

PS4TLA: Privacy Support for the Total Learning Architecture

Specification Document, Volume 2:

Modeling Factors

Author information

Project lead: Dr. Bart P. Knijnenburg, Clemson University

Project team: David Cherry, Yang He, Reza Ghaiumy Anaraky, Moses Namara, Dr. Erin Ash

Technical point of contact: Andy Johnson, Advanced Distributed Learning (ADL) Initiative

"Mention of any commercial product in this paper does not imply DoD endorsement or recommendation for or against the use of any such product. No infringement on the rights of the holders of the registered trademarks is intended."

**CLEARED
For Open Publication**

Dec 04, 2018

Department of Defense
OFFICE OF PREPUBLICATION AND SECURITY REVIEW

Recommended citation: Knijnenburg, B.P., He, Y., Cherry, D., Ghaiumy Anaraky, R., Namara, M., Ash, E., and Johnson, A. (2018) "Privacy Support for the Total Learning Architecture: Modeling Factors". PS4TLA Specification Document, v0.2, Clemson University, Clemson, SC. Available at: <http://www.usabart.nl/PS4TLA/spec2.pdf>.

Executive summary

The purpose of this document is to make recommendations for implementing User-Tailored Privacy (UTP) into Total Learning Architecture (TLA)-based systems and to inform ADL and other TLA performers about the Modeling Factors that need to be considered in the context of this implementation. The set of recommendations put forth in this document will allow ADL and other TLA performers to build a user-tailored privacy decision-support system that supports users in making better privacy decisions.

The first section **makes the case for UTP**. This section presents the argument that technical solutions are insufficient to ensure that users will feel comfortable and capable of disclosing the personal information needed by TLA to provide learning recommendations. UTP complements technical solutions by supporting users' privacy management practices. Moreover, this section notes that the concept of Privacy by Design (the topic of Specification Vol. 1) cannot resolve privacy problems where users have a plurality of conflicting preferences regarding their privacy and the benefits of disclosure. In these cases, a personalized approach is needed. This section also demonstrates that transparency and control put an undue burden on the user to manage their own privacy settings—a complicated task, especially in complex learning management systems, such as those that will likely be based on TLA. UTP can alleviate some of this burden by automating part of the privacy decision-making process. Finally, this section demonstrates that privacy nudging implies a paternalistic and one-size-fits-all approach that will not benefit all users equally. UTP avoids paternalism by providing personalized nudges.

The second section subsequently provides a **definition** of UTP and presents it as a framework that measures users' characteristics and behaviors; models their decisions using machine learning algorithms; and then adapts the privacy settings, justifications, user interface, and/or personalization to match the user's preferences and the context of the decision. This section also highlights how UTP can be implemented to support privacy-setting practices in TLA's learning applications and social learning practices.

The subsequent sections detail the three stages of UTP's "measure, model, adapt" framework. The section on **measuring privacy** covers the antecedents of TLA users' privacy decision-making practices: the data, the user themselves, the recipient, and other system specific and purpose specific factors. This section demonstrates that users' privacy behaviors are multi-dimensional. It notes that cultural, demographic, and personality differences create large variations in privacy practices among users, but that this variability can often be captured by a concise set of "privacy profiles". It subsequently notes that users' concerns and behaviors towards the recipients of their data is governed by trust, and that recipients can often be grouped into a number of groups or "circles". Finally, it surveys the impact of time, location, and other contextual factors on users' privacy decisions.

The section on **modeling privacy** explains how these measurements can serve as inputs for a privacy prediction algorithm that can model TLA users' privacy behaviors. This section surveys the relative usefulness of input data available in TLA and the use of various modeling algorithms.

It also notes how matching the users' current privacy practices may not always be the best modeling strategy; in certain cases, UTP can solidify users' privacy management practices by recommending privacy behaviors that are complementary with their current behaviors, while in other cases UTP can completely move beyond users' current behavioral/attitudinal patterns. It also covers a number of potential modeling problems, such as the choice of cost functions, trade-offs with other user goals, overfitting, and cold start problems. This section concludes with the recommendation of taking a layered and gracefully degrading approach to privacy modeling in TLA.

The section on **adapting privacy** covers ways UTP can personalize the privacy settings of a TLA-based application, the justification it gives for requesting certain information, its privacy-setting interface, and its learning recommendation practices. This section discusses how UTP can take a proactive automated approach or a more conservative suggestion-based approach to personalize privacy settings. It explains how UTP can tailor justifications to the user, adjust them to the decision context, and optimize their timing. It suggests that in more complicated privacy management situations UTP can adjust the user interface of a TLA-based application to emphasize or deemphasize certain privacy management functionalities. It also suggests that TLA Processors can make use of privacy-enabling recommendation algorithms to provide micro-, macro-, and meta-adaptations. Finally, this section discusses how proactive and conservative adaptation strategies need to be balanced in order to reduce users' burden but at the same time give them sufficient control and reduce undue persuasion.

The final section outlines the higher-level **goals** of UTP in a TLA environment. This section acknowledges that to support the user, UTP must reconcile their various, potentially conflicting goals. It also covers the notion of UTP support for teaching the user about privacy, e.g., using privacy tips and self-actualization-based suggestions. This section also suggests that UTP's support can extend beyond the current user, and help activity providers, researchers, and supervisors to improve personalization practices, educational research, and personnel-related decision-making, respectively. Finally, it discusses practical and ethical considerations of reconciling the potentially conflicting goals of these various stakeholders in the privacy decision-making process.

For the final document, we will seek consensus among TLA performers regarding the operational characteristics (see Specification Vol. 1) and the implementation of user-tailored privacy. This will allow us to make specific and concrete recommendations regarding privacy support for TLA.

Introduction

Purpose

The Total Learning Architecture (TLA) is a set of specifications to enable the development of next-generation learning systems. TLA crucially depends on the collection of an extensive amount of user data to provide social learning capabilities and personalized learning recommendations [101,287]. Users will only agree to such extensive data collection if they feel that their privacy is adequately protected. This means that the TLA will need a plethora of settings that allow users to customize their desired level of privacy. Unfortunately, though, there is ample evidence that users often have difficulties navigating privacy settings. User-Tailored Privacy is thus an approach to privacy that models users' privacy concerns and provides them with adaptive privacy decision support [166,167,175,180]. By providing user-tailored support, it reconciles the need for extensive customizability with users' lack of skills and motivation to manage their own privacy settings. The purpose of this document is to allow ADL and other TLA performers to build a user-tailored privacy decision-support system that supports users in making better privacy decisions.

