2018 IEEE Symposium on Privacy-Aware Computing

The Potential for User-Tailored Privacy on Facebook

Moses Namara
School of Computing
Clemson University
Clemson, SC, USA
mosesn@clemson.edu

Henry Sloan
Nyack High School
Nyack, NY, USA
henryksloan@gmail.com

Priyanka Jaiswal
School of Computing
Clemson University
Clemson, SC, USA
pjaiswa@clemson.edu

Bart P. Knijnenburg
School of Computing
Clemson University
Clemson, SC, USA
bartk@clemson.edu

***Abstract*—Research shows that Facebook users differ extensively in their use of various privacy features, and that they generally find it difficult to translate their desired privacy preferences into concrete interface actions. Our work explores the use of User-Tailored Privacy (UTP) to adapt Facebook's privacy features to the user's personal preferences. We developed adaptive versions of 19 Facebook privacy features, and for each feature we test three adaptation methods (Automation, Highlight and Suggestion) that can be used to implement the adaptive behavior. In a "think-aloud" semi-structured interview study (N=18), we show participants paper prototypes of our adaptive privacy features and ask participants to judge the presented adaptive capabilities and the three adaptation methods that implement them. Our findings provide insights into the viability of User-Tailored Privacy. Specifically, we find that the optimal adaptation method depends on the users' familiarity with the privacy feature and how they use them, and their judgment of the awkwardness and irreversibility of the implemented privacy functionality. We conclude with design recommendations for the implementation of User-Tailored Privacy on Facebook and other social network platforms.**

Keywords—privacy, social media, Facebook, user-tailored privacy, privacy on social media

I. Introduction

As one of the most used social network sites, Facebook has a plethora of privacy controls and features in place to give its users more control over their privacy settings [1]. While these features are certainly comprehensive, previous research has shown that users have varying individual privacy preferences [2], [3], they have a hard time translating their desired privacy levels into concrete interface actions [4], and often avoid the hassle of utilizing the available controls despite their stated interest in having control over their private information [5].

Given these difficulties, advocates of User-Tailored Privacy (UTP) suggest making it easier to manage one's privacy by automatically tailoring a system's privacy settings to the user's preferences [6] in order to find a right fit between user's desire for privacy and their actual privacy experiences [1]. Facebook's privacy management functionality goes beyond simple "settings" though—Facebook provides a multitude of privacy features [7], and users have been shown to utilize distinct, coherent subsets of these features [3]. The goal of UTP is thus to support and/or complement these privacy management strategies, which arguably provides users with just the right amount of control and useful privacy-related information so as not to be overwhelming or misleading.

This work was supported by a Facebook Emerging Scholar Award and a DoD Award W911QY-16-C-0105.

Making Facebook's privacy features user-tailored may not be as futuristic as it may sound. Indeed, Polisis [8], a generic framework that provides automatic privacy policy analysis, suggests that Facebook already utilizes user data to tailor its services and personalize its interface to its users.

While there is a considerable amount of research on making privacy functionality adapt itself to users' preferences [9]–[13], the successful implementation of user-tailored privacy features is not an easy task [14]. Assuming that it is possible, though, we are still left with two important research questions: *which* features should be tailored to the user's preferences, and *how* should such adaptations be implemented? To answer these research questions, we explore user's reactions to user-tailored versions of 19 Facebook privacy features. For each user-tailored feature, we consider three interface adaptation methods that implement the user-tailored behavior. Our study comprises a think-aloud style evaluation of these privacy features with 18 participants.

Our results confirm previous research findings that users gain the most benefits when social network sites give them the privacy they desire [1], and that this can be accomplished by tailoring the privacy features to the user's privacy preferences. More specifically, though, we find (1) that participants have profoundly varying opinions on the different adaptations, (2) that participants prefer different types of adaptation methods for different privacy features, depending on their familiarity with the privacy feature and the perceived severity/irreversibility of the privacy mechanism represented by the feature.

In the remainder of this paper, we first present related work and our research questions. We then describe how we conducted our semi-structured interviews and present the results. Finally, we discuss our findings and conclude with design implications and suggestions for future work towards the successful implementation of user-tailored privacy.

II. Related Work

A. Facebook Users' Privacy Behaviors

Social network sites offer a wide variety of mechanisms to protect users' "privacy boundaries" [2], [7]. Research shows that Facebook users differ substantially in the extent to which they employ these privacy protection mechanisms [3], and that users' experience can be enhanced if the protection offered by the system matches their privacy needs [1]. Unfortunately, though, users often fail to effectively manage their privacy on social networks [4], [15], [16]. A reason for these privacy management failures is that social network users' privacy decisions—like most decisions—often fall prey to heuristic influences such as the neat appearance and design of a website, the difficulty of mentally picturing the

978-1-5386-8442-9 /18/$31.00 ©2018 IEEE
DOI 10.1109/PAC.2018.00010

consequences of identity theft ("availability heuristic"), others' privacy decisions ("social proof"), and the available privacy options to choose from ("context non-invariance") [14], [17]. And while researchers have developed various ways to increase the transparency and control of the privacy functionality of social network sites [18]–[21], two "paradoxes" remain with any privacy feature that requires users to control their own privacy: the first is the "transparency paradox", which states that privacy notices that are sufficiently detailed to have an impact on people's privacy decisions are often too long, detailed and complex for people to read [22]; the other is the "control paradox", which states that while users often claim to want full control over their data, they often avoid the hassle of actually exploiting this control [5]. Consequently, several scholars have recently questioned the effectiveness of putting users in full control over their privacy [22]–[24].

B. *User-Tailored Privacy*

One way to alleviate users' privacy decision-making burden is through User-Tailored Privacy (UTP). Knijnenburg et al. [14] define UTP as an approach that provides decision support by *measuring* users' privacy preferences and behaviors, using the measurements to create a personalized *model* and finally *adapting* the user interface to the predicted privacy preferences by changing the default privacy settings (see Fig 1).

For the *measure* part of UTP, user preferences and behaviors have been found to differ among users. They can be drawn from personal and contextual factors such as the data requested ("what"), user ("who"), system/recipient of information("whom"). For example, Wang et al. [25] found that people are comfortable disclosing their interests, groups, religions and links on their social network pages but are least comfortable disclosing their e-mails, street addresses and phone number. Similarly, Dong et al [26] found that time (weekday or weekend, daytime or evening) are important determinants of user's willingness to disclose their location. Such characteristics and contextual factors can be quantified and used as input for modelling user privacy.

A considerable amount of research has focused on the *model* part of UTP. For example, Wisniewski et al. [3] investigate the dimensionality of the privacy behaviors of 308 Facebook users, and extract 11 behavioral strategies. Clustering users on these strategies, they find 6 privacy management profiles: Privacy Maximizers, Selective Sharers, Privacy Balancers, Self-Censors, Time Savers/Consumers and Privacy Minimalists. Based on these profiles, Wilkinson et al. [27] proposed a "user-tailored privacy-by-design" approach: they created a more prominent version of each privacy feature, and integrated them in a user-tailored manner by only using the more prominent versions of the features that fit the user's profile.

