User-Tailored Privacy by Design

Saadhika Sivakumar, Daricia Wilkinson, David Cherry, Bart P. Knijnenburg
Clemson University
Clemson, SC, USA
ssivaku@clemson.edu, dariciw@clemson.edu, dcherry@clemson.edu, bartk@clemson.edu

Elaine M. Raybourn
Sandia National Laboratories*
/ADL Initiative
Orlando, FL, USA
emraybo@sandia.gov

Pamela Wisniewski
The University of Central Florida
Orlando, FL, USA
pamwis@ucf.edu

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

ABSTRACT
The "privacy by design" philosophy addresses privacy aspects early in the design and development of an information system. While privacy by design solutions often provide considerable advantages over "post hoc" privacy solutions, they are usually not customized to the needs of individual users. Further, research shows that users differ substantially in their privacy management strategies. Thus, how can we support such broad privacy needs in a comprehensive and user-centered way? This paper presents the idea of *user-tailored privacy by design*, a design methodology that combines multiple privacy features into a single intelligent user interface. We discuss how this methodology moves beyond the "one-size-fits-all" approach of existing privacy by design solutions and the narrow focus on information disclosure of existing user-tailored privacy solutions. We illustrate our approach through an implementation of user-tailored privacy by design within Facebook based on six privacy management profiles that were discovered in recent work, and subsequently extend this idea to the context of the Total Learning Architecture (TLA), which is a next generation learning platform that uses pervasive user monitoring to provide highly adaptive learning recommendations.

Author Keywords
Privacy; Social Networking Sites; Training Systems; Design; Personalization

INTRODUCTION
Privacy by design (or PbD; see [9] for an overview) is a design philosophy in which privacy aspects are addressed early in the system design and development process, rather than after the system has been developed ("post hoc privacy"). While post hoc privacy solutions typically try to mitigate privacy problems that exist within a system, privacy by design tries to avoid problems from occurring at all [46]. Some recent criticisms of the PbD philosophy is that some of the principles are simply too vague to implement in practice [49] and that, while putting privacy at the forefront of design, PbD does not address variations in the privacy needs of all users.

Research shows that users differ substantially in the strategies they use for managing their online privacy [63–66]. Therefore, a key research question posed by this research is whether it possible to move the PbD philosophy beyond the a one-size-fits-all approach to provide more tailored support for these different privacy management strategies? In this paper we propose *User-Tailored Privacy by Design*, a design methodology that combines multiple privacy features (such as withholding information, blocking, and selective sharing) in an intelligent system that can tailor these features to best support users' preferred privacy management strategies.

Building upon recent work that identifies six Facebook privacy management profiles [63], our work describes Facebook re-designs for each profile, and suggest a way in which these re-designs can be adapted to the user's profile on the fly. Furthermore, in an effort to extend these findings and ideas to a domain beyond Facebook, we also apply user-tailored Privacy by design to develop guidelines for the "Total Learning Architecture" that is being developed by the Advanced Distributed Learning Initiative [44].

Our work concludes with a discussion of methods for discovering the user's privacy management profile, as well as alternative adaptation strategies that attempt to move users beyond their current strategy.

RELATED WORK
In this section we cover existing research focused on networked privacy, privacy by design, and user-tailored privacy, and identify the gaps in this research that our work attempts to cover.

Managing Networked Privacy

Networked privacy is a complex topic that has broad implications for all users, ranging from influencing their usage and acceptance of various online platforms, such as Social Networking Sites (SNS) [8,12,16,62], to altering their intended interactions with others as well as their outcomes or goals [62]. Given the impetus privacy has on outcomes end users want to achieve (or risks they prefer to avoid), it is reasonable that users and researchers alike devote considerable discourse to the topic of privacy protection [5,33,41,51]. Even though users often report being highly concerned about their privacy [6,12,16], many users still seem to misunderstand their own privacy settings [51], while others continue to use online platforms despite expressing negative privacy experiences [6].

Alternatively, even when many online platforms give users the ability to maintain their preferred privacy settings [24], users do not always exercise this option in way that is consistent with their self-reported desires [33]. For instance, Facebook gives comprehensive privacy control to users, but users rarely take advantage of all of the privacy features available to them [11,63]. Yet, researchers continue to attempt to alleviate users' privacy concerns by trying to give users more *control* over what data they wish to share, and by providing them with more *information* about the implications of their decisions [1,5,36,41,53]. These researchers have argued that such control and transparency mechanisms empower users to regulate their privacy at their desired levels [10,33,67], especially when these mechanisms are carefully integrated into the system and support a plethora of privacy management strategies [65]. However, the complexity of most socio-technical systems makes increasing transparency and control an unwieldy solution; for instance, Facebook's privacy controls have been labeled "labyrinthian" by Consumer Reports [12].

As such, networked privacy researchers continue to try to find a mag cic bullet; some researchers have explored using *privacy nudging* to relieve some of the burden of privacy decision-making from users. Carefully designed nudges make it easier for people to make the right choice, without limiting their ability to choose freely [54]. Example nudges include justifications [2,7,27,40], defaults [2,23,26,31], sentiment and audience feedback, and timers [55,56]. A problem with these nudges is that their "one-size-fits-all" approach makes normative assumptions about the value of privacy [49], taking a paternalistic stance that implicitly reduces users' control over their privacy settings [52].

Privacy by Design

The question remains: How should privacy solutions for networked technologies address the complexity of privacy control, without falling into the trap of overly paternalistic nudges? A popular solution to this problem is *Privacy by Design* (PbD), a set of design principles revolving around the idea that it is better to build privacy into the core functionality of a system, rather than adding information and control mechanisms at the end of the development process. Existing privacy by design implementations demonstrate that building privacy into the core of a system allows users to protect their privacy in more diverse and more intuitive ways than a traditional "sharing matrix" in which users specify who gets to see what [20,32,38].

While PbD has gained considerable interest among privacy researchers, it is not without criticism. One critique is that the methodology often too abstract and thus hard to implement in a specific context [50]. Part of this critique relates to the lack of a broader integration with other considerations that need to be addressed in the software development cycle [48], but an arguably more important part of it relates to the difficulty of creating a design solution that is suitable for *all* users of a system. In fact, one of the most cited results in privacy research is the finding that people differ extensively in their desire for privacy [17,18,58,59]; yet, this is not addressed by PbD. Worse yet, research shows that users' disclosure behavior is multi-dimensional [28] (i.e., users differ not just in the *amount* of information that they disclose, but also in the *kind* of information that they are most and least likely to disclose), and that they employ inherently different strategies to limit their disclosure [65]. Given this heterogeneity of users' privacy preferences [47], a *user-tailored* approach to privacy is preferred.