Scope

This document describes the modeling factors that need to be considered when building UTP into learning systems in general—and specifically TLA—as well as their potential use to support users' privacy decisions.

This document is written to support both the current development of the TLA specifications, as well as current and future implementations of these specifications in real-life distributed learning systems. The described uses of UTP may therefore go beyond any currently envisioned specification and implementation of TLA.

Technical aspects of UTP are not discussed in this document. Most notably, this document does not concern specific algorithms for the implementation of UTP. These algorithms are no different from the types of algorithms used in generic recommender systems.

Where possible, the document contains concrete recommendations. Further recommendations will be added after an intensive discussion with ADL and other TLA performers during the development of the final version of this document.

Definitions and Abbreviations

The **Total Learning Architecture (TLA)** is a set of specifications to enable the creation of a next-generation Learning Management System (LMS). These specifications consist of a set of web service specifications and APIs for sharing learning-related user data in a consistent way, thereby

allowing the integration of learning applications (User Facing Apps created by Activity Providers) ranging from eBooks to Massively Open Online Courses (MOOCs) into comprehensive personalized e-learning solutions [287]. The TLA specifies ubiquitous data collection (e.g. by integrating a wide variety of learning applications, interfacing with social media activity, and tracking smartphone sensors) and user modeling (e.g. by collecting highly detailed learner runtime activity) to enable highly personalized and pervasive (On The Job, Just In Time) training recommendations, calculated by the TLA Providers [101]. Moreover, the TLA specifications calls for an Open Social Learner Model (OSLM) that allows learning materials, activities, and outcomes to be shared across learners (enabling peer interactions) and learning systems (allowing for an extensible learning environment) [336]. This document describes how certain characteristics of the TLA specification—and of distributed learning systems implementing these specifications—have an impact on users’ privacy concerns.

User Tailored Privacy (UTP) is an approach to privacy (often implemented as an adaptive system) that measures user privacy-related characteristics and behaviors, uses this as input to model their privacy preferences, and then adapts the system’s privacy settings, user interface, contextualized help, and recommendation practices to these preferences [166,167,175,180].

Overview

Privacy threats have been shown to be an important barrier to the adoption of personalized systems [2,23,56,104,189,266,291,315,328], and it is therefore of utmost importance that such threats are minimized in any TLA-based system. From a privacy perspective, the social capital-based advantages of freely sharing learner profiles are at odds with the fact that these learner profiles may be protected by laws like FERPA, since these profiles are also used for sensitive employment decisions regarding placement, selection and promotion. On top of this, the envisioned international deployment of TLA introduces prominent cultural variation in privacy concerns and social etiquette [51,67,78,83,203]. Because of this, users of TLA-based distributed learning systems must carefully navigate a multi-dimensional array of privacy concerns, carefully balancing the benefits and risks of disclosing or allowing access to their personal information. However, users of complex information systems have been consistently incapable of effectively managing their own privacy [8,142,144,171,199,217,222], leaving them vulnerable to perceived and real privacy threats.

UTP is an approach to privacy that provides adaptive decision support to help users manage their privacy. Broadly speaking, user-tailored privacy first predicts users’ privacy preferences and behaviors based on their known characteristics. It can then use these predictions in various ways. One way is to provide automatic default settings, request orders, or suggestions in line with users’ disclosure profiles. Another is to provide justification-style nudges in situations where they are needed, but to only show them to users who can be expected to react favorably to them, so that they will not cause privacy scares in the other users. Yet another solution is to adapt the privacy-setting user interfaces to make it easier to engage with the available privacy management tools. Finally, UTP can adapt the recommendation algorithm by selectively

restricting the types of personalization a system can engage in based on the collected user data, thereby preventing potential unwanted inferences to be made.

This document first makes a case for UTP by highlighting the shortcomings of other paradigms for privacy support (Section 1). It then provides a comprehensive definition of UTP and explains various ways in which it can be implemented in TLA (Section 2). Subsequent sections explain how UTP can measure user- and context-related factors that influence users' privacy concerns and behaviors (Sections 3), how to model these concerns and behaviors (Section 4), and how to subsequently adapt the system to this model (Section 5). The document concludes with a discussion of the various privacy-related goals that TLA can support (Section 6).

Where possible, concrete recommendations are made. Further recommendations will be added after an intensive discussion with ADL and other TLA performers.

1 Making a case for User-Tailored Privacy

The TLA specifications envision an interconnected set of learning activity providers, supported by a personalization architecture consisting of TLA Processors and a TLA Data Core. This architecture allows for meta-adaptations (cross-provider individualized recommendations to switch from one learning activity to another), macro-adaptations (recommendations to determine the next learning activity within a single learning activity provider), and micro-adaptations (personalized changes to learning content within a single learning activity). These adaptations are supported by the collection of vast amounts of data, and users may not always be comfortable with these data collection practices.

While it is common knowledge in the realm of personalization that more data leads to better personalization outcomes [311], not all data collection practices are equally useful for personalization purposes [166], not all personalization practices are equally valued by the user [112], and users' preferences regarding data collection practices vary extensively as well [188]. Arguably, then, TLA should allow users to decide which data collection practices they find sufficiently harmless and useful.

However, given the presumed complexity of TLA-based systems (with multiple learning activity providers collecting various types of data, and using it for numerous different adaptation purposes), it may be difficult for users to manage their privacy in such a system. Hence, TLA should not only afford users the ability to manage their privacy but actively support these privacy management practices as well.

This section makes a case for User-Tailored Privacy as a means to support privacy management practices by highlighting the shortcomings of other means of providing privacy support. This does not mean that these other means should not be implemented; indeed, User-Tailored Privacy builds upon—and works in tandem with—these other privacy support mechanisms.

1.1 Technical solutions

Engineers tend to employ technical solutions to privacy problems. When it comes to personalized systems, these solutions broadly fall into two categories:

- Architectures, platforms, and standards designed specifically to minimize data leakage. These include distributed architectures, portable user profiles, and client-side personalization techniques that provide limited access to and “linkability” of user data [29,66,122,190,191].
- Algorithmic techniques for data protection. These include anonymization, obfuscation, differential privacy, and homomorphic encryption [49,89,192,243,267].