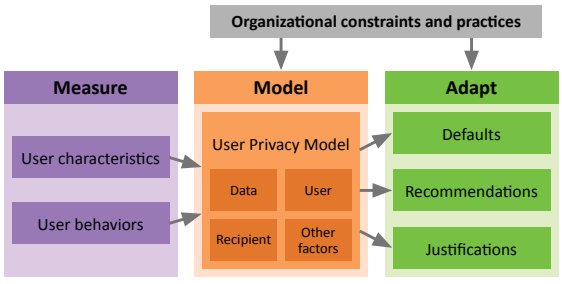

Fig. 1. A schematic overview of User-Tailored Privacy.

Fang and LeFevre [10] propose a privacy wizard that can automatically assign privileges to a user's Facebook friends. The wizard iteratively asks users to assign privacy "labels" to selected friends. This input is then used to construct a classifier that automatically assigns privileges to the rest of the user's friends. On evaluation of the wizard with privacy preference data collected from 45 real Facebook users, the study found that the privacy wizard can generate highly accurate privacy settings with minimal user input.

Dong et al. [9] created a generic Privacy Prediction Model and applied it to users' Twitter and Google+ followers to generate recommendations of whom to follow back. Their model shows that follower/following characteristics of the user and the follower, as well as the overlap between them, are valuable predictors in determining whom the user is likely to follow back.

In the realm of location-sharing, several researchers have attempted to model users' sharing preferences using machine learning algorithms [13], [28]–[31]. Most work in this field agrees that the time, place, and recipient of the shared location are the most important factors in predicting users' intention to share.

Moving beyond social networks, Liu et al.[32] analyzed the privacy settings of 4.8 million smartphone users, and found that while people's mobile app privacy preferences are diverse, a relatively small number of profiles can be identified to help simplify their privacy decision making process. Similar profile-based solutions have been proposed for social network privacy behaviors [3], location sharing [33], and IoT privacy settings [34].

C. *Testing Adaptation Methods*

The works mentioned in the previous section have made important contributions that can support users' privacy decision-making practices in a user-tailored way. However, fewer works have carefully examined and tested the *adapt* part of UTP [14]. Those few works have shown that adaptations are generally welcomed by users. For example, in a field study of a personalized privacy assistant with 72 participants, Liu et al. [12] found that 78.7% of recommendations made by the assistant were adopted by users. Likewise, Knijnenburg and Jin [35] found that users in their study accepted between 62.5% and 98.7% of the presented location-sharing recommendations.

In the latter study, the percentage of accepted recommendations depended on the length of the list of recommendations, and the way they were presented (the system either highlighted the recommended behaviors or hid the ones that were not recommended). This raises the important point that there exist various "adaptation methods", i.e., ways in which suggested behaviors can be presented to the user. Beyond hiding and highlighting, one can automatically implement the behavior, e.g. by automatically changing the default setting [36], or give explicit suggestions on what behaviors to implement [37].

The optimal adaptation method remains an open question [14]—a question we seek to answer in this paper. Moreover, given that platforms like Facebook have a plethora of privacy features that can all potentially be adapted to the user's

preferences, we also investigate whether each of these features should be tailored to the user at all. In line with this argument, we pose the following two research questions:

RQ1: Which features should be tailored to the user's preferences?

RQ2: How should such adaptations be effected?

Unfortunately, while there is a plethora of research on the user experience of personalized systems (cf. [38]), there is a dearth of research on when such personalization is desired in the first place, and if so, what degree of automation would be desired [39]. Sheridan argues that this likely depends on the situation but provides little concrete guidance as to which situations are most suitable for personalization [40].

To fill this gap, we propose a conceptual understanding of personalization as a tradeoff between ease of use and control: Automation reduces the need to engage in a task by oneself, at the cost of relinquishing some control over the task. Given this ease of use/control tradeoff, we argue that users' desire for tailoring the privacy features depends on their awareness and usage of these features. As for the privacy features that the user most frequently uses, one could argue that tailoring these features would significantly make privacy management easier to perform. On the other hand, users may not want to relinquish control over these most-used features. Similarly, users may not be interested in using the privacy protection mechanisms that they use infrequently, which would suggest that tailoring these features is an assault on their desire to control what privacy management practices they engage in. On the other hand, users may avoid certain features simply because they consider them cumbersome, which would suggest that tailoring such features would actually be beneficial, due to the increased ease of use. Finally, users may feel uncomfortable about adaptations to features that reside outside their awareness, as such adaptations would take place outside the purview of their control. On the other hand, such adaptations may make it easier for them to discover new privacy protection mechanisms.

In sum, it is not clear which features should be tailored to the user's preferences—it crucially depends on users' perceptions of the effort related to engaging with various privacy features and their desire for control over these features. Hence, our study attempts to answer RQ1 by investigating the ease of use/control tradeoff to determine which features should be tailored to the user's preferences, and which features should remain untailored.

A similar ease of use/control tradeoff applies in the context of RQ2. Specifically, adaptations at the highest degree of automation ("the computer decides on everything and acts autonomously") are easiest to use, as they require no active input from the user. However, if a privacy feature automatically makes decisions on the user's behalf, this may be regarded as a severe reduction in control. Hence, we predict that the highest degree of automation will only be suitable for those privacy features that users find effortful to use, but do not desire a high level of control over.

According to Sheridan and Verplank [39], lower degrees of automation could involve a system explicitly suggesting an option to the user without automatically executing it. Following up on such active suggestions is more taxing, but ultimately gives users control over the actual privacy feature. Lower degrees of automation are thus arguably most suitable in situations where control is desirable, and a certain amount of effort is justifiable, e.g. for the purpose of educating users about a privacy feature. Active suggestions are especially suitable for such situations, because the suggestion can be explained to the user.

Finally, intermediate degrees of automation may involve implicit suggestions, e.g. by highlighting suggested actions [35]. Highlighting recommended actions keeps the user in control over these actions, so it does not decrease the physical burden, but it does increase the ease of use by indicating to the user what action is likely most suitable. Moreover, highlighting is an implicit suggestion, which is likely less taxing than an explicit suggestion. Highlighting is thus arguably most suitable for privacy features that users do not want to fully automate, but where explicit suggestion would require considerable effort to accomplish. In sum, it is likely that users will prefer higher degrees of automation for some privacy features and lower degrees for other features. We predict that the optimal degree of automation depends on their ease of use/control tradeoff. Hence, our study attempts to answer RQ2 by investigating this tradeoff to determine which privacy features—among those that should be tailored to the user's preferences—should be tailored at a higher degree of automation, and which features should be tailored to a lower degree of automation.

PARTICIPANT GENDER, AGE GROUP, AND EXPERIMENTAL TREATMENT (FEATURES SHOWN).

ID	Gender	Age group	Features shown
A	F	18-21	1,4,9,15,16,17
B	M	21-25	4,7,9,10,13,17
C	M	21-25	4,7,8,10,11,13
D	M	18-21	4,6,8,11,14,19
E	F	18-21	5,6,12,14,18,19
F	M	18-21	4,5,8,11,12,18
G	F	35-40	4,8,11,16,18,19
H	M	18-21	2,5,7,12,16,18
I	M	25-30	2,3,5,12,14,18
J	M	25-30	1,6,9,13,14,17
K	M	25-30	1,5,6,10,14,16
L	M	25-30	5,9,10,13,16,17
M	F	25-30	2,3,7,9,12,13,17
N	F	20-25	2,4,5,7,8,11
O	M	20-25	2,5,7,12,15,16
P	M	20-25	1,2,4,15,16,17
Q	M	25-30	1,4,6,14,17,18
R	M	25-30	5,6,12,14,18,19

III. METHODOLOGY

To answer our research questions: *which* features should be tailored to the user's preferences, and *how* should such adaptations be implemented? We created 19 mockups of "user-adaptive" versions of Facebook privacy features. Implementing each adaptive feature with three different adaptation methods (Automation, Highlight and Suggestion) at varying levels of automation. We carried out a series of semi-structured user interviews with 18 participants, showing them paper prototypes of our adaptive privacy features, and asking them to judge the presented adaptive capabilities and the three adaptation methods. In this section, we describe our participant recruitment and interview procedures.

