accurate solution for any given tree size and number of clusters.

All of the resulting decision trees have *Cluster* as the root node. This justifies our approach, because it indicates that the *Cluster* parameter is a very effective for predicting users' decisions. It also allows us to split the decision trees at the root node, and create different "smart profile" for each subtree/cluster. Note that for some solutions two clusters end up with the same decision tree, which effectively reduces the number of profiles by 1.

For the 2-cluster solutions (the blue line in Figure 9), the highest accuracy is 72.66%, which is a 14.0% improvement over the best single "smart default" setting. However, this tree has an average of 121.5 nodes per profile. In comparison, the most parsimonious solution has only one node ("disable all") for one of the clusters, and three nodes ("disable sharing with third parties") for the other cluster (see Figure 10). This solution still has an accuracy of 69.44%, which is still an 8.9% increase over the best single "smart default" setting.

For the 3-cluster solutions (the orange line in Figure 9), the highest accuracy of 73.47% is achieved by a set of trees with 26.67 nodes on average (a minimal improvement of 1.1% over the best 2-cluster solution, but with simpler trees), while the most parsimonious solution has a "disable all" and an "enable all" tree, plus a tree that is the same as the most parsimonious smart default setting (see Figure 8). This solution has an accuracy of 72.19%, which is a 4.0% increase over the most parsimonious 2-cluster solution.

The 4-cluster solutions (the grey line in Figure 9) all result in "over-clustering": All solutions based on the 4-cluster *Cluster* parameter result in two profiles with the same subtree, effectively

Fig. 10. The most parsimonious 2-profile attitude-based solution (2 nodes/profile, accuracy: 69.44%). Parameter value abbreviations correspond to the “code” column in Table 1.

Fig. 11. A 3-profile solution example of attitude-based clustering (18.33 nodes/profile, accuracy: 73.26%). Parameter value abbreviations correspond to the “code” column in Table 1.

resulting in a 3-profile solution. The accuracy of these solutions is actually lower than the accuracy of similar 3-cluster solutions, so we will not discuss them here.

The 5-cluster solutions (the yellow line in Figure 9) are also “over-clustered”, resulting in 4 profiles. The highest accuracy of 73.56% is achieved by a set of trees with 26 nodes—this is about the same accuracy and parsimony as the most accurate 3-cluster solution. The same holds for the most parsimonious 5-cluster solution, which has a similar accuracy and parsimony as the most parsimonious 3-cluster solution.

The accuracy of the 6-cluster solutions (which result in either 4- or 5-profile solutions) is lower than the accuracy of similar 5-cluster solutions. Therefore, we will not discuss these results further.

Reflecting upon the attitude-based clustering results, we observe in Figure 9 that there is indeed a tradeoff between accuracy and parsimony: The most parsimonious results are less accurate, but the most accurate results are more complex. Moreover, the 2-profile solutions are about 5% less accurate than the 3-profile solutions at any level of complexity. The 4-profile solutions do not improve the solution much further, though.

The 3-profile solution with an average of 18.33 nodes per profile and 73.26% accuracy provides a nice compromise between accuracy and parsimony. Part of this decision tree is shown in Figure 11: it contains one “disable all” profile, one “enable all” profile, and a more complex profile with 55 nodes that disallows sharing with third parties and allows remote and local storage depending on the purpose (not further shown).

Fig. 12. Parsimony/accuracy comparison for agglomerative clustering.

6.5 Agglomerative Clustering

The attitude-based clustering approach requires knowledge of users' attitudes towards the household IoT information-sharing scenarios, which may not always be available. We developed an alternative method for finding "smart profiles" that follows a hierarchical bottom-up (or agglomerative) approach, using users' decisions only. This method first fits a separate decision tree for each participant, and then iteratively merges these trees based on similarity. In our previous work [5], only 10 out of the 200 users in the dataset had unique trees fitted to them (all others had an "enable all" or "disable all" tree), making the merging of trees a rather trivial affair. Our current dataset has many more participants, and is more complex, making the agglomerative clustering approach more challenging but also more meaningful.

In the first step, 283 participants' decision trees predict "enable all", 414 participants' decision trees predict "disable all", while the remaining 436 participants have a multi-node decision tree.

In the second step, a new decision tree is generated for each possible pair of participants in the "multi-node group". The accuracy of the new tree is compared against the weighted average of the accuracies of the original trees. The pair with smallest reduction in accuracy is merged, leaving 435 clusters for the next round of merging. If two or more candidate pairs have the same smallest reduction in accuracy, priority is given to the pair with the most parsimonious resulting tree (i.e., with smallest number of nodes). If there are still multiple pairs that tie on this criterion, the first pair is picked. The second step is repeated until it reaches the predefined number of clusters, and the entire procedure is repeated with 20 random starts to avoid local optima.

To fit the trees, we use the J48 classifier with a Confidence Factor ranging from 0.01 to 0.25 with increments of 0.01. Surprisingly, smaller tree sizes result in a *higher* accuracy for agglomerative clustering (see Figure 12). This suggests that without extensive trimming, our agglomerative approach arguably overfits the data, resulting in a lower level of cross-validated accuracy.

Fig. 13. The best 4-profile agglomerative clustering solution (2 nodes/profile, accuracy: 79.40%). Parameter value abbreviations correspond to the “code” column in Table 1.

Fig. 14. The best 5-profile agglomerative clustering solution (2.4 nodes/profile, Accuracy: 80.35%). Parameter value abbreviations correspond to the “code” column in Table 1.

The best 4-cluster solution has an average of 2 nodes per profile and an accuracy of 79.40%—a 24.53% improvement over the “smart default”, and a 7.9% increase over the most accurate 5-cluster/4-profile attitude-based clustering solution. The decision trees are shown in Figure 13: Aside from the “enable all” and “disable all” profiles, there is a “disable sharing with third parties” profile and a “local storage only” profile.