User-Tailored Privacy

In contrast to the nudging and PbD research, user-tailored privacy solutions (for an overview see [25]) acknowledge the wide variety in users' privacy preferences, and attempt to automatically tailor the privacy settings of the system to fit these preferences. While user-tailored privacy is a nascent area of research, several researchers have demonstrated its potential benefit. In the area of location sharing, Ravichandran et al. [42] demonstrated that a small number of default policies can accurately capture most users' location-sharing preferences. Similarly, in the area of smartphone app permissions, Liu et al. [35] show that three profiles may be sufficient to capture users' permission preferences (they later developed an approach with 7 profiles [34]). Finally, in an SNS context, Fang and LeFevre [13] demonstrate how a "privacy wizard" can simplify privacy settings in a way that is simple to understand and use, while Watson et al. [57] find that using multiple default settings does not significantly improve their fit beyond a single, optimized default setting.

Beyond profiles, some existing work has developed and evaluated more advanced, *personalized* techniques to predict the privacy settings of each individual user. Sadeh et al. [45] use a *k*-nearest neighbor approach to predict location-sharing preferences, while Pallapa et al. [39] leverage users' interaction history to determine the privacy required in future user-to-user sharing situations.

In summarizing these works, a common theme emerges in that most of the user-tailored privacy research mainly focuses on personalized approaches to managing information

disclosures and/or selective information sharing through friend lists or circles. This work thus ignores the fact that users of systems developed using the privacy by design philosophy have the ability to employ privacy management behaviors that go beyond selective information sharing. For example, in the case of SNSs, Wisniewski et al. [24,61] demonstrate that users can also manage their privacy in terms of relational boundaries (e.g. friending and unfriending), territorial boundaries (e.g., untagging or deleting unwanted posts by others), network boundaries (e.g. hiding one's friend list from others), and interactional boundaries (e.g. blocking other users or hiding one's online status to avoid unwanted chats). These privacy behaviors extend beyond the sharing matrix—they are enabled in Facebook's interface by a variety of designed privacy features.

Indeed, Wisniewski et al.'s subsequent work on these privacy behaviors demonstrates that users substantially differ in the extent to which they use each behavior [63–66]. Consequently, we argue that user-tailored privacy should move beyond user-tailored settings for managing information disclosure, towards tailoring the design of the interface itself.

In traditional user-tailored privacy, adaptation is applied exclusively to the "sharing matrix"—the specification of what should be shared with (and/or withheld from) whom. Once the profiles have been determined, the implementation of the adaptation is somewhat trivial: it is merely a user-tailored specification of the settings in the sharing matrix. Tailoring the design of the interface itself is a much more complex matter, which requires aspects taken from nudging, such as hiding, highlighting, or improving the accessibility of privacy by design functionality in line with each user's unique privacy management strategy.

In this paper we attempt to take on this task using "user-tailored privacy by design" (UTPbD); a design methodology that combines the positive aspects of nudging and transparency and control. Specifically, we acknowledge the evidence that different users learn and interact with SNSs differently [65,66], leverage existing research that structures these different privacy management strategies [63], and use nudging (by changing the salience and defaults of certain privacy controls) as a means to support these strategies.

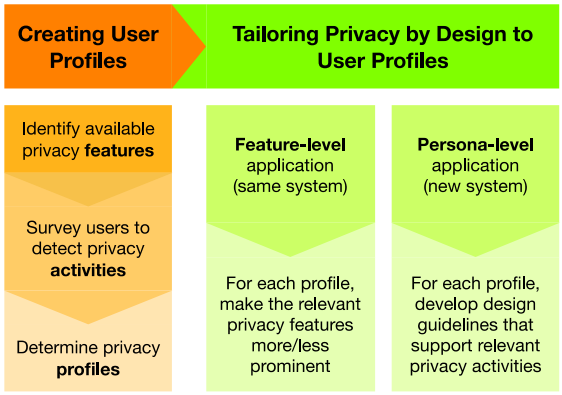

Figure 1: The User-Tailored Privacy by Design framework

We apply our UTPbD framework to by creating user-tailored Facebook redesigns for an existing classification of Facebook users [63]. Furthermore, we use the same classification as personas in the development of privacy features for the Total Learning Architecture (TLA).

USER-TAILORED PRIVACY BY DESIGN FRAMEWORK

In this section, we develop a generalizable framework for the implementation of user-tailored privacy by design that researchers and practitioners can use as a design methodology for their own systems (see Figure 1). This design methodology combines privacy by design and user-tailored privacy in an attempt to solve the shortcomings of each of these individual approaches to privacy. Specifically, we introduce adaptiveness to privacy by design, thereby moving beyond its one-size-fits-all nature, and we apply user-tailored privacy to a wider set of privacy features, thereby moving beyond its focus on selective information sharing. Our design methodology consists of two steps: creating user profiles and tailoring privacy by design to these profiles. The profiling step is adapted from Knijnenburg et al. [28,30]; the tailoring step is an original contribution. Each step is explained in more detail below.

Creating User Profiles

The first step is to develop a classification of users, resulting a set of *privacy profiles* (Figure 1, left). Profiling is an increasingly popular practice in the field of usable privacy [3,4], and we adapt the methods developed by Knijnenburg et al. [28,30] to create these profiles.

User profiling starts by identifying the privacy features available in the application. Researchers should be careful to not just focus on features that give users control over their disclosure boundaries (i.e. the sharing matrix), but also the features that given them control over relational, territorial, network, and interactional boundaries [24,61]. Identifying a broad variety of privacy features will elevate the final solution beyond user-tailored sharing settings to an actual user-tailored privacy by design solution.

The next activity is to survey users of the application regarding their use of these privacy behaviors. Ideally, a usage extent (e.g. "How often do you use this feature?" — 1 = Never, 7 = Always) is measured for each feature. The answers to the survey are then submitted to an Exploratory and Confirmatory Factor Analysis procedure in order to reduce their dimensionality (see [28,63] for technical details). This creates a higher-level set of "privacy activities" (e.g. on Facebook: Timeline moderation), each consisting of a number of related privacy behaviors (e.g. deleting content from one's Timeline, hiding a story on one's Timeline, and reporting Timeline posts as spam).

Finally, these privacy activities are submitted to a Mixture Factor Analysis procedure, which classifies users into distinct classes based on their activities (see [28,63] for technical details). The activity pattern in each class describes a privacy profile. In most cases, one can assign a meaningful

label to these profiles, too (e.g. a profiles in which "limiting access" and "withholding info are the most prominent activities can be labeled as the "self-censor" profile).

Tailoring Privacy by Design to User Profiles

The second step is to develop privacy by design solutions that fit each of the identified profiles (Figure 1, right). There are two ways to apply user-tailored privacy by design: direct application and extrapolated application.