A. Recruitment and Participants

Between October and December of 2017, we recruited adult self-reported Facebook users with the purpose of collecting their feedback on our adaptive privacy features and

adaptation methods. They were recruited through flyers around a university campus and the surrounding area, and via email using university student email listservs. 18 participants each completed the 45-minute interview session; their demographics are shown in Table I.

B. Interface Mockups

The instrument for our study was a set of paper-based mockups of the 19 Facebook privacy features listed in Table II. The choice of these features is inspired by Wisniewski et al. who map out an exhaustive set of boundary regulation mechanisms on various social network sites in [2], and identify Facebook features implementing these boundary regulation mechanisms in [3]. We called our system "Fakebook" and used cartoon-style renderings to have the participants focus on the presented mechanism rather than the specific graphical implementation and the feasibility of the adaptive technology. For each feature, we created a mockup of the default non-adaptive version currently available on Facebook, plus three adaptive version, with each version implementing a different adaptation method: automation, highlight or suggestion. These three adaptation methods implement varying degrees of automation [39], and are further discussed below.

1) Automation

The Automation adaptation method implements adaptations without first requesting permission from the user. This adaptation method has the highest degree of automation, as it can operate completely outside of the user's awareness. In our implementation, the user is not explicitly notified of the automatic adaptation, but they are able to see that automated action has occurred when they arrive at the location where they would have done the action themselves. For example, when a user is untagged from a post, a participant shown the Automation method would see the tag removed and replaced with a message informing them that they were automatically untagged (see Fig 2).

The Automation method substantially reduces the *onus* of privacy decision-making but can feel like a significant loss of *control* [38], [39]. Indeed, Vihavainen et al.[41] studied the implications of full automation on social interaction on social network sites (SNS) and found that the loss of granular control leaves users feeling powerless to adjust the specifics of what is being disclosed. Optimization of such details of disclosure is a task that users still feel cannot readily and correctly be transferred from them to a system. Hence, in our designs, the message indicating the automated action has an undo button, allowing the user to reverse the action. The undo button makes the Automation method more similar to the other adaptation methods (which require user intervention before the adaptation is enacted) and is predicted to increase the perceived control of this adaptation method, without harming its inherently unobtrusive nature.

2) Highlight

The Highlight adaptation method increases the visual prominence of the action that the adaptive procedure predicts the user would want to take. This can be done either through a color change, or by giving the recommended action a more prominent location on the screen. In our implementation, we give the recommended action a yellow background color, and change its ordering in the list of options, if appropriate. The Highlight method implements a moderate degree of automation: it gives users a clear indication as to what action they should consider—reducing their cognitive *load* without reducing their *control*.

As some privacy features in the Facebook interface are hidden behind a button or a menu, our Highlight implementation can also highlight the element that gives access to the adapted feature. For example, when a user is missing important basic information such as political views (See Fig.3), a highlight on this missing information and of the feature that enables users to edit this basic information could be necessary. The highlight provides guidance to users in cases where the adapted feature is not prominent.

3) Suggestion

The Suggestion adaptation method displays an "agent" (virtual character) that verbally suggests a recommended action to the user. Our implementation is based on Facebook's "Privacy Dinosaur", which the Facebook platform currently uses to display "Privacy Check-up" notifications to the user. The Dinosaur provides suggestions in a general form of, "I think you should...", increasing the personal nature of the interaction (see Fig 4). The provided options are "Ok" and "Rather Not", allowing the user to either accept or reject the recommended action. Users were told that if they selected "Ok", the setting would automatically be changed however they would still be taken to the appropriate setting as well. By asking for an explicit decision, this adaptation method implements our lowest degree of automation.

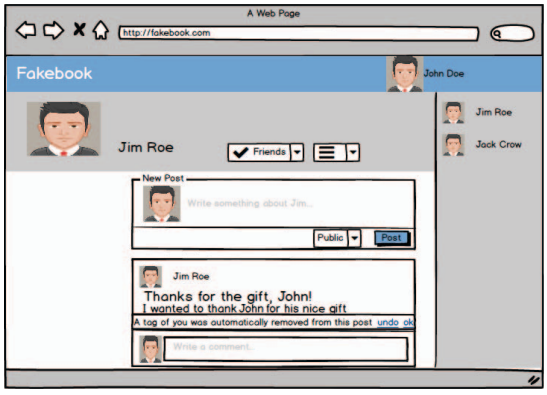

Fig. 2. Mockup of the Automation version of feature 13 "untag yourself from posts".

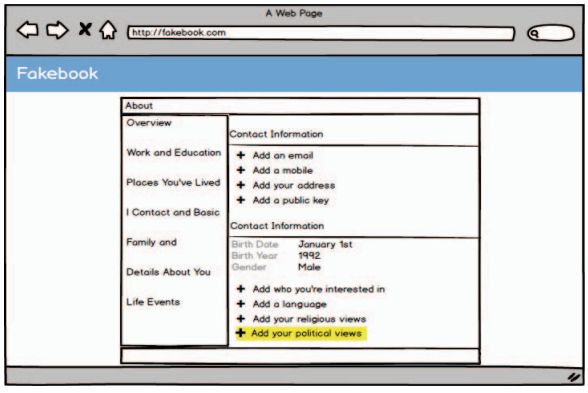

Fig. 3. Mockup of the Highlight version of feature 19: "Add/remove personal information e.g date of birth, language, political views".

Personalized and anthropomorphic agents have been shown to have beneficial effects on the acceptability of recommendations [42]. That said, the suggestions will likely be perceived as relatively intrusive: they take up space and time, potentially creating an undue *onus*. On the other hand, the explicit suggestions provide a safer alternative to the other methods, as they give the user explicit *control* over the adaptation.

C. Interview Procedure

Each interview session lasted about 45 minutes, and participants were compensated with a $5 Starbucks gift card for their time. An IRB-approved interview protocol was adopted to ensure consistency across all sessions. After obtaining informed consent from participants, the sessions were audio recorded and later transcribed. The interview with participant O was conducted remotely using video conferencing and screen-sharing, while the remaining interviews were conducted face-to-face.

After building rapport with participants and introducing them to the study, they answered two questions for each of the privacy features in Table II. The first question asked how familiar participants were with each feature (using a 5-point scale: not at all familiar–extremely familiar) and the second question asked how frequently they used each feature (5-point scale: always–never).

TABLE I. FACEBOOK PRIVACY FEATURES TESTED IN THIS STUDY.