The best 5-cluster solution has an average of 2.4 nodes per profile and an accuracy of 80.35%—a 26.02% improvement over the “smart default”, but only a 1.2% improvement over the 4-cluster agglomerative solution. The decision trees are shown in Figure 14: It has the same profiles as the 4-cluster solution, plus an “allow automation and alerts, but don’t track my presence or location in the house” profile.

Finally, the best 6-cluster solution⁴ has an average of 3.17 nodes per profile and an accuracy of 80.68%—a 26.54% improvement over the “smart default”, but no substantial improvement over

⁴There is another solution with slightly fewer nodes per profile (2.67) and a slightly lower accuracy (80.60%).

Fig. 15. The best 6-profile agglomerative clustering solution (3.17 nodes/profile, Accuracy: 80.68%). Parameter value abbreviations correspond to the “code” column in Table 1.

the 5-cluster agglomerative solution. The decision trees are shown in Figure 15: It has the same profiles as the 5-cluster solution, plus a profile that allows local storage for anything, plus remote storage for any reason except for user profiling (i.e., to recommend other services or to give the user insight into their behavior).

6.6 Fit-Based Clustering

We now present a “fit-based” clustering approach that, like the agglomerative approach, clusters participants without using any additional information. Instead, it uses the fit of the tree models to bootstrap the process of sorting participants into different clusters. The steps of our algorithm are as follows:

- *Random Starts:* We randomly divide participants into k separate groups, and learn a tree for each group. This is repeated until a non-trivial starting solution (i.e., with distinctly different trees per group) is found.
- *Iterative Improvements:* Once each of the k groups has a unique decision tree, we test for each participant which of the k trees best represents their 12 decisions. If this is the tree of a different group, we switch the participant to this group. Once all participants are evaluated and put in the group of their best-fitting tree, the tree in each group is re-learned with the data of the new group members. This then prompts another round of evaluations, and this process continues until no further switches are performed.
- *Repeat:* Since this process is influenced by random chance, it is repeated 1,000 times in its entirety to find the optimal solution. Cross-validation is performed in the final step to prevent over-fitting.

We perform this approach to obtain 2-, 3-, 4-, and 5-cluster solutions. To fit the trees, we use the J48 classifier with a Confidence Factor ranging from 0.01 to 0.25 with increments of 0.01. The best results are summarized in Figure 16.

For the 2-cluster solutions (the blue line in Figure 16), the highest accuracy is 76.72%—a 20.33% improvement over the “smart default” setting and a 5.6% improvement over the most accurate 2-cluster attitude-based solution. However, this tree has an average of 151.5 nodes per profile. The

Fig. 16. Parsimony/accuracy comparison for fit-based clustering.

Fig. 17. The most parsimonious 3-profile fit-based solution (7 nodes/profile, accuracy: 79.80%). Parameter value abbreviations correspond to the “code” column in Table 1.

most parsimonious solution is exactly the same as the most parsimonious 2-cluster attitude-based solution (see Figure 10), but with a higher accuracy (74.43%).

For the 3-cluster solutions (the orange line in Figure 16), the highest accuracy of 80.81% is achieved by a set of trees with 65.33 nodes on average. This is a 26.74% improvement over the “smart default”, a 10.0% improvement over the most accurate 3-cluster attitude-based solution (but at a cost of lower parsimony), and a 5.2% improvement over the best 2-cluster fit-based solution. The most parsimonious solution, on the other hand, has seven nodes on average, with an accuracy of 79.80%, thereby still outperforming all other 3-profile solutions. The decision trees for this solution are shown in Figure 17.

Fig. 18. The most parsimonious 4-profile fit-based solution (9.25 nodes/profile, accuracy: 81.88%). Parameter value abbreviations correspond to the “code” column in Table 1.

For the 4-cluster solutions (the grey line in Figure 16), the highest accuracy of 82.41% is achieved by a set of trees with 58.25 nodes on average. This is a 29.25% improvement over the “smart default”, a 3.8% improvement over the 4-cluster agglomerative solution (but at a cost of lower parsimony), and a 2.0% improvement over the best 3-cluster fit-based solution. The most parsimonious solution, on the other hand, has 9.25 nodes on average, with an accuracy of 81.88%. It still outperforms all other 4-profile solutions, but the agglomerative solution is more parsimonious. The decision trees for this solution are shown in Figure 18.

For the 5-cluster solutions (the yellow line in Figure 16), the highest accuracy of 83.35% is achieved by a set of trees with 51.4 nodes on average. This is a 30.05% improvement over the “smart default”, a 3.8% improvement over the 5-cluster agglomerative solution (but at a cost of lower parsimony), and a 1.1% improvement over the best 4-cluster fit-based solution. The most parsimonious solution, on the other hand, has 4.2 nodes on average, with an accuracy of 82.92%. It still outperforms the 5-profile agglomerative solution, but it is slightly less parsimonious. The decision trees for this solution are shown in Figure 19.

6.7 Discussion of Machine-Learning Results

Figure 20 shows a comparison of the presented approaches. The X-axis represents the parsimony (higher average tree size per profile = lower parsimony); the Y-axis represents the accuracy. While the “smart default” setting makes a significant 15.3% improvement over the naive default setting (“disable all”), we observe that having multiple “smart profiles” substantially increases the prediction accuracy even further. The fit-based clustering algorithm performs the best out of all the approaches, followed by agglomerative clustering and attitude-based clustering.

The most parsimonious 2-profile fit-based solution (with an accuracy of 74.43%) is the *simplest* of all “smart profile” solutions: one profile is simply “disable all”, while the other profile is the same as our OneR solution: “disable sharing with third parties”. In fact, these profiles are so simple, that one might not even want to bother with presenting them to the user: In our current interface (see Figure 3), these defaults are incredibly easy for users to implement by themselves.