In a *direct application*, the privacy controls of the system for which the profiles were developed are tailored in a way that changes their salience or default setting depending on the profile of the current user. This is the most straightforward application of user-tailored privacy by design, because there is a direct mapping between profiles and features: profiles are defined by increased (or sometimes decreased) privacy activities (e.g. "self-censors" on Facebook may be more likely to withhold information and limit access to their content), activities consist of underlying behaviors (e.g. limiting access may consist of reducing the visibility of Timeline posts, or posts on others' Timelines in which the user is tagged), and each behavior is implemented by an existing feature (e.g. Facebook has a specific control for reducing the default visibility of posts on the user's Timeline). Research on information disclosure shows that the salience [14,21,22,29,67] and default setting [26,37] of privacy controls significantly influences users' engagement with such controls. User-tailored privacy by design can thus be implemented for each profile by emphasizing features that are more likely to be used by users with that profile, which will make this behavior easier to engage in.

When user-tailored privacy is implemented effectively, the system will tailor its interface to the privacy profile of the user. How does the system assign the correct privacy profile to the current user? There are multiple ways of doing this. The simplest method is to allow the user to simply select the profile themselves. This will work best if there is a limited number of profiles, each with a semantically descriptive label. Another method is to simply try certain profiles, and observe to which profile application the user reacts most favorably (this is akin to the idea of "website morphing" [19] or "bandit testing" [60]). A more sophisticated method assign users to profiles based on demographics (cf. [28], provided that users disclose these demographics, of course). Finally, the most sophisticated technique tracks users' privacy behaviors as they use the system, and then assigns a profile dynamically (cf. [25]). In the examples we present below, we stop short of implementing one of these tailoring procedures. Instead, we focus on creating designs that make it easier to engage in the behaviors related to each profile.

This brings up the question: Why not circumvent the adaptiveness altogether, and make *all* behaviors easier to engage in for *every* user? First of all, emphasizing all possible privacy management features of a system would significantly clutter its interface and spoil its design aesthetic—this is something that we expect only the users with the most privacy-sensitive profile are willing to put up with. Secondly, due to the persuasive nature of salience and default settings, an undue emphasis on all available privacy management features would inadvertently "nudge" users to engage in more (and/or different) privacy-related behaviors than they normally would, thereby tipping the privacy-functionality balance unduly towards privacy. Indeed, Sunstein and Thaler (authors of the seminal work on nudging [54]) argue that developers/designers have a moral obligation to implement nudges in a balanced manner [52].

Aside from direct application of user-tailored privacy, our methodology also allows for *extrapolated application*. In extrapolated application, the user profiles identified in one system are used as "personas" to the develop privacy design guidelines for a different system (often a system that is new, or not yet implemented) that has (or is envisioned to have) similar privacy features. Personas are a design tool first introduced by Cooper as a means to focus design practice on key segments of the audience of a system. Like profiles, personas are an increasingly popular tool in the field of usable privacy [3,4]. Personas serve a more conceptual purpose compared to profiles—this is necessary because there may not be a direct mapping between the privacy functionality of the system on which the profiles are based, and the system to which these profiles are subsequently applied. As such, the extrapolated application of privacy profiles in new or not yet implemented systems often takes the form of design guidelines for privacy by design features that may or may not be tailored to the user.

APPLYING THE UTPBD FRAMEWORK

We instantiate our proposed UTPbD framework within two different contexts in order to tangibly illustrate its application. First, we extend Wisniewski et al.'s [63] work, which completed the first stage of our UTPbD framework by creating six privacy management strategy profiles of Facebook users. We build upon this work by taking directly applying PbD principles to the privacy controls within Facebook as they map to the user profiles.

Second, we generalize our UTPbD framework beyond Facebook and SNSs by extrapolating Wisniewski et al.'s privacy profiles [63] to a new type of online platform: the "Total Learning Architecture" or TLA that is currently being developed by the Advanced Distributed Learning Initiative [44]. TLA is an architecture for "next-generation" learning systems. It comprises an open source set of specifications that describe how development patterns, interfaces (APIs), and data models can be implemented to facilitate sharing analytics about learners and their learning process across different platforms, systems and technologies [44]. Because TLA specifications are still in the process of being formalized, and thus, no actual systems currently, PbD is a well-suited methodology to be applied at this early juncture of design. Our goal is to show the value of applying our user-tailored approach to PbD for TLAs.

Privacy Behaviors on Facebook

Wisniewski et al. [63] identified a total of 32 privacy behaviors that Facebook users can perform. They then followed the Knijnenburg et al. [28,30] approach to uncover eleven privacy activities on Facebook (see Table 1 in [63]):

1. *Altering News Feed* includes hiding a story, changing a subscription, and unsubscribing from status updates.
2. *Moderating Timeline* consists of deleting or hiding content, and reporting content as spam.
3. *Reputation management* happens through untagging, or asking a friend to take down an unwanted photo or post.
4. *Limiting access* is effected by reducing the default visibility of information shared through one's Timeline, and of posts on others' Timelines in which one is tagged.
5. *Blocking people* consists of blocking a user, or adding them to the "restricted" list.
6. *Blocking apps/events* happens by blocking invitations to install an app or join an event.
7. *Restricting chat availability* is effected by going "offline" on Facebook Chat, or by changing the default visibility in Facebook Chat to invisible.
8. *Selective sharing* happens when the user posts a photo or status message to a custom friend list.
9. *Custom friend list creation* consists of categorizing a new or existing friend into a custom friend list.
10. *Withholding contact information* includes withholding one's cell phone number, other phone number, IM screen name, and street address.
11. *Withholding basic information* consists of withholding one's interests, religion and political views.

Facebook Privacy Profiles

Next, Wisniewski et al. identified six privacy management profiles that summarize the distinctly different ways in which users manage their privacy (see Figures 3-5 in [63]):

1. *Selective Sharers* limit the audience with whom their share information. They limit the default visibility of posts on their Timeline and posts in which they are tagged. Additionally, they create custom friend lists, and use these to share content selectively.
2. *Self-Censors* use few of the privacy features that allow for selective sharing, but instead protect their privacy by withholding information from anyone.
3. *Time Savers/Consumers* use Facebook in a way that allows them to consume relevant information without being bothered by unwanted conversations or status updates. For them, privacy is less about the right to withhold information, and more about the "right to be left alone". Consequently, they use privacy strategies to selectively read posts without being bothered by others.
4. *Privacy Maximizers* employ every privacy activity available to them, except for limiting access to posts others' make on their Timeline or tag them in.
5. *Privacy Balancers* exhibit moderate levels of privacy management behaviors across the entire spectrum of privacy features. They do not engage in privacy management to the same extent as Maximizers, but are considerably more active than Minimalists.
6. *Privacy Minimalists* use only a few common methods to protect their privacy.

FEATURE-LEVEL UTPBD FOR FACEBOOK

In this section, we propose design solutions based on the existing Facebook privacy management functionality that are tailored to the six profiles uncovered by Wisniewski et al. [63]. These UTPbD solutions (for an overview, see Table 1) make it easier for users with a certain privacy management strategies to engage in the privacy management behaviors that are associated with a particular user profile.