#	Description
1	Restrict the audience that can view your photo albums
2	Block or unblock an app or game
3	Ignore future event requests from a friend
4	Block or unblock people from seeing your timeline posts
5	Place friends into custom lists
6	Turn the chat on/off
7	Add/remove your contact information
8	Restrict the audience of a post to friends on a custom list
9	Delete a post
10	Hide a post
11	Turn on/off game and app notifications and invites
12	Restrict who can look you up using your email address or phone number
13	Untag yourself from posts
14	Place friends on the "restricted" list
15	Give feedback and/or report a post
16	Limit the default audience that can view your posts
17	Restrict who can posts on your timeline, and who can see what others post on your timeline
18	Follow or unfollow a friend
19	Add/remove your personal information e.g. date of birth, languages, political views

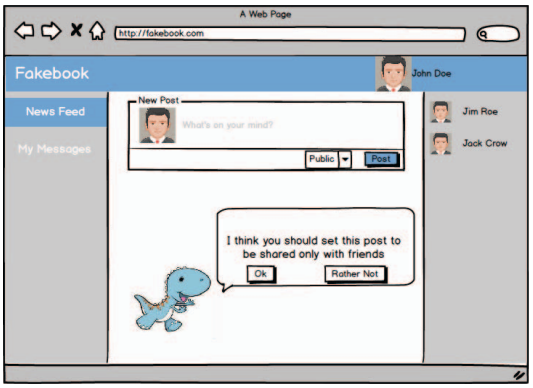

Fig. 4. Mockup of the Suggestion version of feature 8: "restrict the audience of a post to friends on a custom list".

Next, participants were presented with a paper-based user interface mockup of a randomly selected privacy feature. They were given a scenario to fully understand the use of the feature. The scenario was: *"You are John Doe from Fresno, California. You are 22 years old, and regularly use Facebook for business and leisure. You are currently looking for a job and are trying to keep a clean Facebook account. You would like to <use privacy feature> to achieve <some goal>"*.

Participants were then first shown the default non-adaptive version of the feature, and asked if they were aware of the feature, and how often they used it (on Facebook). If they had used the feature before, they would be asked for what purpose they used the feature. If they were completely unfamiliar with the feature the scenario would again be invoked to help them better understand the use of the feature.

Next, participants were shown a randomly selected adaptive version of the same feature, and asked for their opinion on the presentation, functionality, pros, cons and comfort with the adaptive feature, and the method with which it was implemented.

This procedure was repeated for a total of six times per participant. The subset of features shown to each participant is listed in the last column of Table I; we ensured that all participants encountered each adaptation method at least twice (semi-random) with a different privacy feature, and endeavored to cover all the privacy features equally among all the participants.

After completing six features, participants were given an exit survey, asking them to select their preferred adaptation method (Automation, Highlight, Suggest, or As is) for each of the features. This helped us gain a broader overview on whether participants would want to use any of these adaptive features beyond the in-depth interview, and if so, which adaptation method they would prefer. The findings for each feature are presented in Table III. Note that the exit survey was only completed by 10 of our 18 participants.

TABLE II. THE OVERALL DISTRIBUTION OF PREFERENCE FOR EACH ADAPTATION METHOD PER PRIVACY FEATURE.

Feature #	Automatic	Highlight	Suggestion	As is
1	1	0	6	3
2	2	5	3	0
3	3	2	4	1
4	0	0	6	4
5	1	1	6	2
6	0	3	3	4
7	0	3	3	4
8	0	2	5	3
9	0	0	2	8
10	2	3	3	2
11	1	3	3	3
12	2	2	3	3
13	1	3	3	3
14	0	3	5	2
15	2	5	1	2
16	1	1	4	4
17	2	1	2	5
18	1	2	5	2
19	2	1	3	4

IV. FINDINGS

Based on our analysis of the interview data, we find that Suggestion was the most preferred adaptation method, followed closely by Highlight, with Automation being the least preferred. However, we find that the preferred adaptation method for each specific feature largely depends on the user's awareness and usage of the feature, and in some cases on whether the feature results in awkward or irreversible privacy behaviors. We discuss these findings in detail below.

A. Automation

1) Automation and Frequency of Use

We find that participants generally dislike the Automation method, especially for features they never use or are unaware of. As participant M stated when shown the Automation version of the privacy feature that enables one to block app invites (feature 11 in Table II):

"I was not aware you can block game app invites because I have not explored Facebook properly. Maybe if I knew this particular feature existed, I would prefer doing it manually than automatic because you never know who is getting automatically blocked."

On the other hand, participants are more accepting of the Automation method for privacy features they use frequently, just as Participant C stated about the automatic removal of a tag (feature 13)

"It saves me a lot of time and […] effort because I do not have to look through 100 posts that all my friends have tagged me in […] In terms of situations where I am applying for a job or like applying for school or something maybe taking those precautionary measures has a certain cognitive load on me, so it kind of takes that off […]. It follows along the line of 'prevention is better than cure' […] So it kind of prevents a wrong, rather than have a wrong thing out there and then cure it. […] Better safe than sorry!"

Nevertheless, participants stressed the need for additional *control* over the automated feature, e.g., they would want to be able to turn it on or off. When shown the Automation version of the audience selector tool used to control who can see a photo album (feature 1) participant A expressed:

"I feel like it should be a choice for people to have stuff like this automated for you. I personally would not care for it because I feel it does not save you that much time and I can set my intended audience in a few seconds."

This indicates that the *ease of use* is an important reason to like the Automation method, and that the absence of cognitive load reduces the need for fully automated adaptations.

Furthermore, participants are worried about the accuracy of the adaptation for features they use only occasionally. For example, participant B, who only occasionally uses the "block people" feature (feature 4), argued:

"It means I am relying on the system to detect someone that I know needs blocking. So essentially, I am believing the system understands me perfectly. Maybe to some degree the system can learn what kind of people I block [but] I am not so sure that it's just learnable like that."

2) The Presumed Irreversibility of Automation

With the Automation method, participants are wary that the system will reduce their ability to make their own privacy decisions. Combined with the fear that the system might get their privacy preferences wrong, they worry that the Automation method will implement privacy behaviors that are irreversible, leading to persistent negative consequences. As participant A put it when shown the Automation version of the blocking app invites feature (feature 8):

"*I am kind skeptical of the automatic option that it automatically picks who's going to see this, because again, it could pick the wrong person and I do not notice, and then you are not able to know who sees the picture type of thing."*

Similarly, participant M stated about automatically blocking app invites (feature 11):

"Say you have a close friend who is into this game stuff, then automatically blocking him would not be nice, you would lose a friendship there."

Finally, participant B made a comment about automatically blocking people (feature 4) that really shows how this fear is related to a potential loss of control:

"Let's say for example, I block two people that posted something about politics, so if the system understands that ok he does not like things regarding politics. I think that's kind of assumed just because two things were related to politics. I really want to know what algorithm it uses to understand my character in terms of what kind of people I block."

While our implementation of the Automation method gives users the possibility to undo the automated action, this did not alleviate participants' concerns. Many stated that it might already be too late to undo the automated action by the time they take note of it. For example, participant I on the possibility to undo the Automation of the friends list management feature (feature 5) stated that:

"I should not have to undo. It should not do unless I tell it to. Some things cannot be undone. […] What if it assumes that this person is my friend, yet he is my boss and I happen to share an inappropriate post with that person? Now am fired! Sure, you can undo the setting, but you cannot undo the damage."