Distributed architectures preserve privacy, but are limited and slow

In Specification Vol. 1 (Section 4) we recommended a system architecture that would divide the personalization tasks between the TLA Processors (meta-adaptations), Learning Activity Providers (macro-adaptations), and the client’s device (micro-adaptations). In this architecture, Processors and Providers have access to the TLA Data Core with user data. Several studies have suggested mechanisms that could be used to “shield” the Data Core from the Processors and Providers using privacy-preserving personalization protocols [66] or hierarchical personalization mechanisms [29]. While these mechanisms effectively preserve privacy, they are slow, and limited in their capabilities (preventing state-of-the-art personalization algorithms).

Client-side personalization is limited, and vulnerable to data loss and theft

A step beyond distributed architectures is to perform all necessary calculations for personalization on the user's own device [53,148,239]. These client-side personalization methods are often limited to content-based personalization, although collaborative filtering strategies exist [49,296,334]. The inference methods that can be used in client-side personalization are limited, though, since there is no direct access to other users’ data. Research shows that users indeed prefer client-side methods as a means to alleviate privacy concerns [190,315], and we recommended the use of client-side personalization for micro-adaptations in Specification Vol. 1 (Section 4). We note, however, that users of client-side personalization systems may be concerned about potential loss or theft of their device [191].

Providing anonymity or pseudonymity is difficult and not always desired

A possible mitigation of privacy concern with TLA-based systems is to allow users to remain anonymous. Fully anonymous interaction is difficult though, since the idea of a sustained personalized learning experience relies on the systems’ ability to recognize users across interactions [286]. Alternatively, users can be allowed to interact under a pseudonym [21,192]. However, it may be possible for other users or third parties to re-identify users based on the data that is being collected by the system, even if this data is devoid of identifiers [241], or based on the output of the system [48]. A means to overcome this problem is *differential privacy*, a model that inserts noise into the user profile, making it harder to re-identify anonymized data records [89,220,267,275,377]. Such methods do however reduce the quality of the predictions [230].

Another problem with anonymity/pseudonymity is that it is not always desired from a social perspective. As mentioned in Specification Vol. 1 (Section 2.1), having users interact under their real name reduces rude or abusive behavior [60]. In formal and diplomatic learning/training settings, pseudonymity is thus not desired. In addition, research has shown that identifiability

enhances users' trustworthiness perceptions for both in-group and out-group members [318], making it easier for users to assess their trust in other users based on both reciprocity and trustworthiness perceptions. In systems that provide anonymity, trustworthiness is harder to judge, so users' trusting behavior is based on in-group reciprocity only. This reduces users' ability to create trust-based relationships with out-group members, thus effectively "shrinking" their social network.

Encryption is slow and may not be trusted

A final technical solution to provide privacy-preserving personalization is encryption. By encrypting data before sending it to the personalization system, the system can provide recommendations without "seeing" the data itself [243]. This solution relies on *homomorphic encryption*, an encryption method that preserves certain mathematical properties of the original data, thereby allowing personalization algorithms to operate on the encrypted data the same way it would on the unencrypted data. This encryption method is slow, though [243]. Moreover, it may be hard for users to understand how a system can provide accurate recommendations without having access to the unencrypted data, and they may therefore not trust this technical solution.

Privacy-preserving technologies do not apply to social networking

The aforementioned privacy-preserving technologies may shield users' data from personalization providers. However, TLA-based systems may collect user data for purposes other than personalization as well. For example, as outlined in Specification Vol. 1 (Section 5.2), TLA-enabled learning applications can provide *social learning experiences*, which allow users to interact with each other on a learning platform [153]. Social learning is often powered by sharing learning status or outcomes between users. An activity provider may thus collect and even expose some user data to other users on the platform to enable social learning. In these cases, there are no purely technical solutions that enable social learning without exposing some personal data to the app or other users.

Privacy-preserving technologies do not necessarily increase disclosure

Another problem with the aforementioned privacy-preserving technologies is that while they protect users' data from being exposed to personalization providers, this in itself does not necessarily mean that users will disclose more information. Indeed, work on client-side personalization shows that users' perception of the privacy afforded by this technology is modest, and has only a very slight impact on their subsequent sharing decisions [190,191].

One reason for the underwhelming effect of privacy-preserving technologies could be the “ironic effect of transparency”. Most privacy-preserving technologies operate in a manner that is invisible to the user. Hence, the only way for these technologies to increase disclosure is for the application to notify users of the fact that their privacy is being respected. Such notice may, however, have an unexpected effect. For example, marketers [1,46,105] have discovered that displaying a privacy label on an e-commerce website—a supposed vote of confidence in the site’s privacy practices—may *decrease* instead of increase purchases. Moreover, privacy policies have been shown to incite privacy concerns rather than easing them [268]. Finally, disclosure justifications have been shown to have a negative effect on users’ disclosure and self-anticipated satisfaction [169]. Arguably, the main reason for these effects is that users perceive justification messages as a warning sign. Similarly, then, notifying users of the use of privacy-preserving technologies may have the opposite effect, and make users *more* rather than less wary about their privacy.

Recommendation: Complement technical privacy-preserving solutions with user-centric solutions

Privacy is primarily an attitude, and privacy management is an inherently human behavior. As such, technical privacy-preserving solutions can never solve all of TLA’s privacy problems. Given the apparent shortcomings of technical privacy-preserving solutions, TLA should also—or rather, primarily—employ user-centric solutions to these problems. The remainder of this section addresses existing user-centric solutions and explains why user-tailored privacy is best among these solutions.

1.2 Privacy by design

Privacy by design 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 [54,201,285,290,308]. Specification Vol. 1 followed a privacy by design process, identifying design suggestions for the operational characteristics of TLA that would minimize privacy concerns. However, Specification Vol. 1 also explicitly noted (in Section 6), that it will be inevitable for TLA-based systems to have a wide array of privacy settings. There are several reasons for this, which are outlined below.

In some cases, privacy is in direct opposition with system functionality

Specification Vol. 1 outlined several means to reduce users’ concerns regarding TLA’s data collection practices. The only way to truly avoid any concerns, though, is to not collect any data at all. This is of course not realistic: personalization and social learning are not feasible without data collection [23,312]. So, while there exist privacy-by-design practices that can increase privacy without reducing functionality or vice versa, at some point privacy and functionality are

in direct opposition, and a choice has to be made regarding how much data collection is justified to provide a certain level of functionality.