PbD for Selective Sharers

Selective Sharers want to limit the audience with whom they share information. In our design for these users, we propose setting the default access for posts to their Timeline and posts in which they are tagged to "friends only," making it easier to assign friends to custom friend lists (Figure 2), and make the selective sharing options that Facebook provides when submitting a new post more prominent (Figure 3).

Figure 2: A more prominent design for friend list management. Users can directly classify friends into a list.

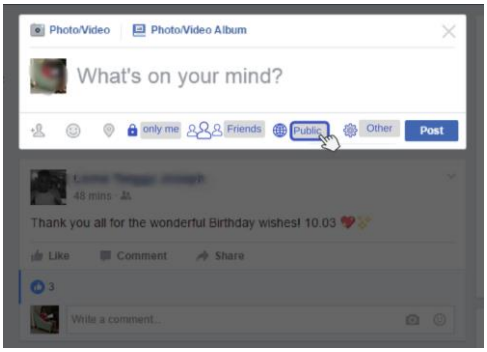

Figure 3: A more prominent design for selective sharing. Users can directly change the audience of a post with toggle buttons, without having to use the standard drop-down list.

Selective Sharers are also relatively more likely to block apps, events and people, so we suggest to put the blocking functionality directly in the notifications list (Figure 4).

Figure 4: A more prominent design for blocking apps, events, and people, that is displayed directly in the notifications.

Additionally, *Selective Sharers* are relatively more likely to moderate posts on their Timeline, alter their News Feed, and manage their reputation by untagging. The default Facebook interface hides the features related to these activities under a dropdown at the top-right side of a post; we make them more directly accessible to *Selective Sharers* by placing additional buttons next to this dropdown (Figures 5 and 6)

Figure 5: A more prominent design for News Feed and reputation management. Users can easily unfollow a user or a page, hide a post, and untag themselves.

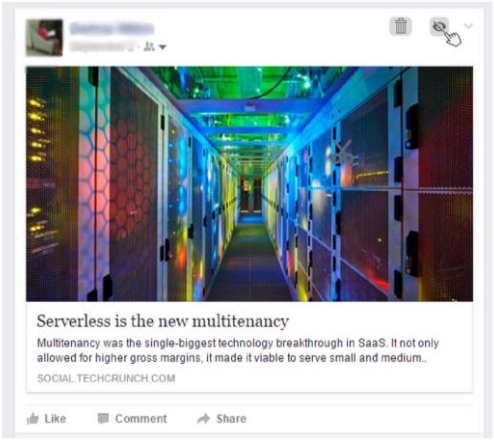

Figure 6: A more prominent design for Timeline moderation. Users can easily delete or hide posts on their Timeline.

Finally, *Selective Sharers* prefer to limit their availability in Facebook's chat. We thus automatically set these users to "offline" on chat when they log in, but make it easy to change their availability by using a toggle button (Figure 7).

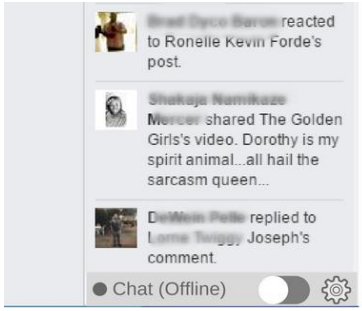

Figure 7: A more prominent design for restricting chat. Users can use the toggle to go online or offline in Facebook chat without having to use the standard options pop-up.

PbD for Self-Censors

Unlike *Selective Sharers*, who share abundantly but selectively, *Self-Censors* do not make distinctions between friends but instead prefer not to share their basic and contact information with anyone. For these users we set the default visibility of personal information (e.g. phone number, address, interests, religious and political views) to "only me" (Figure 8).

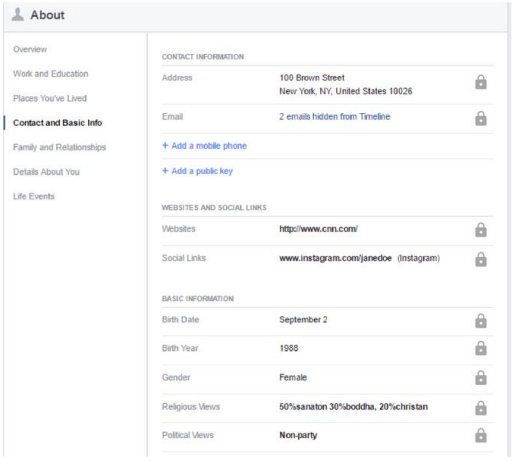

Figure 8: The default visibility for Contact and Basic Info is set to "only me".

At the same time, we reduce interface clutter for these users by removing the friend list management functionality from the dialog that pops up when the user hovers over a friend's name (Figure 9). Should the user want to categorize this friend after all, then they can still find this functionality by going to the friend's page.

Figure 9: A *less* prominent design for friend list management (left). The features that enable users to categorize friends into lists are removed from the dialog that pops up when the user hovers over a friend's name in the original interface (right).

PbD for Time Savers

Time Savers use privacy strategies that enable them to selectively read posts without being bothered by unwanted chat messages or status updates. To facilitate this behavioral pattern, set their chat availability to "offline" by default, and make it easier to go offline on chat at any time by means of a toggle button (Figure 7).

Time Savers also tend to alter their News Feed by deleting stories and/or hiding posts (Figure 5), but do not engage much in reputation management, so that feature may remain in the dropdown list rather than being displayed as a button.

Time Savers also occasionally moderate their own Timeline, so we also highlight the dropdown menu that allows them to remove or hide posts in their Timeline (Figure 10), but unlike for *Selective Sharers* we do not pull these features out of the dropdown menu (as in Figure 6), because these behaviors are not as strong for *Time Savers* as they are for *Selective Sharers*.

Finally, since time savers rarely create or use custom friend lists, we deemphasize this feature by removing the functionality from the friend popup dialog (Figure 9).

Figure 10: A somewhat more prominent design for Timeline moderation. More emphasis is put on the dropdown menu where users can unfollow a user or a page, and untag themselves in a post.

PbD for Privacy Maximizers

Privacy Maximizers display the widest variety of privacy behaviors, i.e. they utilize almost all of the available privacy features. Therefore, increasing the accessibility of all of the eleven aforementioned privacy activities (except for the features that allow them to limit access to posts others' make on their Timeline or tag them in) will help *Privacy Maximizers* to maintain their preferred settings. Specifically, we make it easier to manage custom friend lists (Figure 2), share posts selectively (Figure 3), alter their News Feed, manage their reputation through untagging (Figure 5), moderate posts on their Timeline (Figure 6), block apps, events and people (Figure 4), withhold personal information (Figure 8), and restrict chat accessibility (Figure 7).

Note that the combination of these features is likely going to result in a considerably more cluttered interface. However, we argue that *Privacy Maximizers* are likely to have such strong privacy concerns that they prefer this cluttered interface over the standard Facebook interface (cf. [55,56]).