Others mentioned that having to always check to make sure the system got their preference right would only increase the cognitive load. For example, on the Automatic version of the friend list management feature (feature 5), participant I stated:

"Sure, you can undo the setting but […] doesn not that even cause more or the same amount of work? I thought the point of this was to make it easier, but this makes it harder. Now I have to go through and check to make sure all is good."

This responsibility could even spill over into their other social network activities. As Participant A stated about the Automation of the audience selector (feature 8):

"Personally, if it says 'do you want to share with friends' I would not undo [it], because most of the time I share with friends anyway. [But] if it got it wrong, it would make me be conscious about the text I post, making me read it over and over again."

3) Automation for Actions with No Consequences

Our findings suggest that Automation is only appropriate when the automated action has no big consequences for the use. As participant L stated about the Automation of hiding a post (feature 10)

"*If you are already going so far as to like make decisions about automatically hiding posts or what not which Facebook already does in the backend obviously why are u telling the user about that in the first place"*

Participants expected that the adaptation would have to be very accurate for there to be no negative consequences. For example, participant A stated that she would be comfortable with the Automation of audience selection (feature 8) if it was very advanced:

"*I guess if it could […] guess who is in the picture and what the picture is about, then you can set it to an audience. But I do not think the technology is probably there yet. It's like if a picture has 5+ people, it would probably analyze 'Oh you are at a party.' then you probably should [share it with] your friends [only].' If this was automatic I think I would rather have the automatic [version]."*

Similarly, participant B would only be comfortable with the system automatically blocking people (feature 4) if it were very accurate:

"Once I believe that the system is […] very good at understanding the kind of people I block, then I would be comfortable. But if you are asking me to use it right now, am not so sure the system knows my character very well. I would use this fully if I have proof the system is good at its job of understanding what kind of people I block. Having automatic blocking kind of gives me the assurance that I am going to look clean in the eyes of the people."

B. Highlight

1) Highlight for Unobtrusive Awareness

Participants appreciated the Highlight method for its ability to unobtrusively raise users' awareness about a privacy feature. When shown the Highlight version of the friend list management feature (feature 5) participant L stated:

"*I think it's not obstructive to seeing the rest of the screen but it also gives a visual cue to say like we recommend this choice or information*"

Similarly, when shown the Highlight version for adding/removing contact information (feature 7), participant C stated:

"*I will have to agree 100%, that it definitely makes me more aware, because otherwise I am just seeing plain text, and I do not really care about what information I am putting out there. This kind of makes me […] more aware of what I am putting out there and what it's asking me for […] It helps me be aware or control my privacy to a degree better."*

The same participant also remarked that the Highlight method is less cluttered and less *taxing* than the Suggestion method because it allows him to *"selectively choose to ignore the highlight feature."*

Also comparing Highlight to Suggestion, participant K stated regarding the "hide a post" feature (feature 10):

"if I scroll through a bunch of posts to hide, suggestions would be annoying but if I was scrolling through and it had a highlight then that would be ok. If it was highlighted yellow or something then that would draw my attention."

2) A U-shaped Relation with Familiarity

We find that participants' preference for Highlight depends on their familiarity with the privacy feature. On the one hand, expert users of a certain privacy feature may find Highlight a redundant adaptation method, and prefer full Automation instead. For example, participant A regarding the reporting of a post as spam (feature 15) stated that:

"*It's a redundant adaptation to have. I understand that your trying to raise awareness that 'oh this is the spam button,' but […] if you wanted to report it in the first place then you would report it as spam, but if you did not want to mark it as spam regardless then you would not.*"

On the other hand, participants could easily get confused with the Highlight method if they are unfamiliar with a privacy feature, resulting in a perceived loss of *control*. They are instead more likely to prefer a Suggestion that provides some more information. When shown the Highlight version of the friend list management feature (feature 5) participant L stated this downside:

"It cannot really show a justification for why it is being highlighted over something else […] I mean that's less information being given to the user."

In sum, our findings suggest that there is a nuanced U-shaped relationship between the participant's familiarity with a privacy feature and their preference for the Highlight method: while it may be redundant for expert users of the feature, and confusing for novice users, it unobtrusively provides an optimal level of awareness to those who occasionally use the feature.

C. Suggestion

1) Convenience or a Nuisance?

All but one of the participants prefer the Suggestion method for at least one of the presented privacy features. Like Highlight, Suggestion raises users' awareness about the privacy feature. For example, participant E stated the following about the Suggestion version of the adaptive "restricted" list feature (feature 14):

"yes, restricted lists are not what I always think of. I think of blocking more than restricted list and thus having the suggestion pop-up brings it more to mind."

Suggestions are convenient, because they provide a shortcut to the functionality. Participant J was shown the Suggestion version of the adaptive untag feature (feature 13) and he stated:

"I do not have to go through the settings[..]I do not have to click the drop down and find anything in the settings menu. I am given a clear choice about the tag to either keep it or remove [..] It really focuses me in on the thing that might be important."

Similarly, when shown the Suggestion version of the hide post feature (feature 10), participant C brought up both increased awareness and convenience benefits:

"I would definitely save a lot of time because this would pop up and I would click "ok" for the posts that I do not care

about [...] and it will take care of all similar posts. It's going to catch my attention more than a hide icon."

Several participants appreciated the idea of getting privacy advice from a virtual character. For example, when participant A was shown the Suggestion version of limiting the default audience that can view one's posts (feature 16), she expressed:

"I think it's a cute dinosaur for starters. It really does grab your attention because if I am about to post and something like this pops up, I am definitely going to look at it [...] so it reminds you before you post."

Similarly, when shown the Suggestion version of the feature that turns game and app notifications and invites on/off (feature 11), participant C stated:

"I prefer the suggestion, because comics and pictorial representations are more than just text […] comically depicted speech bubbles kind of engage my mind and bring my immediate focus and attention into this […] it's drawing me towards fixing the need of the hour."

On the other hand, some participants did not like the virtual character, suggesting it was somewhat childish, and not serious enough for the topic of privacy. For example, participant F commented on the Suggestion version of the follow/unfollow a friend feature (feature 18):

"It looks like a blue bunny almost […] It's a little childish I guess [...] I think a little more professional presentation would be in order."

Some participants also suggested ways in which the virtual character could be improved or made better e.g. participant A, feature 16:

"I guess it would be better if you can have an option to change or customize it to maybe something like a privacy dog or a self-resembling avatar to [make it] seem like I am reminding myself."

Furthermore, we find that participants tended to dislike the Suggestion method for features they use frequently. This is because too many suggestions require considerable attention from the user to successfully be dealt with. As participant C continues to explain about the hide post feature (feature 10):

"I do not want to see more than two of these at a particular instance for two consecutive posts […] It gets repetitive […] I do not want to see a suggestion saying the exact same thing on three consecutive posts, even if those posts are things that I do not care about."

Similarly, participant F commented about the follow/unfollow feature (feature 18):

"It would be okay if it was every once in a while. I really do not want it to be like 'oh you should do this this this and this!' [...] But I think [I would like it] if every once in a while, it was like 'you have not spoken to this person in three years maybe u should unfollow'."