Users differ a lot in how they think about privacy

This inherent tradeoff between privacy and functionality has no universal solution, though, because users differ extensively in how they think about privacy (see Specification Vol. 1, Section 1). Evidence for this can be found in the area of personalization, where privacy concerns prevent some users from using personalized applications [315], while others are willing to give up privacy in return for personalization benefits [113,248,266], such as content relevance, time savings, enjoyment and novelty [125,132]. And for some, privacy concerns may not even matter, as long as the benefits are clear [165]. Based on these results, it is clear that some users will go at great lengths (in terms of disclosure) to use TLA’s personalized learning recommendations, while others will disclose a very limited amount of personal information, regardless of the inherent privacy protection or the personalized benefits TLA can provide.

***Recommendation: Where privacy cannot be designed,
provide decision support***

In conclusion: while privacy by design attempts to avoid privacy issues to begin with, privacy is at times in direct opposition with system functionality, and in these cases users often have a desire (and in some cases a legal right or expectation) to engage in a “privacy calculus” [205,206], i.e., to make personal decisions regarding the inherent privacy-functionality trade-offs. In these situations, a system is required to have some controls (e.g. “privacy settings”) as a means to effect this tradeoff.

Likely, a TLA-based learning ecosystem will have numerous components that involve such trade-offs, and therefore require privacy controls [272]. Consequently, users are expected to make a substantial number of privacy decisions when using TLA-based learning systems, and these decisions are not only numerous, but also impactful, since both the sensitivity and the potential benefits involved in these privacy decisions are likely to be substantial.

Given that TLA-based learning systems potentially involve difficult privacy decisions, these systems should support users in their decision-making practices. Specification Vol. 1 (Sections 6.1–6.3) discussed existing privacy support mechanisms that have been developed under the prevailing paradigms of “notice and choice” and “privacy nudging”. Below we revisit these paradigms from a theoretical perspective and contrast them to user-tailored privacy.

1.3 Notice and choice

Where privacy cannot be designed—because it is in direct opposition with system functionality, and users differ in the amount of privacy they prefer to trade off for functionality—privacy experts argue that users must be given controls (e.g. “privacy settings”) as a means to effect this trade-off, as well as a certain amount of information that will help them operate these controls [44,55,207,283,319,371]. This idea of “notice and choice” is also at the heart of existing or planned regulatory schemes [95,353].

Providing notice and choice is not a trivial task. For example, Lederer et al. [207] state that existing systems often make it difficult for people to manage their privacy, and current system designs inhibit people’s abilities to both understand the privacy implications of their use and to conduct socially meaningful actions through them. They outline the pitfalls related to people’s understanding of privacy implications (e.g., understanding the scope of privacy implications and information flow) and pitfalls related to meaningful privacy action (e.g., emphasizing configuration over the action, lack of coarse-grained control, and inhibiting established privacy practice). Therefore, developers should beware of potential issues while designing notice and choice privacy management approaches.

Beyond this, there are more fundamental problems with notice and choice: Researchers have identified a number of paradoxes that make notice and choice very difficult to implement.

Sufficient notice is complex, and the effect of notices is fickle and fleeting

Nissenbaum [245] postulates the Transparency Paradox: simple privacy notices do not contain sufficient information to support people’s privacy decisions, while 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.

Moreover, as mentioned earlier, notice may actually decrease disclosure if it focuses users on their concerns [169], even if the notice is supposed to indicate a positive privacy protection practices [1,46,105,268]. Indeed, John et al. demonstrate that users may process privacy-related information in unexpected ways [142]. Specifically, they show that professional looking sites may garner higher privacy concerns than unprofessional looking sites because the former design reminds users of privacy. Because of this, it is unsurprising that the framing of privacy notices can increase or decrease disclosure, even if the information itself remains the same [12].

Finally, the effect of notice may be very fleeting: Adjerid et al. demonstrate that even the slightest distraction can easily nullify any effect of privacy notices [12]. Indeed, while many people claim to read online privacy policies [135,236], many do not actually read them [13,27,28,117,141,157,297,332], or do not read closely enough to understand them [258]. In other words, requiring each TLA-based system to have a detailed privacy policy will likely not help users in their privacy decision practices.

Users may not always take control over their privacy, and control can lead to misplaced confidence

In Specification Vol. 1 (Section 6.2) we already mentioned the Control Paradox, which states that users claim they want full control over their privacy but often do not actually take control, even when it is offered [68]. Indeed, many users tend to pay little attention to privacy seals [204], social navigation cues [31], privacy assurances [234], and permission requests [98].

Moreover, Brandimarte et al.[39] conducted a study to demonstrate that users disclose more when they feel more in control over their disclosure, even if the disclosure risk is higher. They demonstrate a similar effect when giving users explicit granular control.

Finally, Acquisti and Gross showed that users may not always have the requisite knowledge to understand the potential implications of their sharing decisions. They show that information about an individual's date and place of birth can be exploited to predict his or her Social Security number [86]. If users are unaware of such potential inferences, giving them control to disclose their information can, in fact, make them more vulnerable to privacy risks.

In other words, offering TLA users ample privacy controls does not guarantee that they will use these controls, and if they do, it may not actually help them protect their privacy.

Users' privacy decisions are not always rational

In Specification Vol. 1 (Section 1.1) we noted that many researchers no longer believe that users always make "calculated" privacy decisions, but often employ heuristic decision strategies instead. As such, people's privacy decision can be influenced by:

- The perceived immediacy of benefits, versus the delayed nature of risks ("temporal discounting" [4])
- Information on others' privacy decisions ("social proof" [8])
- The order of sensitivity in which decisions are being made ("foot in the door" and "door in the face" [8])
- The overall professionalism of the privacy-setting user interface ("affect heuristic" [142])
- The available options to choose from ("context non-invariance" [39,171])
- The default setting and phrasing of privacy-related requests ("default" and "framing" effects [9,144,168,200])
- If data protection was mentioned to be higher or lower in comparison to the current situation ("anchoring effects" [12]).

When users are given control unconditionally, these effects may disturb their decision practices. In the context of privacy nudging (and User-Tailored Privacy), however, these effects are used in a targeted manner to guide users' decisions.