PbD for Privacy Balancers

Privacy Balancers display moderate levels of privacy management. Designing for these users is particularly hard; we cannot simply highlight all features like we do for *Maximizers*, because they do not seem to have similarly strong privacy concerns. Our solution is to make certain key privacy features more prominent. Specifically, we increase the accessibility of the settings for restricting their chat availability (Figure 7), Timeline moderation (the "light" version, Figure 10), altering post on their News Feed and managing their reputation (Figure 5), and blocking apps, events, and people (Figure 4).

PbD for Privacy Minimalists

Privacy Minimalists report the lowest levels of privacy management behavior among all user classes. For these users we keep the Facebook interface "as is", except that we remove the friend list assignment functionality from the friend popup dialog (Figure 2).

PERSONA-LEVEL UTPBD FOR TLA

Privacy behaviors within TLA-based systems

As an architecture that enables pervasive user monitoring, integration of various learning applications, and data sharing among different users, the TLA provides an excellent use case for the development of PbD solutions [43,44]. Indeed, the developers of TLA argue that "security and privacy considerations should be interwoven into the software development process from the very beginning. Engaging design assurance experts focused on securing and integrating subsystems into the final system will help address privacy and security concerns. (p. 71)" [43]. The TLA is envisioned to have a wide variety of users that span a broad spectrum of privacy attitudes. We therefore consider the implementation of privacy by design in TLA as a use case for extending our UTPbD methodology from an existing application (i.e. Facebook) to an application that is currently under development.

Selective Sharers	require a more restrictive default sharing setting, more prominent capabilities for friend list management and selective sharing, and a button to block apps, events and people in their notification window.
Self-Censors	do not use selective sharing capabilities (hence some friend list management features could be hidden), but benefit from their basic and contact info to be shared with "only me" by default.
Time Savers	require more prominent News Feed moderation features, and their chat availability should be set to offline by default.
Privacy Maximizers	require all of the functionality described above.
Privacy Balancers	require more prominent controls to alter their News Feed and timeline, a toggle to easily go offline in chat, and a button to block apps, events and people in their notification window.
Privacy Minimalists	use very few privacy features, so the current Facebook interface would be sufficient. Since they do not use selective sharing capabilities, some of the friend list management features could be hidden.

Table 1: Overview of the UTPbD solution for Facebook.

A system implementation leveraging the TLA specifications is envisioned to be an integrated, interoperable network of existing learning technologies that use pervasive user monitoring to provide meta-adaptive learning recommendations to a wide array of end-users [43]. Such a system would notify users regarding learning opportunities, track their progress, and provide users with personalized recommendations based on their personal goals and organizational needs. Social aspects, such as sharing recommendations, learning progress and achievements with peers and superiors, are also envisioned to play an important role in this process. Sharing this information is envisioned to expand exposure to learning content, encourage users to work harder, and increase organizational awareness of workers' skills and capabilities. Meta-adaptation—recommendations that cross technical boundaries and are able to identify differences in how learning systems address learner needs—necessitates the sharing of data between and among individual systems [15].

Yet, the TLA specifications have yet to be implemented in an actual learning ecosystem. Therefore, the approach taken by Wisniewski et. al. [63] for creating user profiles based on Facebook users' past privacy behavior within the system is not feasible for this stage of the TLA, and limits our ability to apply UTPbD. However, we argue that Wisniewski et al.'s profiles are sufficiently generic in nature that they may be used to develop UTPbD solutions for applications in domains other than SNS. To explore this presumption, we use the same six profiles to develop user-tailored design guidelines for the TLA. We anticipate that users of TLA systems will also likely be users of social media, such as Facebook, and therefore, can identify with the profiles of *Selective Sharers*, *Self-Censors*, *Time Savers*, *Maximizers*, *Balancers*, and *Minimalists*. As such, we suggest several ways in TLA-based systems can address the concerns of such different types of users.

PbD for Selective Sharers

While *Selective Sharers* are selective in deciding with whom their personal information should be shared, they seem to have less concern about disclosing their information to system itself. In the context of a TLA system, they are expected to provide their skills, interests, training schedule and achievements with the system, which allows them to benefit from the TLA's advanced personalization facilities. Note though, that in a network of training applications, they may be selective regarding the applications that they are willing to use (cf. blocking apps/events).

Selective Sharers would be more restrictive regarding the social aspects of TLA. Specifically, they would be likely to carefully manage who is within their network (cf. friend list management, blocking people), what training outcomes are posted publicly (cf. Timeline moderation), and who gets to see the information that is collected or generated by the training systems (cf. selective sharing) or that other people share about them (cf. reputation management). This means that for *Selective Sharers* the TLA systems should reduce the default visibility of personal information (cf. limiting access), and have privacy features that allow them to hide certain information from the public, and share it only with select groups of contacts within their network, such as their direct coworkers. Since the TLA is likely to impact promotion decisions, *Selective Sharers* may want to keep training results and job aptitude scores strictly confidential, and share them with their supervisors only.

Finally, *Selective Sharers* tend to consume selectively as well; they would be relatively more likely to limit communication while using the system (cf. restricting chat, altering News Feed). These activities are also prominent for *Time Savers*, so we discuss them in more detail below.

PbD for Self-Censors

Self-Censors tend to manage privacy by withholding information, and so TLA systems should allow these users limit the extent to which the system collects and tracks information about their skills, interests, training schedule and achievements are shared with the system (cf. withholding basic info, Timeline moderation). This lack of data collection is expected to hinder the personalization aspect of TLA, which is geared towards giving recommendations based on the user's activities and interests. As such, it is important to make these users aware of the fact that their learning recommendations will consequently not be tailored to their personal situation and preferences. If *Self-Censors* indeed

opt out of getting personalized learning recommendations, then the TLA system should still be able to provide relevant non-personalized recommendations, as well as a sufficiently powerful browsing functionality. This functionality could for instance give users a default set of the most popular learning tools for a specified task. This would be useful feature for *Self-Censors*, since it can expose them to relevant items without the need for extensive tracking.

Systems in the TLA architecture also track learning progress and outcomes as a means to give users credit for their learning activity, which may eventually influence job-related decisions. If *Self-Censors* refuse to share their learning outcomes with the system, this could prevent them from getting credit altogether, and impede their career goals. It would thus be best if such learning outcomes were still tracked, but shared only with direct supervisors, and only at a granular level (e.g. no detailed learning activity report, but only an overall assessment of learning performance at a level of detail that is sufficient for making promotion decisions).

For the social aspects of TLA, *Self-Censors* should be allowed to prevent their information from being shared with their network. Unlike *Selective Sharers*, *Self-Censors* would likely not customize sharing with specific groups, but instead allow hide it from the entire network. Note that *Self-Censors* also tend to hide their contact information; this indicates that they prefer to protect their "real world" privacy as well. Real world social functionality, such as suggestions for group training, should thus also be avoided.