The same is true for the 4-profile agglomerative clustering solution (see Figure 13) and the 5-profile agglomerative clustering solution (see Figure 14): these profiles involve little more than a single high-level setting, which users can likely easily make by themselves.

Fig. 19. The most parsimonious 5-profile fit-based solution (4.2 nodes/profile, accuracy: 82.92%). Parameter value abbreviations correspond to the “code” column in Table 1.

The 5-profile fit-based solution is the *most accurate* of all “smart profile” solutions. The most parsimonious 5-profile fit-based clustering solution (Figure 19) has an accuracy of 82.92%. It has the following five profiles:

- Enable all
- Enable local and remote storage, but disable third-party sharing
- Enable local storage only
- Enable local storage for everything except location-tracking, enable remote storage for everything except location- and presence-tracking, and disable third-party sharing
- Disable all

The fourth profile in this list specifies an interaction between between **Storage** and **Purpose**—something that is not possible in our current manual settings interface (which only allows interactions between **Who**, **What**, and **Purpose**). The next section will present a slightly altered interface that accommodates these profiles.

There is another 5-profile fit-based solution with a slightly higher accuracy (83.11%) and a reasonably simple tree (5 nodes/profile on average). This solution is shown in Figure 21. In this solution, the third profile (“enable local storage only”) is replaced by a slightly more complex profile (“enable local storage only, but not to recommend other services”). This profile specifies an additional interaction between **Storage** and **Action**. The next section will present a settings interface that accommodates this profile as well.

Other usable solutions are the 3-profile fit-based solution (Figure 17) or the 4-profile fit-based solution (Figure 18). However, like almost all of the less parsimonious solutions, these profiles involve higher-order interaction effects, e.g., between **Storage**, **Purpose**, and **Action** and between **Storage**, **Purpose**, and **Who**. Consequently, a rather more complex interface is needed to accommodate these default profiles.

7 PRIVACY-SETTING PROTOTYPE DESIGN USING MACHINE-LEARNING RESULTS

In Section 5, we developed a prototype interface that household IoT users can use to manually set their privacy settings (see Figure 3). Our machine-learning analysis (Section 6) resulted in a number of interesting solutions for “smart profiles” that would allow users of this interface to set their privacy settings with a single click (i.e., a choice of profile). While some of these profiles can

Fig. 20. Summary of all our approaches.

be integrated in our prototype (e.g., the most parsimonious 2-profile fit-based solution and the 4-profile and 5-profile agglomerative solutions) other profiles have an interaction effect between variables that are modeled as independent in our current prototype interface (e.g., the two 5-profile fit-based solutions presented in Figures 19 and 21).

In this section, we therefore present two modified prototypes that are designed to be compatible with these two 5-profile solutions. These two solutions are not the most accurate, but they produce a parsimonious set of profiles that require only minimal alterations to our interface design. They thus provide the optimal tradeoff between reduction accuracy, profile parsimony, and interface complexity.

7.1 Interface for the 5-Profile Fit-Based Solution with an Accuracy of 82.92%

This machine-learning solution (Figure 19) requires an interaction between the *Storage* parameter and the *Purpose* parameter—two parameters that are controlled independently in the prototype in Figure 3. Our solution is to slightly alter the interface, and add the profile selection page at the beginning of the interface (see Figure 22):

Fig. 21. A good 5-profile fit-based clustering solution (5 nodes/profile, Accuracy: 83.11%). Parameter value abbreviations correspond to the “code” column in Table 1.

- *Screen 1*: On this screen, users choose their most applicable default profile. For some users, the selected profile accurately represents their preferences, while others may want to adjust the individual settings manually.
- *Screen 2*: After clicking ‘Next’, users are given the option to select ‘Storage/Sharing & Device/Sensor Management’ or ‘Data Use’.
- When users select either ‘Storage/Sharing & Device/Sensor Management’ they first get to set their sharing preferences for ‘local storage’, ‘remote server’ and ‘third party sharing’ (*Storage*). Each of these can independently be set to *enabled* or *disabled*, but users can also click on ‘More’.
- *Screen 4*: When users select ‘More’, they can manage *Who-What-Purpose* combinations for that particular storage/sharing option.
- *Screen 5*: When users select ‘Data Use’ on screen 2, they get to enable/disable the use of the collected data for various secondary purposes (*Action*).

7.2 Interface for the 5-Profile Fit-Based Solution with an Accuracy of 83.11%

The alternative machine learning solution presented in Figure 21 requires an additional interaction between the *Storage* parameter and the *Action* parameter. This requires us to slightly alter the interface again (see Figure 23):

- *Screen 1*: The profile selection screen remains unchanged, with the exception that the ‘Local Storage Only’ profile is replaced by the more complex ‘Local Storage & No Recommendations’ profile.
- *Screen 2*: After clicking ‘Next’, users first get to set their sharing preferences for ‘local storage’, ‘remote server’, and ‘third party sharing’ (*Storage*). Each of these can independently be set to *enabled* or *disabled*, but users can also click on ‘More’.
- *Screen 3*: When users select ‘More’, they are given the option to select either ‘Device/Sensor Management’ or ‘Data Use’.
- *Screen 4*: When users select ‘Device/Sensor Management’, they can manage *Who-What-Purpose* combinations for that particular storage/sharing option.
- *Screen 5*: When users select ‘Data Use’, they get to enable/disable the use of the collected data for various secondary purposes (*Action*) for that particular storage/sharing option.

Fig. 22. Design for 5-Profile solution presented in Section 7.1. From top left, screen 1 is the profile selection page, screen 2 is the slightly altered landing page of our manual settings interface, screen 3 is the slightly altered Data Storage page, screen 4 (bottom left) is the Device/Sensor Management page, and screen 5 is the Data Use page.