2) An Opportunity for Explanation

For features that participants are unfamiliar with, Suggestion has an added advantage: the opportunity to explain the privacy feature and the adaptation to the user. These explanations give users a reason for the Suggestion, thereby actively helping them learn something about Facebook privacy. Combined with the ability to either follow or ignore the suggestion, such explanations may help users feel more in control of their privacy. For example, when participant H was shown the "follow or unfollow a friend" privacy feature (feature 18), he stated:

"I think it would be helpful if it gave a reason, the tough part is counting on the person to follow through. But I think people value their privacy and I think it would be successful because there [are] lots of fake accounts."

Similarly, when participant I was shown the privacy feature that restricts who can look him up using his email address or phone number (feature 17), he stated:

"I feel like if you are going to suggest something to me, you should give me a reason."

3) Awkward Suggestions and Social Norms

Our findings show that suggestions can break certain social norms, especially when applied to private behaviors that carry a negative social perception, such as deleting posts and unfollowing users. For example, when participant I was shown the "follow or unfollow a friend" privacy feature (feature 18), he stated:

"I might like someone's posts a lot but not follow them. Thus the system can suggest that I follow them based on those likes. However it should not make a suggestion that I unfollow anyone, because common sense dictates [that you should not suggest to me to unfollow people]."

Similary, participant L stated about the deletion of a post (feature 9):

"I do not want Facebook to suggest what I should delete because that would be a weird decision to make for me."

Indeed, some participants mentioned that the Privacy Dinosaur adds to the awkwardness of Suggestions that carry a negative social perception. For example, when participant I was shown a privacy dinosaur that came with the suggestion to restrict who can look him up using his email address or phone number (feature 17), he stated:

"Why is the dinosaur giving me a suggestion and not just insight? […] I do not think the dinosaur should suggest. I think the dinosaur should just give me options, or [tell me] what different options do or something […] But giving me a suggestion without actually giving me reason why it's suggesting would probably be a reason I would be uncomfortable [with it], because I would feel like this dinosaur knows more than it's giving me information about."

This comment also suggests that explanations can potentially reduce the awkwardness of Suggestions by carefully explaining the reasoning behind them. Without explanations, though, certain suggestions felt unsolicited or even rude. As participant H expressed about the Suggestion to turn on/off game and app notifications and invites (feature 11):

"The system could notice how much I have been clicking 'NO'. It would then be helpful to have a suggestion that says, 'we noticed you say NO a lot, do you want to block the app invite?' It's kind of a call to action I guess."

D. No Adaptation

1) No Adaptation Rather Than a Different Method

We have already discussed several situations where participants preferred the traditional untailored privacy features to our user-tailored alternatives. This preferrence was most pronounced for seemingly irreversible actions (especially when participants saw such features paired with the Automation method, e.g. participant I, feature 5: "Sure, you can undo the setting, but you cannot undo the damage") and for actions with a negative social perception (especially when participants saw such features paired with the Suggestion method, e.g. particiapant L, feature 9: "I do not want Facebook to suggest what I should delete because that would be a weird decision to make for me.") In both cases, participants did not prefer a different adaptation method, but rather opted for no adaptation at all.

2) The User Is the Best Adaptation Algorithm

Beyond this, the preference for 'no adaptation' also seemed to correlate with participants' trust in the system's ability to learn the user's preference (again, this was most pronounced when participants saw the Automation method, e.g. participant B, feature 4: "It means I am relying on the system to detect someone that I know needs blocking"), and finally, this preference seemed to correlate with participants' familiarity with the privacy feature. For example, when shown the Suggestion to have the chat feature turned off (feature 6), Participant J stated:

"I feel like if I turn off the chat it's because I want to be temporarily without notifications and I will come back and turn it back on later. But I think more likely I will just put my phone on silent [...] I want chat all the time—like, that's my main use of Facebook. I would not want some automatic process to turn it off. And if it suggested I turn it off, I would not listen."

In sum, when participants distrusted the algorithm behind a certain adaptive privacy feature, or when they were already intimately familiar with the privacy feature, they essentially considered themselves to be a better adaptation algorithm than the system. Hence, in these cases they preferred the traditional untailored version of the privacy feature.

V. DISCUSSION

Our findings summarized in Table IV answer and shed an interesting light on our research questions. We find that the preferred adaptation method for the different privacy features depends on users' awareness and usage of those features (RQ2). Since different Facebook users are (un)familiar with different features, this means that the preferred adaptation method for each feature differs per user. The adaptation method itself should thus be tailored to the user as well.

TABLE III. PREFERRED ADAPTATION METHODS GIVEN ADAPTATION EFFECTS AND USER PRIVACY FEATURE AWARENESS OR USAGE

Awkward/ Irreversible?	**Awareness/Usage?**		
	Unfamiliar/ Do not use	***Occasional use***	***Frequent use***
Yes	As is	Highlight	As is
No	Suggestion	Highlight	Automation

Moreover, we find that the preferred adaptation method may sometimes not be suitable, in which case users end up preferring the untailored version (RQ1). This limits the extent to which user-tailored privacy can be implemented on Facebook.

A. Unfamiliar/Infrequently-Used Features

When Facebook users are unfamiliar with a privacy feature, they prefer the Suggestion method, mainly because our implementation of Suggestion (the "Privacy Dinosaur") allows for the adaptive behavior to be explained. The infrequent use and unfamiliarity makes the *load* of using them more cognitive rather than physical. With proper explanation, the Suggestion method actually reduces this load.

Moreover, its superior level of *control* turns the Suggestion method into a 'privacy education' tool that introduces users to a privacy feature they were previously unaware off. Normally, introducing users to a new privacy feature can be daunting or confusing: because the user is unfamiliar with the feature, they may not know how to interact with it (for example: if the user has never 'blocked' another user, they may not know when it would be appropriate to do so). The adaptive behavior solves this problem, though, by not only introducing the feature to the user, but also suggesting to the user how to interact with it, thereby reducing the cognitive load. In effect, the adaptive nature of the Suggestion makes it a very accessible tool for education.

However, users do *not* prefer the Suggestion method when it gives the wrong suggestions that they are likely to find awkward such as blocking a friend. Such a suggestion is considered to be against the norm of social interaction. Therefore, rather than opting for one of the other adaptation methods (which lack the desired explanation of the adaptive behavior), users prefer that the privacy feature remains untailored.

B. Occasionally-Used Features

When Facebook users use a feature occasionally, they may prefer the Highlight method. This preference is mainly a compromise: Suggestion would significantly be a destruction for features that are used with some regularity (the privacy dinosaur would show up too frequently), while Automation would significantly reduce *control* (users are not familiar enough with these features to comfortably allow the system to take over altogether).

C. Frequently-Used Features

When users use a feature frequently, users prefer Automation, suggesting that they are willing to give up some *control* in return for a reduction in the *effort required for proper* privacy management. Frequent users already know what to do with a feature, so their main effortful load is rather physical than cognitive. In effect, neither Highlight nor Suggestion would sufficiently reduce this load. Moreover, users seem to have an intuitive understanding that their frequent use of a feature likely improves the quality of the adaptive behavior. This gives them a certain amount of 'indirect' control over the Automation method.

However, users do *not* prefer the Automation method when the resulting automated privacy decision feels

irreversible. For example, Facebook users would not appreciate the system automatically unfriending or blocking their friends, deleting their posts and setting their post audiences. Even though our implementation of Automation provides a clear mechanism to 'undo' the decision, making every decision technically reversible, users are uncomfortable when a system automatically implements a decision that 'feels' irreversible without asking the user.