Privacy by Design for Time Savers

Time Savers willingly provide their personal information to the system, but unlike *Selective Sharers* they tend to share it rather indiscriminately. *Time Savers'* main privacy management strategy is to minimize the amount of communication they have while using the system, both when it comes to direct communication (cf. restricting chat) and indirect communication (cf. altering News Feed).

A relatively direct design implication that can be derived from this privacy management strategy is that *Time Savers* should have the ability to opt out of social connectivity features such as chat or status updates if the TLA implementation has such functionality.

A more indirect implication could be that the system should allow *Timer Savers* to consume relevant recommendations without being bothered by too much interaction. This may require features like allowing them to curate their list of suggested recommendations (cf. altering News Feed) and allowing them to switch off push notifications and emails sent out by the system.

PbD for Privacy Maximizers

Privacy Maximizers employ almost all of the combined privacy management activities of *Selective Sharers*, *Self-Censors*, and *Time Savers*. This means that all of the functionality described above should be available for Privacy Maximizers, which results in a system with features for reducing the collection and sharing of information, increasing the opportunity for curation, and allowing users to opt out of active notifications and social features.

PbD for Privacy Balancers

As mentioned earlier, *Privacy Balancers* are difficult to design for: while they do not portray particularly high levels of privacy concern, they do employ a variety of privacy functionality, but only to a limited extent. Their most prominent privacy activities involve curation (cf. altering News Feed, Timeline moderation, reputation management), blocking (cf. blocking apps, events, and people), and avoiding direct interaction.

We therefore suggest that *Privacy Balancers* to get the same functionality as *Time Savers* (i.e. allowing users to opt-out of active notifications and social features), plus some functionality to block specific learning applications and people, and to moderate some of the content of the system. Completely withholding of personal information is not necessary for *Privacy Balancers*, nor do they require any mechanism to carefully specify selective sharing of information with specific groups of people.

PbD for Privacy Minimalists

While *Privacy Minimalists* constitute the least privacy-sensitive privacy profile, it is important to contemplate design solutions for them as well. Particularly, for *Privacy Minimalists* the system needs to be designed in a way that unfettered personalization can take place. Tailoring to *Privacy Minimalists* means removing all possible barriers to data sharing, communication, and recommendation.

Selective Sharers	require sophisticated functionality to curate and selectively share their personal information and training outcomes with specific applications and groups of people.
Self-Censors	require mechanisms for curation, non-personalized mechanisms for the selection of learning material, and highly restricted forms of sharing learning outcomes.
Time Savers	should be able to opt out of active notifications and social features.
Privacy Maximizers	require all of the functionality described above.
Privacy Balancers	require mechanisms for curation, blocking, and avoiding direct interaction.
Privacy Minimalists	require systems that allow them to maximally benefit from their adaptive and social functionalities.

Table 2: Overview of the UTPbD solution for TLA.

DISCUSSION, LIMITATIONS, AND FUTURE WORK

Our work addresses the non-trivial question: "What adaptations can be made to support the user's privacy management strategy in a user-tailored way?" We leveraged an existing classification of Facebook users' privacy management behaviors, and tailored existing Facebook privacy features to the six profiles of this classification.

Facebook's original privacy functionality largely operates in the background (i.e. features are accessible through menus and dropdowns), so in our designs we bring features to the foreground based on the needs of each profile. We argued that this increases both their salience and their accessibility, thereby providing an "adaptive nudge".

We furthermore applied our user-tailored privacy by design approach to TLA, a next-generation learning architecture that is currently still in the conceptual stage. We successfully transferred the Facebook privacy management profiles to TLA by abstracting them to the level of personas. These personas allowed us to argue about the privacy features that TLA-based applications need to implement.

Our work is not without limitations. First of all, while our tailored PbD solutions are based on outcomes of extensive user research, it would be useful to bring our work "full circle", and test the suggested designs with the relevant groups of users. An experiment could be conducted to see how effective the profile-based PbD solutions are compared to the traditional Facebook interface (and potentially a "maxed out" interface with all PbD solutions enabled for every user) in terms of perceived privacy control, privacy threat, and overall system satisfaction. This study could also study the best method for classifying users into profiles: users could pick the profile by themselves, or we could implement some adaptive procedure for detecting the correct profile.

Our work also makes the normative assumption that UTPbD systems should tailor the privacy functionality to the user's current privacy practices. While this avoids nudging users into using features they do not want to use [52], one could question whether e.g. *Privacy Minimalists* fall into that profile because of a conscious decision, or because they are simply not aware of the privacy features that are available in the system [63]. In the latter case, one could argue for a version of UTPbD that highlights and makes accessible those features that the user is currently *not* using, in an effort make them more aware of these features, and more conscious of what they can do to maintain their privacy. Would this "antithetical user tailoring" method increase privacy awareness, or would the highlighted features simply be ignored, while reducing users' overall satisfaction? Future work can conduct a controlled experiment testing these two approaches against each other in order to find out.

A final limitation is that we make a theoretical jump by using the Facebook privacy profiles as personas for TLA. This limitation is hard to overcome, because TLA is still in a conceptual state. As TLA gets implemented in real learning applications, we can do a study to observe its users' privacy behaviors, and develop profiles based on this data. The similarity between the current Facebook profiles and the profiles we will detect in TLA will give us a good indication of the effectiveness of applying UTPbD at the persona-level in new networked applications.

CONCLUSION

In this paper we have presented a framework of *user-tailored privacy by design* as a means to provide privacy management support for a diverse set of users with a wide variety of privacy management strategies. Our work covers the implementation of privacy by design for each of the six privacy management profiles developed by Wisniewski et al. [63], and furthermore extends this approach to privacy recommendations for the Total Learning Architecture (TLA) that is being developed by the Advanced Distributed Learning (ADL) initiative.

Beyond existing privacy by design solutions, our work acknowledges the inherent heterogeneity of users' privacy preferences and management strategies, and attempts to give users the privacy they want by tailoring the design solutions to these preferences and strategies. Moreover, beyond existing user-tailored privacy solutions that focus on user-tailored sharing settings, our work attempts to tailor to the user's management of relational, territorial, network, and interactional boundaries. As such, our work provides a more comprehensive adaptation strategy based on user-tailored privacy by design solutions.