7.3 Reflection on Design Complexity

The interfaces presented in this section have an additional ‘layer’, compared to the original interface presented in Section 5. This additional layer makes setting the privacy settings manually more difficult, but it is necessary to accommodate the complexity of the smart profiles uncovered by our machine learning analysis. On the one hand, this demonstrates the value of developing a parsimonious machine learning model—the more accurate but more complex profiles that

Fig. 23. Design for 5-Profile solution presented in Section 7.2. From top left, screen 1 is the profile selection page, screen 2 is the slightly altered Data Storage page, screen 3 follows the 'More' button to offer access to screen 4 (bottom left, the Data Use page) and screen 5 (bottom right, the Device/Sensor Management page).

comprise some of the solutions in Section 6 are not only more difficult to explain to the user, they also contain more complex interactions between decision parameters, forcing the manual settings interface to become even more complex. A simple smart profile solution avoids such complexity in the interface.

On the other hand, one should not over-simplify the profiles, lest they become overly generic and inaccurate in representing users' privacy preferences. Indeed, when we make our smart profile solutions more accurate, fewer users will need to make any manual adjustments at all, so we can allow some additional complexity in the interface.

8 LIMITATIONS AND FUTURE WORK

In this section, we discuss the limitations of our work, our plans to evaluate the presented interfaces, and other future work.

One limitation of our work concerns our ethical approach to privacy support. The decisions regarding the scenarios presented in our work encompass a tradeoff between privacy and utility. Our article adheres to a philosophy that considers neither privacy nor utility a morally superior goal: it is up to the end user to decide upon the balance between these two goals. We acknowledge that end users vary widely in how they evaluate this balance; hence, the goal of our work is to help each user find the optimal balance that is right for them. From this perspective, we believe that consistent and coherent privacy profiles are a better starting point for users than taking an “everything off” or “everything on” approach. We acknowledge, though, that other (more paternalistic) moral standpoints exist, and that our work does not adhere to these standpoints.

We also note that participants in our study made decisions about hypothetical rather than “real-life” scenarios. However, compared to most other privacy studies, our study asks participants about very specific IoT scenarios, measuring their attitudes and behaviors in the context of these scenarios. The hypothetical nature of the scenarios is thus a conscious tradeoff here: it is impossible to measure privacy in 4,000+ scenarios without presenting them on a screen.

Another limitation is that the effectiveness of our approach is contingent upon on the benevolence of the IoT infrastructure provider, and that it is easy for an “evil” provider to tweak the profiles in their favor (cf. [26]). To that point, we note that the provider of the IoT infrastructure may very well be a separate entity from the device vendors (similar to how smartphone OS providers are often separate from smartphone app developers). In that case, it is not in the infrastructure provider’s interest to “oversell” the IoT functionality, but rather to increase users’ comfort with the use of their platform.

A limitation regarding our machine-learning approach is that it assumes a perfect assignment of users to profiles. However, in our current approach, users of the profile-based interface make their own choice as to which profile they want to apply. If they do not make the correct choice, then this introduces additional uncertainty, and the accuracy of our approach will be substantially lower than described in our article. This limitation highlights the importance of the parsimony/accuracy tradeoff: Users benefit from parsimony in the context of our study, because parsimony makes for simpler profiles, which are easier to understand and hence easier to choose from. At the same time, though, these more parsimonious profiles are likely going to be less accurate, which means that users need to make more manual adjustments to the profile-based settings.

To further explore this tradeoff between parsimony and accuracy, and also to evaluate the usability of the proposed privacy-setting interface prototypes in this article, we are in the process of developing a user study to test these interface prototypes. In this study, we compare several default/profile solutions: all disabled by default, all enabled by default, two variants of a single smart default (Figure 4 and Figure 8), two variants of smart profiles (Figure 14 and Figure 23). The study will also consider different levels of complexity for the manual settings interface.

Aside from this user study, we also intend to expand our work in the direction of conducting a cross-platform study on IoT privacy. This study would aim at understanding how user perspectives about privacy differ across different IoT platforms, such as wearable IoT, household IoT, and public IoT.

9 CONCLUSION

The motivation behind our research was to reduce the information and choice overload associated with the plethora of choices that users might face while setting their privacy settings in a

household IoT environment. While our work is certainly not the first to propose a new interface for privacy settings (cf. [40] and [54]), our work is unique in its development of a novel *data-driven process* that leveraging an understanding of user decision making from a contextual perspective of a technology to aid the design of more efficient privacy-setting interfaces.

Specifically, whereas it is standard practice to develop privacy-setting interfaces based on user feedback from user testing on existing interfaces, we gather users' feedback on IoT scenarios *before* developing the interface, and employ statistical and machine learning analyses to develop a settings interface and a series of 'smart profiles' to aid users in their privacy-setting task. The use of scenarios solves the challenge of creating a privacy-setting interface for technologies that are still under development. Moreover, in using scenarios, the presented procedure avoids typical decision externalities such as default effects, framing effects [6], and decision-context effects [26] that tend to obfuscate users' behaviors in more naturalistic studies.

Beyond our previous work [5], the current article describes a more thorough method for analyzing potential solutions that span a variety of perspectives on the balance between the number and accuracy of shortcut solutions (profiles), the complexity of the manual settings interface, and users' ability to comprehend and control these systems. Our approach provides several options, but it is up to the designer to decide what the best option is. In some cases, this may involve implementing several potential solutions and comparing them in a user test or a controlled experiment.