VI. Design Implications

We offer the following insights for social network designers interested in implementing user-tailored versions of privacy features. While our study focuses on the Facebook platform, we argue that our insights are sufficiently generic to also apply to other social networks (or indeed, other information systems in general).

A. Selectively Automate Privacy Features

Our findings suggest that designers can automate privacy features to relieve some of the user responsibility in privacy decision-making. However, they are advised to only do so for features that users use frequently, and to avoid automating any privacy behaviors that are perceived as having irreversible consequences. Given the large variation in privacy feature usage among Facebook users [3], this means that the selective Automation of privacy feature should itself be tailored: the system should find out which features each user frequently uses, and only automate those features.

Accuracy is of utmost importance when fully automating privacy features: Many participants in our study portrayed a lack of trust in the system's ability to accurately tailor its privacy settings to their preferences. Unless the underlying algorithm is extremely accurate, users will likely believe that they themselves are much better at managing their own privacy (even though research shows this often not to be the case! [4], [15], [16]).

B. Selectively Apply Highlights

Designers can use subtle highlights recommending certain privacy behaviors as a means to assist users in making better decisions, but also to help raise their awareness of certain privacy features that they may have forgotten about. Designers can capitalize on the subtle awareness-raising capabilities of this adaptation method by using it primarily for privacy features that users only use occasionally. Again, this means that the application of the Highlight method should itself be tailored to the user.

C. Selectively Make Suggestions

Facebook already has a Privacy Dinosaur that makes privacy-related suggestions, so designers have the opportunity to leverage this functionality to make adaptive privacy suggestions or design a similar virtual character for other social networks/information systems.

The virtual character should be designed not only to make privacy recommendations, but also to explain those recommendations: several participants in our study suggested—unprompted—to include explanations of the adaptive behavior in the dinosaur's suggestion. Designers should avoid the potential awkwardness of suggesting privacy behaviors with negative social connotations (e.g. blocking or unfollowing people), though. That said, a good explanation can alleviate some of these concerns.

The opportunity for explanations also makes the Suggestion method particularly useful for introducing the user to privacy features they are unfamiliar with. Again, this means that the application of the Suggestion method should be tailored to the user's awareness of the various privacy features.

VII. Limitations and Future Work

An obvious limitation of our study is that our adaptive privacy features were mere paper mockups, using cartoon-style renderings with less visually distracting features as compared to the actual Facebook. This might have given them a less realistic appearance, but also made it easier for the participants to concentrate on the presented adaptation mechanism and envision the use of the adaptive privacy features without getting hung up on design details. Moreover, whereas in real life such adaptive features would likely make the occasional mistake, our presented scenarios assumed that the adaptation methods presented to participants worked with 100% accuracy. That said, participants questioned the idea that the adaptive system would always get their privacy preferences right, and frequently brought this up as a potential reason to prefer the traditional untailored privacy features. As such, the potentially reduced accuracy of the presented adaptations in real-world systems is likely to significantly impact users' perceptions and may result in a reduced preference for adaptive privacy functionality.

On the other hand, we note that most existing work on adaptive privacy features evaluates their accuracy only, without testing the user experience of the resulting system or the usability of the mechanism by which the privacy recommendations are presented to the user (Liu et al. [12] and Knijnenburg and Jin [35] are notable exceptions). Our paper demonstrates that the method by which the recommendations are presented has a strong influence on the user experience. Hence, we encourage researchers and developers of adaptive privacy features to conduct usability and user experience tests.

Our study design relied on users' self-reported evaluations of the paper-based mockup designs we showed them. While this allowed users to critically reflect upon the consequences of the user-tailored functionality and the three adaptation methods, users did not have the opportunity to interact with the privacy features in a social network interface. This precludes us from making strong claims about the usability of the adaptation methods, and it may even mean that users' preferences for these methods change once they have the opportunity to interact with them. Thus, future research should explore the usability of different adaption mechanisms in an interactive test environment.

We also limited ourselves to a subset of prominent Facebook privacy features as previously identified by Wisniewski et al. [3]. They cover only a limited subset of the available privacy features and are restricted to the features on the Facebook platform. That said, we made sure that the selected features span the various "boundary protection mechanisms" covered in [2]—a work that also demonstrates that these mechanisms exist in various forms across a variety of social network sites.

Despite these limitations, the answers to our research questions constitute a clear pattern of user preferences, with Table III mapping out which situations call for adaptive privacy features, and which adaptation method would likely be preferred. We argue that these insights are sufficiently generic to apply to any social network site, or indeed any information system that may benefit from adaptive privacy features. In future work, researchers, developers and designers can leverage these insights for the development adaptive privacy features in research prototypes or real-world social networking sites.

VIII. Conclusion

In this work, we contribute to the advancement of User-Tailored Privacy (UTP) by studying users' opinions on adaptive Facebook privacy features, as well as three potential adaptation methods (Automation, Highlight, and Suggestion) that can be used to adapt Facebook's privacy features to the user's personal preferences. We find that participants generally dislike the full Automation method, except for privacy features they use frequently and perceive as inconsequential, where it can alleviate some of the behavioral onus and effort of managing one's privacy. The Highlight method is appreciated for its ability to unobtrusively raise users' awareness about a privacy feature and is thus most suitable for features users only use occasionally. Finally, the Suggestion method is preferred as a means to teach users privacy features they are unfamiliar with, unless this results in awkward suggestions of behaviors with negative social connotations.

As the familiarity with and usage of the various privacy features differs extensively per user, we argue that the choice of adaptation method itself needs to be tailored to the user as well. Overall, our results demonstrate the viability of UTP, and we believe that our insights will help researchers, designers and developers in their future endeavors developing user-tailored privacy interfaces and experiences.

Acknowledgements

This research was supported in part by a Facebook Emerging Scholar Award and by the DoD Award W911QY-16-C-0105. The authors would like to thank the participants in this research for their tremendous contribution.

References

[1] P. Wisniewski, A. K. M. N. Islam, B. P. Knijnenburg, and S. Patil, "Give Social Network Users the Privacy They Want," in Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, Vancouver, Canada, 2015, pp. 1427–1441.

[2] P. Wisniewski, A. K. M. Islam, H. R. Lipford, and D. Wilson, "Framing and Measuring Multi-dimensional Interpersonal Privacy Preferences of Social Networking Site Users," Communications of the Association for Information Systems, vol. 38, no. 1, Jan. 2016.

[3] P. J. Wisniewski, B. P. Knijnenburg, and H. R. Lipford, "Making privacy personal: Profiling social network users to inform privacy education and nudging," International Journal of Human-Computer Studies, vol. 98, pp. 95–108, Feb. 2017.

[4] M. Madejski, M. Johnson, and S. M. Bellovin, "A study of privacy settings errors in an online social network," in Fourth International Workshop on Security and Social Networking, Lugano, Switzerland, 2012, pp. 340–345.

[5] R. Compañó and W. Lusoli, "The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas," in Economics of Information Security and Privacy, New York, NY, 2010, pp. 169–185.

[6] B. P. Knijnenburg, "Privacy? I Can't Even! Making a Case for User-Tailored Privacy," IEEE Security Privacy, vol. 15, no. 4, pp. 62–67, 2017.