ACKNOWLEDGMENTS

This work was supported by the U.S. Advanced Distributed Learning (ADL) Initiative (Contract W911QY-16-C-0105). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ADL Initiative or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes. *Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

REFERENCES

1. Alessandro Acquisti and Ralph Gross. 2006. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In *Privacy Enhancing Technologies* (Lecture Notes in Computer Science), 36–58. https://doi.org/10.1007/11957454_3
2. Alessandro Acquisti, Leslie K John, and George Loewenstein. 2012. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research* 49, 2: 160–174. https://doi.org/10.1509/jmr.09.0215
3. Alessandro Acquisti, Bart P. Knijnenburg, Norman Sadeh, and Allison Woodruff. 2015. 2nd Annual Privacy Personas and Segmentation (PPS) Workshop. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security*, xviii–xix.
4. Alessandro Acquisti, Anthony Morton, Norman Sadeh, and Allison Woodruff. 2014. Workshop on Privacy Personas and Segmentation (PPS): Call for Papers. Retrieved April 12, 2014 from

https://cups.cs.cmu.edu/soups/2014/workshops/privacy.html

5. Manar Alohaly and Hassan Takabi. 2016. Better Privacy Indicators: A New Approach to Quantification of Privacy Policies. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*.
6. Laura Becker and Key Pousttchi. 2012. Social Networks: The Role of Users' Privacy Concerns. In *Proceedings of the 14th International Conference on Information Integration and Web-based Applications & Services* (IIWAS '12), 187–195. https://doi.org/10.1145/2428736.2428767
7. Andrew Besmer, Jason Watson, and Heather Richter Lipford. 2010. The impact of social navigation on privacy policy configuration. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, 7:1-7:10. https://doi.org/10.1145/1837110.1837120
8. Danah M. Boyd and Nicole B. Ellison. 2007. Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication* 13, 1: 210–230. https://doi.org/10.1111/j.1083-6101.2007.00393.x
9. Ann Cavoukian. 2010. *Privacy by Design*. Information and Privacy Commissioner of Ontario, Canada. Retrieved from http://www.privacybydesign.ca/content/uploads/2010/03/PrivacybyDesignBook.pdf
10. Huseyin Cavusoglu, Tuan Phan, and Hasan Cavusoglu. 2013. Privacy Controls and Content Sharing Patterns of Online Social Network Users: A Natural Experiment. In *ICIS 2013 Proceedings*. Retrieved from http://aisel.aisnet.org/icis2013/proceedings/ResearchInProgress/54
11. Ramón Compañó and Wainer Lusoli. 2010. The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas. In *Economics of Information Security and Privacy*, 169–185. https://doi.org/10.1007/978-1-4419-6967-5_9
12. Consumer Reports. 2012. Facebook & your privacy: Who sees the data you share on the biggest social network? *Consumer Reports*. Retrieved May 13, 2012 from http://www.consumerreports.org/cro/magazine/2012/06/facebook-your-privacy
13. Lujun Fang and Kristen LeFevre. 2010. Privacy Wizards for Social Networking Sites. In *Proceedings of the 19th International Conference on World Wide Web* (WWW '10), 351–360. https://doi.org/10.1145/1772690.1772727
14. B Fogg. 2003. *Persuasive technology : using computers to change what we think and do*. Morgan Kaufmann Publishers, Amsterdam.
15. Jeremiah T. Folsom-Kovarik and Elaine M. Raybourn. 2016. Total Learning Architecture (TLA) Enables Next-generation Learning via Meta-adaptation. In *Interservice/Industry Training, Simulation, and Education Conference Proceedings*.
16. 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. https://doi.org/10.1145/1102199.1102214
17. Harris Interactive inc. 2000. *A Survey of Consumer Privacy Attitudes and Behaviors*. Harris Interactive, Inc., New York, NY. Retrieved from http://www.bbbonline.org/UnderstandingPrivacy/library/harrissummary.pdf
18. Louis Harris, Alan F. Westin, and associates. 2003. *Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits*. Equifax Inc.
19. John R. Hauser, Glen L. Urban, Guilherme Liberali, and Michael Braun. 2009. Website Morphing. *Marketing Science* 28, 2: 202–223. https://doi.org/10.1287/mksc.1080.0459
20. Jaap-Henk Hoepman. 2014. Privacy Design Strategies. In *ICT Systems Security and Privacy Protection*, Nora Cuppens-Boulahia, Frédéric Cuppens, Sushil Jajodia, Anas Abou El Kalam and Thierry Sans (eds.). Springer Berlin Heidelberg, 446–459. https://doi.org/10.1007/978-3-642-55415-5_38
21. David J. Houghton and Adam N. Joinson. 2010. Privacy, Social Network Sites, and Social Relations. *Journal of Technology in Human Services* 28, 1–2: 74–94. https://doi.org/10.1080/15228831003770775
22. Thomas Hughes-Roberts. 2015. Privacy as a secondary goal problem: an experiment examining control. *Information and Computer Security* 23, 4: 382–393. https://doi.org/10.1108/ICS-10-2014-0068
23. Eric J. Johnson, Steven Bellman, and Gerald L. Lohse. 2002. Defaults, Framing and Privacy: Why Opting In ≠ Opting Out. *Marketing Letters* 13, 1: 5–15. https://doi.org/10.1023/A:1015044207315
24. Pamela Karr-Wisniewski, Heather Lipford, and David Wilson. 2011. A New Social Order: Mechanisms for Social Network Site Boundary Regulation.
25. B. P. Knijnenburg. 2015. A user-tailored approach to privacy decision support. Retrieved from http://search.proquest.com/docview/1725139739/abstract
26. B. P. Knijnenburg and A. Kobsa. 2014. Increasing Sharing Tendency Without Reducing Satisfaction: Finding the Best Privacy-Settings User Interface for Social Networks. In *ICIS 2014 Proceedings*.
27. B. P. Knijnenburg and Alfred Kobsa. 2013. Making Decisions about Privacy: Information Disclosure in

Context-Aware Recommender Systems. *ACM Transactions on Interactive Intelligent Systems* 3, 3: 20:1-20:23. https://doi.org/10.1145/2499670