We argue that the exploration of this balance is just as important for many other settings interfaces as well. Indeed, future work could apply the proposed procedure to other privacy-setting domains, such as healthcare privacy, drone privacy, and nanotech privacy. To do this, the researchers have to identify the parameters that are potentially relevant in the domain under investigation. "Who" (e.g., for healthcare privacy: primary-care physician, specialist, first-aid worker, health insurer, employer) and "what" (e.g., for healthcare privacy: demographics, physical data, chronic illnesses, medications, past illnesses, past surgeries, allergies) and "purpose" (e.g., for healthcare privacy: general care, emergency care, epidemic prevention, health research, insurance fraud prevention, billing) are par for the course, but their actual values will of course differ per domain. Moreover, each domain will have its own unique considerations (e.g., for healthcare privacy, who has agency over the data when the user is incapacitated).

Our proposed procedure uses "smart profiles" because IoT privacy decisions are often made upon installation of the IoT system, and rarely revisited. In contrast, in many of the potential application domains for our proposed procedure the privacy decisions are repeated and spread out in time (e.g., in healthcare privacy, privacy decisions have to be made upon each doctor visit). In those domains, fully "adaptive" privacy mechanisms that use "active tracking" (cf. [22], and [30]) are more suitable. Regardless, our approach is still beneficial in these situations, since such an active tracking system still needs a manual interface to make corrections to the settings, as well as a (set of) default setting(s) that is applied upon first use of the system. This would thus still require our static, profile-based approach.

APPENDIX

The legend associated with abbreviations is explained in Table 1.

Fig. 24. Plots of various values of **Storage** vs different Attitudes and Decision.

Fig. 25. Plots of various values of **Action** vs. different Attitudes and Decision.

Fig. 26. Plots of various values of **Purpose** vs. different Attitudes and Decision.

Fig. 27. Plots of various values of **Who** vs. different Attitudes and Decision.

Fig. 28. Plots of various values of **What** vs. different Attitudes and Decision.

Familiarity with smart home devices

How familiar are you with smart home devices?

What kind of device comes to mind when you think of a **smart refrigerator**?

Please mention any particular device(s) or brand(s) you are familiar with. Otherwise, describe in detail what kind of device you are thinking of.

Have you yourself used any smart refrigerator(s)?

I've never used such device

I've used such device, but only occasionally

I use such device several times a month

I use such device several times a week

I use such device every day

-134% done

Continue

Fig. 29. Open-ended questions asked to participants in order to gauge their familiarity with IoT.

What is the Internet of Things?

Here is a little test to make sure you paid attention to the video.

Which of the following things can smart household devices infer about you?

whether I am at home

whether I am sleeping

what room I am in

what is in my fridge

who is entering my house

all of these

Which of the following devices was **not** mentioned on the previous page?

smart TV

smart refrigerator

smart alarm clock

smart HVAC

smart car

smart lighting system

smart washing machine

Which of the following is **not** going to happen with your data?

stored locally

used to improve the service

disclosed to the public

used for advertisement

-144% done

Continue

Fig. 30. Attention check questions asked to participants.

What is the Internet of Things?

Sorry, you got 0 or more of the test questions wrong! Please read the following text again, carefully.

Imagine you are planning to buy a new house and you meet this realtor who shows you a really awesome house. Apart from its beautiful design and great location, it comes with loads of **"Smart" electronic devices**, which can be conveniently controlled using your smartphone:

What kind of devices are we talking about?

- **A smart refrigerator** senses when a product spoils or needs to be replenished, and reminds you when it is time to buy fresh groceries. It also tracks your family's diet, like how much candy your kids are eating..
- **A smart HVAC** with a thermostat that programs itself based on the season, your body temperature, and your daily activities. It makes sure that your house is at a comfortable temperature when you arrive home, and automatically adjusts to when you are away or asleep. It can even adjust the temperature in each room according to the personal preferences of its occupants. The smart HVAC also keeps track of your total energy savings.
- **Smart light fixtures** that turn on automatically when you enter a room. The system adjusts the lights in each room to your preferences and the incoming sunlight. The lights can also wake you up, and adjust the illumination for an optimal experience when you want to read, concentrate, or relax based on your routine. The lights can even synchronize with your smart TV for an immersive effect when you watch movies or play games.
- **A smart TV** that automatically records shows based on your personal preferences. It can detect who is watching, and suggest shows accordingly. It provides regular updates about your personal life (email, weather, upcoming events) and the status of other smart devices. It also allows family members to leave messages for each other that automatically play when the recipient enters the house.
- **A smart washing machine** that knows what kind of clothes you put in it and adjusts the washing cycle accordingly. The washer automatically runs when electricity rates are lowest. If you are not at home, it tumbles clothes in fresh air after the cycle is over. Finally, it has an app that you can use to assign laundry chores to family members, who will be reminded when it is time to load or unload the washer.
- **A smart home security system** that allows you to check the security of your home via several cameras in and around the house, as well as smoke detectors and flood sensors. It automatically locks your doors, and you can set up keyless entry for yourself and your family. You can also talk to people at the door via video, and let them in remotely. The system can identify a walk-in guest as stranger or frequent visitor, and notifies you accordingly in case you are unaware of their presence.
- **A smart alarm clock** that detects your sleeping routines, as well as your movements during your sleep, to give you feedback on your sleeping patterns. It knows when you have to go to work based on your schedule and informs you of current traffic. It gently wakes you by adjusting the lights and turning on your favorite music. In case of bad weather, it can even automatically call you an Uber.
- **A smart assistant** that provides a gateway to online resources. It can play music from online streaming services, order products from online retailers, and answer questions by searching the internet. A smart assistant can also be used to control and learn the status of other smart home devices. You can also use it as an intercom system to talk to other people in your household, and to make hands-free phone calls to others.

What can these devices do?