Notice and choice are unrealistic in complex TLA-based systems

Privacy controls can range from very coarse (turning a system on or off) to very granular (making a decision about sharing each piece of potentially privacy information with each potential recipient), and there have been many debates in the privacy field as to exactly how much granularity is optimal [26,35,168,171,178,283,306,316]. The same is true for the amount of information that is required for adequate notice [245]. In a sense, the optimal level of granularity is yet *another* trade-off—one for the designer of the system to make [171,178].

Likely, the severe impact of disclosure on the functioning of TLA-based learning ecosystems as well as the privacy concerns of their users means that privacy decisions in TLA-based systems will likely have to be both granular and numerous as a means to support users' "privacy calculus" [272]. However, given the issues presented above, one may argue that users do not actually engage in a privacy calculus [205,206]. More specifically, it would be wrong to assume users will read system and application privacy policies, or that they will use privacy controls to safeguard their privacy within a complex system like TLA. And even if they do, it is not certain that they will do so effectively.

Recommendation: Implement UTP to make notice and choice more manageable

The privacy calculus can still be used as a prescriptive model, though, with the risk/benefit tradeoff serving as an objective function for machine learning algorithms [26,124,166]. In this prescriptive approach, privacy is no longer under the direct control of the user but implemented by an algorithm that will take the user's contextualized preferences into account [180]. From this perspective, User-Tailored Privacy is an "automated" version of notice and choice.

By (semi-)automating the decision process, and potentially also the delivery of privacy-related information, UTP can make notice and choice more manageable. Arguably, UTP relieves some of the burden of the privacy decision from the user by providing the right privacy-related information and the right amount of privacy control that is useful, but not overwhelming or misleading [166].

1.4 Privacy Nudging

The cognitive costs associated with considering all the ramifications of disclosure may hamper decision-making, especially in a complex environment like a TLA-based learning ecosystem.

Privacy nudging attempts to relieve some of that burden by making it easier for people to make the right choice, without limiting their ability to choose freely [5,12,25,344,345]. Some confusion exists in the field of privacy regarding the term “nudging”.

Traditionally, nudges have been defined as changes to the structure, framing, and defaults of a decision environment with the intention of making the “right” decision easier to perform [324]. These nudges are covert changes (users generally do not notice them), and they essentially exploit the decision externalities outlined in Section 1.3. Note, also, that these nudges are related to privacy by design, as they answer privacy design decisions: it is virtually impossible to have a “neutral” structure, framing, or default, so nudging is essentially inevitable.

Interestingly, the term nudging in the field of privacy has been used more frequently to describe designs that steer users in a desirable direction in a more overt manner [5]. Symbolic and textual justifications are good examples of these types of nudges. These nudges also provide a shortcut to decision-making by virtue of being persuasive, but unlike traditional nudges they do require users to consciously process them. Some even suggest that nudges should make an individual aware of the biases, lack of information and cognitive overload that may affect their decision [25]. These types of nudges keep users in the loop completely.

Nudges do not work for everyone and may threaten user autonomy

In Specification Vol. 1 (Section 6.3) we elaborated on the relative success of these different types of nudges, and noted that the more overt nudges like justifications [8,31,169,187,263], privacy seals [90,133,234,277,370], and audience/sentiment feedback [140,329,344,345] fail to have a consistent effect on users’ disclosure and privacy concerns.

Moreover, while the traditional nudges have been found to be more effective [8,144,166,168,200], they have typically only been tested for behavioral impact, disregarding the question of whether they reduced users’ privacy concerns or their privacy decision burden [166]. Indeed, researchers worry that defaults may threaten consumer autonomy, especially when they work outside of users’ explicit awareness [302,305]. These researchers argue for “smart nudges” that match the preferences of most users. However, if users vary extensively in their privacy preferences (as ample research, outlined in Specification Vol. 1, Section 1, suggests), then even a “smart nudge” may face ethical objections.

Recommendation: Implement UTP to move privacy nudging beyond one-size-fits-all

Nudges are an interesting way to help users make the right choice without limiting their decision freedom. However, in most privacy settings, the ‘right choice’ is difficult to define, hence nudges are not welcomed by every user. Users differ in their privacy preferences, and a nudge that is effective for one person is not necessarily effective for a different person. Similarly, users’

privacy preferences depend on the who, why, what, and general user mood associated with the decision [69], so the effectiveness nudges depends on these contextual parameters as well.

Given that nudges do not necessarily work for everyone, it imperative to *tailor* them to users based on their privacy preferences. For instance, TLA users who are privacy minimalists may benefit from permissive data-sharing related default settings to fully exploit the benefits of the personalized learning environment, whilst TLA users who are privacy maximizers may benefit from more restrictive defaults, lest they feel that the system invades their privacy. Similarly, TLA users who are privacy minimalists should not be bothered with justifications, whilst for TLA users who are privacy maximizers, justifications can have a significant impact on their privacy concerns and decision burden.

By adapting privacy nudges to users' individual (and contextualized) privacy preferences, User-Tailored Privacy can move these nudges beyond their current one-size-fits-all limitation. This arguably makes these "tailored nudges" more universally effective. Moreover, by using TLA users' own privacy preferences and/or practices as a baseline, UTP side-steps the ethical problems of threatening user autonomy.

2 Definition: What is User-Tailored Privacy?

TLA crucially depends on the collection of an extensive amount of user data to provide social learning capabilities and personalized learning recommendations [101,287]. Users will only agree to such extensive data collection if they feel that their privacy is adequately protected. Section 1 discusses the shortcomings of existing approaches to privacy:

- Privacy vulnerabilities can in some cases be resolved with technical solutions, but these solutions do not guarantee that TLA users gain a sufficiently positive perception about their privacy to feel confident about disclosing their personal information.
- Privacy by design can alleviate TLA users' privacy concerns through a careful selection of the operational characteristics of TLA-based systems (see Specification Vol. 1), but the inherent conflict between privacy and personalization, as well as variations in user preferences, mean that privacy decisions are essentially unavoidable in TLA-based systems.
- Notice and choice are requisite to allow TLA users to make such privacy decisions but easily become overwhelming in systems that make extensive use of user data for social and personalization purposes, such as is envisioned for TLA-based systems.
- Nudges can reduce the burden of making privacy decisions, but their static nature means that they do not work well for all users and that they may threaten user autonomy.