28. B. P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71, 12: 1144–1162. https://doi.org/10.1016/j.ijhcs.2013.06.003
29. B. P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Counteracting the Negative Effect of Form Auto-completion on the Privacy Calculus. In *ICIS 2013 Proceedings*.
30. Bart P. Knijnenburg. 2014. Information Disclosure Profiles for Segmentation and Recommendation. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*.
31. Yee-Lin Lai and Kai-Lung Hui. 2006. Internet Opt-In and Opt-Out: Investigating the Roles of Frames, Defaults and Privacy Concerns. In *Proceedings of the 2006 ACM SIGMIS Conference on Computer Personnel Research*, 253–263. https://doi.org/10.1145/1125170.1125230
32. Marc Langheinrich. 2001. Privacy by Design: Principles of Privacy-Aware Ubiquitous Systems. In *Ubicomp 2001*, 273–291. https://doi.org/10.1007/3-540-45427-6_23
33. Scott Lederer, Jason I. Hong, Anind K. Dey, and James A. Landay. 2004. Personal privacy through understanding and action: five pitfalls for designers. *Personal and Ubiquitous Computing* 8, 6: 440–454. https://doi.org/10.1007/s00779-004-0304-9
34. Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 27–41. Retrieved from https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu
35. Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help? In *Proceedings of the 23rd International Conference on World Wide Web* (WWW '14), 201–212. https://doi.org/10.1145/2566486.2568035
36. Alessandra Mazzia, Kristen LeFevre, and Eytan Adar. 2012. The PViz Comprehension Tool for Social Network Privacy Settings. In *Proceedings of the Eighth Symposium on Usable Privacy and Security* (SOUPS '12), 13:1–13:12. https://doi.org/10.1145/2335356.2335374
37. Craig R. M. McKenzie, Michael J. Liersch, and Stacey R. Finkelstein. 2006. Recommendations Implicit in Policy Defaults. *Psychological Science* 17, 5: 414–420. https://doi.org/10.1111/j.1467-9280.2006.01721.x
38. Deirdre Mulligan and Jennifer King. 2012. Bridging the Gap Between Privacy and Design. *University of Pennsylvania Journal of Constitutional Law* 14, 4: 989.
39. 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: 232–243. https://doi.org/10.1016/j.pmcj.2013.12.004
40. Sameer Patil, Xinru Page, and Alfred Kobsa. 2011. With a little help from my friends: can social navigation inform interpersonal privacy preferences? In *Proceedings of the ACM 2011 conference on Computer supported cooperative work* (CSCW '11), 391–394. https://doi.org/10.1145/1958824.1958885
41. Frederic Raber, Alexander De Luca, and Moritz Graus. 2016. Privacy Wedges: Area-Based Audience Selection for Social Network Posts. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*.
42. Ramprasad Ravichandran, Michael Benisch, Patrick Kelley, and Norman Sadeh. 2009. Capturing Social Networking Privacy Preferences. In *Privacy Enhancing Technologies* (Lecture Notes in Computer Science), 1–18. https://doi.org/10.1007/978-3-642-03168-7_1
43. Elaine M. Raybourn, Nathan Fabian, Warren Davis, Raymond C. Parks, Jonathan McClain, Derek Trumbo, Damon Regan, and Paula Durlach. 2015. Data Privacy and Security Considerations for Personal Assistants for Learning (PAL). In *Proceedings of the 20th International Conference on Intelligent User Interfaces Companion*, 69–72. https://doi.org/10.1145/2732158.2732195
44. Damon Regan, Elaine M Raybourn, and Paula J Durlach. 2013. Personalized Assistant for Learning (PAL). In *Design Recommendations for Intelligent Tutoring Systems: Volume 1-Learner Modeling*, Robert A. Sottilare, Arthur Graesser, Xiangen Hu and Heather Holden (eds.). U.S. Army Research Laboratory, Orlando, FL, 217.
45. 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: 401–412. https://doi.org/10.1007/s00779-008-0214-3
46. Peter Schaar. 2010. Privacy by Design. *Identity in the Information Society* 3, 2: 267–274. https://doi.org/10.1007/s12394-010-0055-x
47. 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: 159–172. https://doi.org/10.1509/jppm.10.114

48. Sarah Spiekermann. 2012. The Challenges of Privacy by Design. *Commun. ACM* 55, 7: 38–40. https://doi.org/10.1145/2209249.2209263
49. Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. 2001. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, 38–47.
50. Luke Stark, Jen King, Xinru Page, Airi Lampinen, Jessica Vitak, Pamela Wisniewski, Tara Whalen, and Nathaniel Good. 2016. Bridging the Gap Between Privacy by Design and Privacy in Practice. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '16), 3415–3422. https://doi.org/10.1145/2851581.2856503
51. Katherine Strater and Heather Richter Lipford. 2008. 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*, 111–119.
52. Cass R. Sunstein and Richard H. Thaler. 2003. Libertarian Paternalism Is Not an Oxymoron. *The University of Chicago Law Review* 70, 4: 1159–1202. https://doi.org/10.2307/1600573
53. Karen Tang, Jason Hong, and Dan Siewiorek. 2012. The implications of offering more disclosure choices for social location sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 391–394. https://doi.org/10.1145/2207676.2207730
54. Richard H Thaler and Cass Sunstein. 2008. *Nudge : improving decisions about health, wealth, and happiness*. Yale University Press, New Haven, NJ & London, U.K.
55. Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A Field Trial of Privacy Nudges for Facebook. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, 2367–2376. https://doi.org/10.1145/2556288.2557413
56. Yang Wang, Pedro Giovanni Leon, Kevin Scott, Xiaoxuan Chen, Alessandro Acquisti, and Lorrie Faith Cranor. 2013. Privacy Nudges for Social Media: An Exploratory Facebook Study. In *Second International Workshop on Privacy and Security in Online Social Media*, 763–770. Retrieved April 26, 2013 from http://www2013.org/companion/p763.pdf
57. Jason Watson, Heather Richter Lipford, and Andrew Besmer. 2015. Mapping User Preference to Privacy Default Settings. *ACM Transactions on Computer-Human Interaction* 22, 6: 32:1–32:20. https://doi.org/10.1145/2811257
58. Alan F Westin, Louis Harris, and associates. 1981. *The Dimensions of privacy : a national opinion research survey of attitudes toward privacy*. Garland Publishing, New York.
59. Alan F. Westin and Danielle Maurici. 1998. *E-Commerce & Privacy: What the Net Users Want*. Privacy & American Business, and PricewaterhouseCoopers LLP.
60. John Myles White. 2012. *Bandit Algorithms for Website Optimization*. O'Reilly Media, Inc.
61. Pamela Wisniewski, A. K. M. Islam, Heather Richter Lipford, and David Wilson. 2016. Framing and Measuring Multi-dimensional Interpersonal Privacy Preferences of Social Networking Site Users. *Communications of the Association for Information Systems* 38, 1. Retrieved from http://aisel.aisnet.org/cais/vol38/iss1/10
62. Pamela Wisniewski, A.K.M. Najmul Islam, Bart P. Knijnenburg, and Sameer Patil. 2015. Give Social Network Users the Privacy They Want. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (CSCW '15), 1427–1441. https://doi.org/10.1145/2675133.2675256
63. Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2016. Making Privacy Personal: Profiling Social Network Users to Inform Privacy Education and Nudging. *International Journal of Human-Computer Studies*. https://doi.org/10.1016/j.ijhcs.2016.09.006
64. Pamela Wisniewski, Bart P. Knijnenburg, and H. Richter Lipford. 2014. Profiling Facebook Users' Privacy Behaviors. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*.
65. Pamela Wisniewski, Heather Lipford, and David Wilson. 2012. Fighting for my space: coping mechanisms for SNS boundary regulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12), 609–618. https://doi.org/10.1145/2207676.2207761
66. Pamela Wisniewski, Heng Xu, and Yunan Chen. 2014. Understanding User Adaptation Strategies for the Launching of Facebook Timeline. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 2421–2430. https://doi.org/10.1145/2556288.2557363
67. Heng Xu, Na Wang, and Jens Grossklags. 2012. Privacy-by-ReDesign: Alleviating Privacy Concerns for Third-Party Applications. In *ICIS 2012 Proceedings*.