These smart devices have embedded location sensors, cameras, and microphones to detect your presence or your exact location in- and outside the house. They can also connect with each other or with your smartphone or smartwatch in order to provide timely alerts or to automate their operations. For example, your security system may connect to your alarm clock to detect when you go to bed and wake up; your light fixtures may use location sensors to know when you leave a room; your refrigerator may connect to your smartphone to provide recipes and grocery lists; your washing machine may connect to your HVAC to slightly increase the room temperature when you hang your laundry to dry.

What happens with your information?

A central feature of the smart devices is that they are connected to the Internet. Sometimes they communicate and store their information locally; at other times they operate via the cloud or a central server. In some cases the stored information is additionally used to give you insight in your behavior, to optimize the current service, to recommend additional services to you, or for advertisement purposes.

How useful is this to you?

As you can see, this house has a lot of novel features. You may find some of these features really useful, or rather invasive at times. On the next pages, you will find 20 scenarios, and you will be asked to evaluate them. Please read each scenario carefully before answering the questions.

-154% done

Continue

Fig. 31. Transcript of video shown to participants if they failed attention checks.

Questions about you

You are done with the information collection part of the system. We now ask you to answer a few questions about you.

The remaining part of this study will take about 10 minutes.

Another question to make sure you are not a robot: Approximately how many minutes will the remaining part of this study take?

10

83% done

Continue

Fig. 32. Attention check question asked to participants.

Scenario 4 out of 13

Your smart assistant connects to your smart phone/watch to detect your location in the house. The data is stored on a remote server and used to give you insight into your behavior.

Would you **disable** this feature?

yes

no

How **uncomfortable** or **comfortable** do you feel about this scenario?

very uncomfortable

uncomfortable

somewhat uncomfortable

neutral

somewhat comfortable

comfortable

very comfortable

How **useless** or **useful** do you find this scenario?

completely useless

useless

somewhat useless

neutral

somewhat useful

useful

very useful

Please select **completely disagree** below.

completely disagree

disagree

somewhat disagree

neutral

somewhat agree

agree

completely agree

How **risky** or **safe** do you find this scenario?

very risky

risky

somewhat risky

neutral

somewhat safe

safe

very safe

How **inappropriate** or **appropriate** do you find this scenario?

very inappropriate

inappropriate

somewhat inappropriate

neutral

somewhat appropriate

appropriate

very appropriate

How **unexpected** or **expected** do you find this scenario?

very unexpected

unexpected

somewhat unexpected

neutral

somewhat expected

expected

very expected

36% done

Continue

Fig. 33. Attention check question shown while participants are answering questions per scenario.

Scenario 1 out of 13

Your smart lighting system uses information collected by your smart alarm clock to detect your location in the house. The data is stored on a remote server and used to optimize the service, as well as shared with third parties.

Data collected by this device will be shared with third-party affiliates.

Would you **disable** this feature?

yes

no

How **uncomfortable** or **comfortable** do you feel about this scenario?

very
uncomfortable

uncomfortable

somewhat
uncomfortable

neutral

somewhat
comfortable

comfortable

very comfortable

How **useless** or **useful** do you find this scenario?

completely
useless

useless

somewhat
useless

neutral

somewhat useful

useful

very useful

What could be specific benefits to you in this scenario? (be as specific as possible; feel free to speculate)

How **risky** or **safe** do you find this scenario?

very risky

risky

somewhat risky

neutral

somewhat safe

safe

very safe

What could be specific risks to you in this scenario? (be as specific as possible; feel free to speculate)

How **inappropriate** or **appropriate** do you find this scenario?

very
inappropriate

inappropriate

somewhat
inappropriate

neutral

somewhat
appropriate

appropriate

very appropriate

How **unexpected** or **expected** do you find this scenario?

very unexpected

unexpected

somewhat
unexpected

neutral

somewhat
expected

expected

very expected

20% done

Continue

Fig. 34. Example of one of the thirteen scenarios presented to the participants.