User-Tailored Privacy (UTP) alleviates these problems through personalization. UTP makes nudges adaptive, by making the user's *own preferences* the basis for the nudge, thereby increasing their adoption and reducing their paternalistic nature. Consequently, UTP makes notice and choice more manageable, by (partially) *automating* the risk-benefit tradeoff inherent in users' privacy decision and presenting the result as a personalized default, justification, or privacy-setting interface design. By providing user-tailored support, it reconciles the need for extensive customizability with users' lack of skills and motivation to manage their own privacy settings.

This section gives a formal definition of user-tailored privacy as a measure-model-adapt framework and explains how this framework can be implemented to provide privacy decision support for various aspects of TLA. Each aspect of the framework is subsequently discussed in more depth in Sections 3–5. Section 6 concludes with the various goals that UTP can serve.

2.1 The UTP Framework: Measure, Model, Adapt

At its core, UTP is an approach to privacy that models users' privacy concerns and provides them with adaptive privacy decision support [166,167,175,180]. With UTP, a system **measures** user privacy-related characteristics and behaviors, uses this as input to **model** their privacy preferences, and then **adapts** the system's privacy settings to these preferences (**Figure 1**). Each aspect of the framework is briefly described here and discussed in more depth in Sections 3 (measure), 4 (model), and 5 (adapt).




Figure 1: A schematic overview of User-Tailored Privacy (UTP)

UTP measures user characteristics and behavior in order to tailor its support to the user and the context of the decision

Existing work has shown that user-tailored privacy critically depends on the evaluation of user- and context-related factors that influence users’ privacy concerns and behaviors. This work has shown that data, user, and recipient are important, but also that for many applications there are additional system-specific factors [86,218,250,261,368] that influence users’ decisions. Due to this context-specific nature of users’ privacy decisions, it stands to reason to also make the User Privacy Model underlying user-tailored privacy context-specific. For TLA, this means that UTP should take into account contextual variables that have been shown to influence TLA users’ privacy concerns and behavior.

One such variables is the data itself. Several researchers have found that people’s privacy concerns are multi-dimensional, meaning that they have different preferences for different types of information [218,250,309]. Furthermore, research shows that these preferences can be summarized into distinct *profiles* of similar users [172,250,360].

The recipient of the information seems to play an important role as well, both in “commercial” and “social” privacy settings [168,173,261], and recipients can also be grouped to simplify the privacy decision problem. For example, on social networks the optimal grouping seems to be Friends, Family, Classmates, Colleagues, and Acquaintances [178], but this clustering might be different for recipients in TLA-based systems.

Furthermore, in certain types of systems, privacy preferences may depend on other contextual factors. For example, researchers have found that time (weekday or weekend, daytime or evening) is an important determinant of users’ willingness to disclose their location [86,368].

Finally, note that user-tailored privacy can operate within organizational constraints, considering existing rules, as well as common practices. This way, user-tailored privacy helps users to select settings that are not only in line with their own preferences, but that also consider the value for the organization that has implemented TLA, as well as existing rules.

UTP models users' privacy decisions using machine learning algorithms and then plans adaptations based on this model

The goal of UTP is to model users' privacy decisions via direct observation of their behaviors or via inference from their attitudes. One way to formalize this, is presented by Dong et al. [87] as a model of the probability of disclosure $p(D)$ by user u of item i to recipient r in context c as a function of the user's disclosure tendency D_u , the sensitivity of the item S_i , the trustworthiness of the recipient T_r , and the appropriateness of the disclosure in this specific context A_c :

$$p(D_{uirc}) = f(D_u, S_i, T_r, A_c).$$

This model can be implemented as a loglinear additive model or a full factorial loglinear model, but more sophisticated implementations can be provided by machine learning algorithms, such as decision trees, Bayesian models, and support vector machines. These algorithms are no different from the standard algorithms discussed in the recommender systems literature.

Once the user's privacy decisions are modeled, UTP should adapt to them. This adaptation procedure could aim to exactly match users' existing decisions, but it could also aim to automate or support synergistic auxiliary behaviors, or even attempt to help users explore new ways to manage their privacy [180].

UTP adapts the privacy settings, justifications, user interface, and/or personalization procedure to users' privacy concerns

There are different ways to provide privacy decision support in an adaptive manner. The most common application of UTP is to adapt the privacy settings or information requests themselves, which alleviates the burden of privacy decision-making. This can be done by making explicit suggestions for certain privacy-related actions that can be taken in a TLA-based system, highlighting the recommended privacy decisions (or hide the ones that are less likely to be chosen), or automatically taking action [180].

Another type of adaptation is to give users personalized justifications for certain settings or information requests [170]. This can involve giving TLA users a relevant explanation about the reasons behind a privacy recommendation, but it can also involve educating them about the risks and benefits involved in a privacy decision. Interface adaptations can also be used, specifically to improve TLA users' ability to control their privacy settings and making certain privacy actions easier to accomplish in the interface of the TLA-based system.

Finally, the TLA processors can adapt their personalization methods to the user’s privacy concerns, selectively restricting the types of personalization a system can engage in based on the collected user data, thereby preventing potential unwanted inferences to be made [343].

Recommendation: Implement UTP pervasively in TLA-based systems

The idea of UTP fits very well within the extensive user modeling approach of TLA. Moreover, given the complexity of TLA, it is likely that user-tailored decision-support is the only feasible solution that allows users to maintain considerable control over their privacy decisions without overburdening them. As such, UTP may be critical to the successful management of privacy issues in the TLA. We therefore recommend implementing UTP pervasively in TLA-based systems. In the remaining two subsections of this section, we discuss various opportunities to do so.

2.2 UTP for Learning Applications

The goal of TLA is to provide a learner with performance support and tailored learning content. To provide this functionality, it requires extensive knowledge of the user’s operating context, and a dynamic model of their knowledge, experience, and learning goals [272]. As discussed in Specification Vol. 1, though, learning is often a creative activity that demands an environment where consequence-free experimentation is encouraged. The continuous tracking of users’ activity can potentially create an expectation of continuous performance review, and thereby quell users’ explorative behaviors. From a learning perspective, TLA faces an important dilemma: extensive tracking allows TLA to support users’ learning behaviors, but the feeling of being tracked can have a negative impact on these behaviors in the first place. In this regard, the goal of UTP is to find the optimal balance between learning personalization and user comfort.