[7] P. Wisniewski, H. Lipford, and D. Wilson, "Fighting for my space: coping mechanisms for SNS boundary regulation," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, New York, NY, USA, 2012, pp. 609–618.

[8] H. Harkous, K. Fawaz, R. Lebret, F. Schaub, K. G. Shin, and K. Aberer, "Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning," EPFL, Feb. 2018.

[9] C. Dong, H. Jin, and B. P. Knijnenburg, "PPM: A Privacy Prediction Model for Online Social Networks," in Proceedings of The International Conference on Social Informatics, 2016, pp. 400–420.

[10] L. Fang and K. LeFevre, "Privacy Wizards for Social Networking Sites," in Proceedings of the 19th International Conference on World Wide Web, New York, NY, USA, 2010, pp. 351–360.

[11] J. Lin, B. Liu, N. Sadeh, and J. I. Hong, "Modeling Users' Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings," in Symposium on Usable Privacy and Security, 2014.

[12] B. Liu et al., "Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions," in Twelfth Symposium on Usable Privacy and Security, Denver, CO, 2016, pp. 27–41.

[13] R. Ravichandran, M. Benisch, P. Kelley, and N. Sadeh, "Capturing Social Networking Privacy Preferences," in Privacy Enhancing Technologies, 2009, vol. 5672, pp. 1–18.

[14] B. P. Knijnenburg, E. M. Raybourn, D. Cherry, D. Wilkinson, S. Sivakumar, and H. Sloan, "Death to the Privacy Calculus?," in Proceedings of the 2017 Networked Privacy Workshop at CSCW, Portland, OR, 2017.

[15] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, "Analyzing facebook privacy settings: user expectations vs. reality," in Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement, New York, NY, USA, 2011, pp. 61–70.

[16] K. Strater and H. R. Lipford, "Strategies and struggles with privacy in an online social networking community," in Proceedings of the 22nd British HCI Group Annual Conference on People and Computers, Swinton, UK, 2008, pp. 111–119.

[17] A. Acquisti and J. Grossklags, "What Can Behavioral Economics Teach Us About Privacy?," in Digital Privacy: Theory, Technologies, and Practices, A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis, Eds. New York/London: Auerbach Publications, 2008, pp. 363–377.

[18] L. Church, J. Anderson, J. Bonneau, and F. Stajano, "Privacy Stories: Confidence in Privacy Behaviors Through End User Programming," in Proceedings of the 5th Symposium on Usable Privacy and Security, Mountain View, CA, 2009.

[19] P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor, "Standardizing privacy notices: an online study of the nutrition label approach," in Proceedings of the 28th International Conference on Human Factors in Computing Systems, Atlanta, Georgia, 2010, pp. 1573–1582.

[20] H. R. Lipford, A. Besmer, and J. Watson, "Understanding Privacy Settings in Facebook with an Audience View," in Proceedings of the 1st Conference on Usability, Psychology, and Security, Berkeley, CA, USA, 2008.

[21] F. Raber, A. D. Luca, and M. Graus, "Privacy Wedges: Area-Based Audience Selection for Social Network Posts," in Twelfth Symposium on Usable Privacy and Security, Denver, CO, 2016.

[22] H. Nissenbaum, "A Contextual Approach to Privacy Online," Daedalus, vol. 140, no. 4, pp. 32–48, Oct. 2011.

[23] S. Barocas and H. Nissenbaum, "On notice: The trouble with Notice and Consent," in Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information, 2009.

[24] D. J. Solove, "Privacy Self-Management and the Consent Dilemma," Harvard Law Review, vol. 126, pp. 1880–1903, 2013.

[25] Y. Wang, G. Norcie, S. Komanduri, A. Acquisti, P. G. Leon, and L. F. Cranor, "'I regretted the minute I pressed share': a qualitative study of regrets on Facebook," in Proceedings of the Seventh Symposium on Usable Privacy and Security, Pittsburgh, PA, 2011, pp. 10:1–10:16.

[26] C. Dong, H. Jin, and B. P. Knijnenburg, "Predicting Privacy Behavior on Online Social Networks," in Ninth International AAAI Conference on Web and Social Media, 2015, pp. 91–100.

[27] D. Wilkinson et al., "User-Tailored Privacy by Design," in Proceedings of the Usable Security Mini Conference, San Diego, CA, 2017.

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

[29] J. Cranshaw, J. Mugan, and N. Sadeh, "User-Controllable Learning of Location Privacy Policies with Gaussian Mixture Models," in Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, San Fancisco, CA, 2011, pp. 1146–1152.

[30] E. Toch et al., "Empirical models of privacy in location sharing," in Proceedings of the 12th ACM international conference on Ubiquitous computing, Copenhagen, Denmark, 2010, pp. 129–138.

[31] J. Xie, B. P. Knijnenburg, and H. Jin, "Location Sharing Privacy Preference: Analysis and Personalized Recommendation," in Proceedings of the 19th International Conference on Intelligent User Interfaces, New York, NY, USA, 2014, pp. 189–198.

[32] B. Liu, J. Lin, and N. Sadeh, "Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?," in Proceedings of the 23rd International Conference on World Wide Web, Republic and Canton of Geneva, Switzerland, 2014, pp. 201–212.

[33] S. Wilson et al., "Privacy manipulation and acclimation in a location sharing application," in Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing, Zürich, Switzerland, 2013, pp. 549–558.

[34] P. Bahirat, Y. He, A. Menon, and B. Knijnenburg, "A Data-Driven Approach to Developing IoT Privacy-Setting Interfaces," in 23rd International Conference on Intelligent User Interfaces, Tokyo, Japan, 2018, pp. 165–176.

[35] B. P. Knijnenburg and H. Jin, "The Persuasive Effect of Privacy Recommendations," in Twelfth Annual Workshop on HCI Research in MIS, Milan, Italy, 2013.

[36] E. J. Johnson, S. Bellman, and G. L. Lohse, "Defaults, Framing and Privacy: Why Opting In ≠ Opting Out," Marketing Letters, vol. 13, no. 1, pp. 5–15, 2002.

[37] Y. Wang, P. G. Leon, A. Acquisti, L. F. Cranor, A. Forget, and N. Sadeh, "A Field Trial of Privacy Nudges for Facebook," in Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems, Toronto, Canada, 2014, pp. 2367–2376.

[38] B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, and C. Newell, "Explaining the user experience of recommender systems," User Modeling and User-Adapted Interaction, vol. 22, no. 4–5, pp. 441–504, 2012.

[39] T. B. Sheridan and W. L. Verplank, "Human and Computer Control of Undersea Teleoperators," MIT Cambridge Man-Machine Systems Lab, Jul. 1978.

[40] T. B. Sheridan, "Human centered automation: oxymoron or common sense?," in IEEE International Conference on Systems, Man and Cybernetics. Intelligent Systems for the 21st Century, 1995, vol. 1, pp. 823–828.

[41] S. Vihavainen, A. Lampinen, A. Oulasvirta, S. Silfverberg, and A. Lehmuskallio, "The Clash between Privacy and Automation in Social Media," IEEE Pervasive Computing, vol. 13, no. 1, pp. 56–63, Jan. 2014.

[42] B. P. Knijnenburg, "A user-tailored approach to privacy decision support," Ph.D., University of California, Irvine, United States -- California, 2015.