REFERENCES

- [1] Alessandro Acquisti and Ralph Gross. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Proceedings of the International Workshop on Privacy Enhancing Technologies*. 36–58.
- [2] Icek Ajzen and Martin Fishbein. 1977. Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin* 84, 5 (1977).
- [3] Naveen Farag Awad and M. S. Krishnan. 2006. The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS Quarterly* 30, 1 (March 2006), 13–28.
- [4] Paritosh Bahirat, Yangyang He, and Bart P. Knijnenburg. 2018. Exploring defaults and framing effects on privacy decision making in smarthomes. In *Proceedings of the USENIX Symposium on Usable Privacy and Security*. Baltimore, MD.
- [5] Paritosh Bahirat, Yangyang He, Abhilash Menon, and Bart Knijnenburg. 2018. A data-driven approach to developing IoT privacy-setting interfaces. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces (IUI'18)*. ACM, Tokyo, Japan, 165–176. DOI: <https://doi.org/10.1145/3172944.3172982>
- [6] Paritosh Bahirat, Qizhang Sun, and Bart P. Knijnenburg. 2018. Scenario context V/s framing and defaults in managing privacy in household IoT. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces Companion (IUI'18)*. ACM, Tokyo, Japan, 63:1–63:2. DOI: <https://doi.org/10.1145/3180308.3180372>
- [7] Ramnath K. Chellappa and Raymond G. Sin. 2005. Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management* 6, 2–3 (2005), 181–202. DOI: <https://doi.org/10.1007/s10799-005-5879-y>
- [8] Nigel Davies, Nina Taft, Mahadev Satyanarayanan, Sarah Clinch, and Brandon Amos. 2016. Privacy mediators: Helping IoT cross the chasm. In *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications (HotMobile'16)*. ACM, New York, 39–44. DOI: <https://doi.org/10.1145/2873587.2873600>
- [9] Cailing Dong, Hongxia Jin, and Bart P. Knijnenburg. 2016. PPM: A privacy prediction model for online social networks. In *Proceedings of the International Conference on Social Informatics*. 400–420.
- [10] Opher Etzion and Fabiana Forunier. 2014. On the personalization of event-based systems. In *Proceedings of the 1st ACM International Workshop on Human Centered Event Understanding from Multimedia*. ACM, 45–48.
- [11] Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web*. 351–360.
- [12] Denis Feth, Andreas Maier, and Svenja Polst. 2017. A user-centered model for usable security and privacy. In *Human Aspects of Information Security, Privacy and Trust (Lecture Notes in Computer Science)*, Theo Tryfonas (Ed.). Springer International Publishing, 74–89.
- [13] Hemant Ghayvat, S.C. Mukhopadhyay, Jie Liu, Arun Babu, Md Alahi, and Xiang Gui. 2015. Internet of Things for smart homes and buildings: Opportunities and Challenges. *Australian Journal of Telecommunications and the Digital Economy* 3 (12 2015), 33–47.
- [14] Nathaniel Good, Rachna Dhamija, Jens Grossklags, David Thaw, Steven Aronowitz, Deirdre Mulligan, and Joseph Konstan. 2005. Stopping spyware at the gate: A user study of privacy, notice and spyware. In *Proceedings of the 2005 Symposium on Usable Privacy and Security*. 43–52. DOI: <https://doi.org/10.1145/1073001.1073006>
- [15] Ralph Gross and Alessandro Acquisti. 2005. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*. 71–80.
- [16] Alexander Henka, Lukas Smirek, and Gottfried Zimmermann. 2016. Personalizing smart environments. In *Proceedings of the 6th International Conference on the Internet of Things*. ACM, 159–160.
- [17] Shuk Ying Ho and Kar Tam. 2006. Understanding the impact of web personalization on user information processing and decision outcomes. *MIS Quarterly* 30, 4 (Dec. 2006), 865–890.
- [18] Robert C. Holte. 1993. Very simple classification rules perform well on most commonly used datasets. *Machine Learning* 11, 1 (01 Apr 1993), 63–90. DOI: <https://doi.org/10.1023/A:1022631118932>
- [19] Kai-Lung Hui, Bernard C. Y. Tan, and Chyan-Yee Goh. 2006. Online information disclosure: Motivators and measurements. *ACM Transactions on Internet Technology* 6, 4 (Nov. 2006), 415–441. DOI: <https://doi.org/10.1145/1183463.1183467>
- [20] Prem Prakash Jayaraman, Xuechao Yang, Ali Yavari, Dimitrios Georgakopoulos, and Xun Yi. 2017. Privacy preserving Internet of Things: From privacy techniques to a blueprint architecture and efficient implementation. *Future Generation Computer Systems* 76 (Nov. 2017), 540–549. DOI: <https://doi.org/10.1016/j.future.2017.03.001>
- [21] Carlos Jensen and Colin Potts. 2004. Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of the 2004 Conference on Human Factors in Computing Systems*. 471–478. DOI: <https://doi.org/10.1145/985692.985752>
- [22] Bart P. Knijnenburg. 2015. *A User-tailored Approach to Privacy Decision Support*. Ph.D. dissertation. University of California, Irvine, Irvine, CA. <http://search.proquest.com/docview/1725139739/abstract>.

- [23] Bart P. Knijnenburg. 2017. Privacy? I can't even! Making a case for user-tailored privacy. *IEEE Security & Privacy* 15, 4 (2017), 62–67.
- [24] Bart P Knijnenburg and Alfred Kobsa. 2016. Taking control of household IoT device privacy. *CCC Sociotechnical Cybersecurity Workshop* (2016).
- [25] Bart P Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71, 12 (2013), 1144–1162.
- [26] Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Preference-based location sharing: Are more privacy options really better? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Paris, France, 2667–2676. DOI : <https://doi.org/10.1145/2470654.2481369>
- [27] Alfred Kobsa, Hichang Cho, and Bart P. Knijnenburg. 2016. 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 (Feb. 2016), 2587–2606. Issue 11. DOI : <https://doi.org/10.1002/asi.23629>
- [28] Hosub Lee and Alfred Kobsa. 2016. Understanding user privacy in Internet of Things environments. In *Proceedings of the 2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)* (2016), 407–412.
- [29] Yao Li, Alfred Kobsa, Bart P Knijnenburg, and MH Carolyn Nguyen. 2017. Cross-cultural privacy prediction. In *Proceedings on Privacy Enhancing Technologies* 2 (2017), 93–112.
- [30] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhammedi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Proceedings of the 2016 Symposium on Usable Privacy and Security*. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu>.
- [31] Franco Loi, Arunan Sivanathan, Hassan Habibi Gharakheili, Adam Radford, and Vijay Sivaraman. 2017. Systematically evaluating security and privacy for consumer IoT devices. In *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy (IoTS&P'17)*. ACM, 1–6. DOI : <https://doi.org/10.1145/3139937.3139938>
- [32] Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. 2012. A study of privacy settings errors in an online social network. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops*. 340–345.
- [33] Helen Nissenbaum. 2004. Privacy as contextual integrity symposium - Technology, values, and the justice system. *Washington Law Review* 79 (2004), 119–158.
- [34] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100–126. DOI : <https://doi.org/10.1111/j.1745-6606.2006.00070.x>
- [35] Judith S. Olson, Jonathan Grudin, and Eric Horvitz. 2005. A study of preferences for sharing and privacy. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems*. 1985–1988.
- [36] Gautham Pallapa, Sajal K. Das, Mario Di Francesco, and Tuomas Aura. 2014. Adaptive and context-aware privacy preservation exploiting user interactions in smart environments. *Pervasive and Mobile Computing* 12 (2014), 232–243.
- [37] Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: Designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'14)*. ACM, New York, 897–908. DOI : <https://doi.org/10.1145/2632048.2632062>
- [38] Charith Perera, Ciaran McCormick, Arosha K. Bandara, Blaine A. Price, and Bashar Nuseibeh. 2016. Privacy-by-design framework for assessing Internet of Things applications and platforms. In *Proceedings of the 6th International Conference on the Internet of Things (IoT'16)*. ACM, New York, 83–92. DOI : <https://doi.org/10.1145/2991561.2991566>
- [39] Joseph Phelps, Glen Nowak, and Elizabeth Ferrell. 2000. Privacy concerns and consumer willingness to provide personal information. *Journal of Public Policy & Marketing* 19, 1 (2000), 27–41. DOI : <https://doi.org/10.1509/jppm.19.1.27.16941>
- [40] Frederic Raber, Alexander De Luca, and Moritz Graus. 2016. Privacy wedges: Area-based audience selection for social network posts. In *Proceedings of the 2016 Symposium on Usable Privacy and Security*. <https://www.usenix.org/conference/soups2016/workshop-program/wpi/presentation/raber>.
- [41] Ramprasad Ravichandran, Michael Benisch, Patrick Gage Kelley, and Norman M Sadeh. 2009. Capturing social networking privacy preferences. In *Proceedings of the 2009 Symposium on Usable Privacy and Security*. 1–18.
- [42] Luke Russell, Rafik Goubran, and Felix Kwamena. 2015. Personalization using sensors for preliminary human detection in an IoT environment. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS'15)*. IEEE, 236–241.
- [43] Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, and Jinghai Rao. 2009. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing* 13, 6 (2009), 401–412.
- [44] R. S. Sandhu and P. Samarati. 1994. Access control: Principle and practice. *IEEE Communications Magazine* 32, 9 (1994), 40–48. DOI : <https://doi.org/10.1109/35.312842>