This is not an easy task. As an “open” learning architecture that provides integrative services such as analytics and recommendations to third-party learning applications, TLA is not only tasked with collecting and managing this knowledge about the user and their context; it is also responsible for the potential propagation of this knowledge to (and between) the Activity Providers that host the learning applications [101]. Users may not want all Activity Providers to have full access to their data, hence substantial privacy management functionality is required. UTP can provide a solution to this complex privacy management task.

Privacy should be taken into account when recommending Learning Activities (meta-adaptations)

Meta-adaptations are individualized recommendations to switch from one Learning Activity to another that are based on the learner’s specific needs and progress. One straightforward way in which UTP can help TLA users manage their privacy, is by taking the privacy requirements of

Learning Activities into account when making such recommendations. Particularly, certain Learning Activities may *require* certain types of data for their operation. If users are not comfortable with sharing these data, the system can avoid recommending these activities. Conversely, if users have a history of sharing their data with certain trusted Activity Providers, the system can restrict its recommendations to Learning Activities from these trusted providers, or at least weigh this aspect in its meta-adaptation procedure. This is particularly useful if there are multiple candidate Learning Activities that serve the same purpose. Finally, since trust is often considered transitive within a social context, if a certain Learning Activity is already used (and trusted) by the user's closest colleagues, this too can serve as a basis for recommendation.

Several studies have looked into evaluating Web sites or applications from a privacy perspective. Tsai et al. augment a product search engine with privacy indicators [330]: when users search for a product, the search engine indicates the quality of the privacy practices of the vendor. They find that users of their search engine actively evaluate the privacy of the vendors, and often end up paying a premium for privacy. Egelman et al. [90] further test three different locations for these privacy indicators: embedded in the search results, as an interstitial screen between search results and product page, and at the top of the product page itself. The aforementioned effects are strongest when the privacy indicators are included in the search results rather than on the product page. Cranor et al. implement the idea of a privacy indicator on the Web in the AT&T Privacy Bird, which compares Web sites' machine-readable privacy policy with users' indicated privacy preferences to provide a notice of whether the site the user is currently visiting adheres to or violates their preferences [74].

Kelley et al. translate this approach to smartphone applications by providing additional privacy-related information [158] or even a "privacy nutrition label" [157] to Android app store users. Harkous et al. implement this approach with data-driven privacy indicators, and augment it with an overview of information that can be inferred from the collected data [116].

Note that while all of the described studies provided users with privacy *information* about the Web site or app, none of these studies took an active approach to provide privacy-based Web site/app *recommendations*. With the advent of smartphones, the recommendation of applications has become a thriving research topic, cf. [15,80,109,294,372], but to our knowledge no work in this field has incorporated the apps' privacy practices into the recommendation process.

Given that it is possible to generate privacy evaluations for applications (as demonstrated by Harkous et al. [116]), it should indeed be possible to incorporate privacy evaluations into the TLA meta-adaptation process. This requires a hybrid recommender system that can merge the output of the meta-adaptation process with the privacy evaluations. A very simple solution would be to combine the predicted rating of the Learning Activity with the "privacy rating" of the activity using a static weighted formula. More in line with the UTP philosophy, one could tailor these weights (and/or the privacy rating itself) to the user's privacy preferences.

UTP can support users’ privacy management practices within each Learning Activity

TLA is an “open” learning architecture that serves third-party learning applications. The adaptive behavior of TLA is not limited to recommending Learning Activities to users (meta-adaptations); it also provides opportunities for adaptive behavior *within* these Learning Activities (macro- and micro-adaptations) [101]. A key benefit of TLA is that it provides these Learning Activities with a centralized “data core”, which means that knowledge about the user no longer has to be collected and managed by each app individually.

In Specification Vol. 1 (Section 4.1) we argued that while the *knowledge* about the user can be managed centrally in TLA, the *adaptation logic* for within-app adaptations are better left to the individual Activity Providers, because they may deem this adaptation logic proprietary business information. Hence, in order to provide macro- and micro-adaptations based on knowledge collected about the user, Activity Providers may require access to the data core through the xAPI. As we have argued, though, users may not want all Activity Providers to have full access to their data. Hence, like in many app ecosystems, users should be given control over their privacy by requiring apps to ask for *permission* to use their data.

As such, we proposed the implementation of an “access control matrix” for xAPI in Specification Vol. 1 (Section 6.2). However, we also noted that it would be difficult for users to manage the settings in this matrix. Several studies have investigated mechanisms to help users manage app permission settings. For example, Wang et al. created a comprehensive permission dialogue box for Facebook applications that allows users to grant or revoke access particular types of information [340]. In follow up work they added tooltips with additional information, and clearly distinguished permissions that are considered sensitive [339]. They find that users make substantial use of the provided controls, but installations do not increase, and even decrease when sensitive permissions are highlighted.

UTP can also provide a solution here, by making user-tailored suggestions for app permissions. For instance, Wang et al.’s initial work [340] also included a version of the permission dialogue box that is somewhat personalized: it highlighted permissions based on a comparison between the permissions requested by the app and the permissions the user has granted to other Facebook users: a request would be marked in green if the user has set that information to “public”, and red if the user has set that information to “friends” or “only me”.

UTP’s Learning Activity permission adaptations can also take the form of “smart default” settings. Several studies have investigated this idea for smartphone apps. For example, Liu et al. [216] developed a profile-based personalized privacy assistant for smartphone app permissions. The app groups users into different profiles based on their privacy preferences. Based on these profiles, the assistant recommends permission settings that the user could change. User study results show that the recommendations were adopted by the majority of users. And that the

recommendations led to more restrictive permission settings without compromising on user comfort with these configurations.

An advantage of using UTP to manage users' Learning Activity privacy settings is that it allows for these settings to be more granular without suffering losses in usability: since UTP takes care of the settings, rather than the user, the existence of a large array of settings and/or the continual reevaluation of these settings no longer needs to be avoided. Wijesekera et al. [355] and Olejnik et al. exploit this advantage by creating a system that makes app permissions *dynamic*, allowing context to dictate whether a certain permission can be granted during a specific session. Their approach yields 80% and 95.7% accuracy, respectively. In both cases, this was a very substantial improvement over static permissions.

Recommendation: Develop a privacy API for TLA, and require Activity Providers to make their policies and settings available through this API

In this section we discussed the use of UTP to recommend Learning Activities based on their privacy practices and to recommend specific privacy settings within each Learning Activity. The former procedure requires UTP to have detailed knowledge of the privacy practices of each Learning Activity, preferably in a standardized format that can be used to compare privacy practices between apps. Similarly, the latter requires UTP to have access to the privacy settings offered by each Learning Activity, and preferably to be able to adjust these settings automatically and on the fly.