- [45] Hong Sheng, Fiona Fui-Hoon Nah, and Keng Siau. 2008. An experimental study on ubiquitous commerce adoption: Impact of personalization and privacy concerns. *Journal of the Association for Information Systems* 9, 6 (June 2008), 344–376. <http://aisel.aisnet.org/jais/vol9/iss6/15>.
- [46] N. Craig Smith, Daniel G. Goldstein, and Eric J. Johnson. 2013. Choice without awareness: Ethical and policy implications of defaults. *Journal of Public Policy & Marketing* 32, 2 (2013), 159–172. DOI: <https://doi.org/10.1509/jppm.10.114>
- [47] Juliana Sutanto, Elia Palme, Chuan-Hoo Tan, and Chee Wei Phang. 2013. Addressing the personalization-privacy paradox: An empirical assessment from a field experiment on smartphone users. *MIS Quarterly* 37, 4 (2013), 1141–1164.
- [48] Max Teltzrow and Alfred Kobsa. 2004. Impacts of user privacy preferences on personalized systems: A comparative study. In *Designing Personalized User Experiences for eCommerce*, Clare-Marie Karat, Jan Blom, and John Karat (Eds.). Kluwer Academic Publishers, Dordrecht, Netherlands, 315–332. DOI: [10.1007/1-4020-2148-8_17](https://doi.org/10.1007/1-4020-2148-8_17)
- [49] Horst Treiblmaier and Irene Pollach. 2007. Users' perceptions of benefits and costs of personalization. In *Proceedings of the ICIS 2007*.
- [50] Blase Ur, Jaeyeon Jung, and Stuart Schechter. 2014. Intruders versus intrusiveness: Teens' and parents' perspectives on home-entryway surveillance. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'14)*. ACM, New York, 129–139. DOI: <https://doi.org/10.1145/2632048.2632107>
- [51] PriceWaterhouseCooper US. [n.d.]. Smart home, seamless life: Unlocking a culture of convenience. <https://www.pwc.com/us/en/advisory-services/publications/consumer-intelligence-series/smarthome.html> [Online; accessed 27-Jan-2018].
- [52] Thibaut Vallée, Karima Sedki, Sylvie Despres, M-Christine Jaulant, Karim Tabia, and Adrien Ugon. 2016. On personalization in IoT. In *Proceedings of the International Conference on Computational Science and Computational Intelligence (CSCI'16)*. IEEE, 186–191.
- [53] Gregg Vanderheiden and Jutta Treviranus. 2011. Creating a global public inclusive infrastructure. In *Proceedings of the International Conference on Universal Access in Human-Computer Interaction*. Springer, 517–526.
- [54] Yang Wang, Liang Gou, Anbang Xu, Michelle X. Zhou, Huahai Yang, and Hernan Badenes. 2015. VeilMe: An interactive visualization tool for privacy configuration of using personality traits. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 817–826. <http://dl.acm.org/citation.cfm?id=2702293>.
- [55] Jason Watson, Andrew Besmer, and Heather Richter Lipford. 2012. +Your circles: Sharing behavior on Google+. In *Proceedings of the 8th Symposium on Usable Privacy and Security*. 12:1–12:10. DOI: <https://doi.org/10.1145/2335356.2335373>
- [56] Meredydd Williams, Jason R. C. Nurse, and Sadie Creese. 2016. The perfect storm: The privacy paradox and the Internet-of-Things. In *Proceedings of the 11th International Conference on Availability, Reliability and Security*. 644–652.
- [57] Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies* 98 (2017), 95–108.
- [58] Ian H Witten, Eibe Frank, Mark A Hall, and Christopher J Pal. 2016. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan-Kaufmann.
- [59] Peter Worthy, Ben Matthews, and Stephen Viller. 2016. Trust me: Doubts and concerns living with the Internet of Things. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS'16)*. ACM, New York, 427–434. DOI: <https://doi.org/10.1145/2901790.2901890>

Received May 2018; revised September 2018; accepted November 2018