To implement these requirements, we recommend the development of a privacy API for TLA. Activity Providers can then be required to make their privacy policies accessible in machine-readable form through this API, much like the P3P specifications do for Web sites [75]. Similarly, they can be required to expose their privacy settings as readable and settable properties, so that UTP can tailor them to the user's privacy preferences.

Without such an API, it will be very difficult to implement automated privacy recommendations for Learning Activities. A potential workaround for the privacy policies could be to analyze the text-based privacy policies using deep learning [115]. For the settings, an "overlay" can be developed that automatically fills out the privacy settings as soon as the user accesses the settings page, much like a form auto-completion tool [34,366].

Compared to these approaches, an enforced privacy API is much preferred. Not only does it streamline the recommendation process; it likely also creates a competitive incentive for Activity Providers to enforce a favorable privacy policy, lest their applications are ignored by the meta-adaptation procedure due to poor privacy practices.

2.3 UTP for Social Learning Experiences

Social learning experiences are an important part of TLA [336], as social learning has been shown to improve performance, both offline [143] and online [153]. However, education is one particular area where legislation is quite strict, even in the United States. As such, sharing learning data without the user’s consent can be considered a violation of the Family Educational Rights and Privacy Act (FERPA). Moreover, as outlined in Specification Vol. 1 (Section 5.2), results from social media privacy studies suggest that users are considerate of their social network privacy [313,346]. Hence, in TLA-based applications, not everyone will want to share all their learning data with teammates, superiors, researchers, and/or other colleagues.

Social privacy management is arguably even more complex than other privacy management aspects of TLA. The reason for this is that on social networks disclosure is influenced by factors that are similar to but conceptually different from privacy, such as in-group identification, social capital, and reputation management [358]. Moreover, beyond disclosure, in social networking there are other privacy boundaries that users may want to protect [151,360,363]. Finally, the social process of connecting and interacting with other learners may be influenced by privacy-based personality traits [254].

UTP can be employed to support social-network sharing through automatic friend grouping and audience selection

Research on social networks shows that users share different types of information to a different extent, to different audiences [172,178,250,373]. This is arguably because many social network users have multiple groups of people in their social network, with each group having different expectations about them [331]. Due to the complexity of this task (e.g., Facebook users have a median of 200 contacts each [299]), users tend to have difficulty estimating the reach of information they share on social media, which can result in over-disclosure [346].

This is not for a lack of control. Indeed, most social networks may have a “sharing matrix” (similar to the “access control matrix” described in the previous subsection) which governs the accessibility of each piece of information on the network, including each individual post. However, in many cases, people tend not to know the true privacy impact of their daily social media interactions [312], and share posts in a manner that is often inconsistent with their own disclosure intentions [222].

Several suggestions have been made to simplify the sharing-setting task on social networks. For example, both Facebook and Google+ have a mechanism that enables users to categorize their contacts as a means to more efficiently determine their selective sharing preferences [150,349]. However, a study found that even though users have a clear understanding of its “circles” feature and know how to use it, many users do not use it to the full extent [82,349].

This lack of engagement with selective sharing features is likely due to the fact that users' self-defined categories are not always useful: when prompted, they tend to categorize contacts into semantically meaningful categories, but these categorizations may not be adequate for making privacy decisions [156]. As such, automatic friend-grouping procedures may generate more accurate categorizations. Along those lines, Knijnenburg et al. present an approach for segmenting recipients into groups that are concise from a privacy perspective [178]. TLA can leverage this approach to determine recipient groups for shared learning experiences.

Beyond this, algorithms can suggest which groups to share certain information with. For example, Knijnenburg and Jin recommend audience-related sharing settings for a location-sharing service [177]. A more sophisticated approach is taken by Fang and LeFevre, who use hierarchical clustering on social network structures to predict the most suitable audience for users' personal information [96]. TLA can use similar techniques to predict with whom the user will want to share their learning data.

UTP can support users' social privacy boundary management strategies

Research has shown that social network users employ privacy management strategies that go beyond selective information sharing. For example, Wisniewski et al. show that social network users protect their privacy across multiple relational boundaries [151,360], which include: managing incoming information (e.g. "newsfeed"), managing connections (e.g. friend requests), managing conversations (e.g. chat, comments), managing others' posts about you (e.g. group activities or accomplishments). Further work shows that social network users can be classified into six profiles when it comes to these privacy boundary management strategies [361].

Arguably, TLA users have similar relationship boundary management strategies when it comes to social learning applications. As implementing these strategies may be difficult and time-consuming, they could be supported in an adaptive manner. Users' privacy profiles can serve as a starting point for such adaptations, which can take the form of automations, suggestions [240], or interface adaptations [356].

UTP can support the selection of co-learners and learning groups

An important part of social learning is the selection of fellow learners with whom to connect, and of groups to join. On regular social networks, relationships and group memberships often mirror people's offline relationships [255], and we largely expect similar patterns in TLA-based social learning applications. Note though, that users' learning performance may be influenced by the performance and behavior of fellow learners. Hence, beyond preexisting offline relationships, users may be advised to consider learning-related factors in the selection of groups and fellow learners.

To support this activity, UTP can help recommend social connections—a practice that is well-studied in the field of “people recommendation” [42,57,111]. Such recommendations could for instance take into account users’ interaction style: As mentioned in Specification Vol. 1 (Section 1.3), not all users interact online in the same manner: FYI communicators prefer to keep in touch with others through posting and reading status updates, i.e., without actually having to interact with them, while people who are not FYI communicators prefer to interact in a more direct manner [254]. Hence, someone with an FYI communication style may be more compatible with others that prefer that communication style, and vice versa, so UTP could take this personal trait into account when recommending fellow learners.

Likewise, and more specific to learning, we mentioned in Specification Vol. 1 (Section 5.2) that users may perform better when associating themselves with others that are better (upward social comparison) or others that are worse (downward social comparison) [320]. Therefore, in situations where learners are to work in pairs, UTP could pair upward social comparators with slightly more advanced downward social comparators.

Finally, when it comes to groups, certain people perform better in smaller groups versus larger groups, and open groups versus closed groups. Adapting to this preference is another thing that can be implemented in UTP.

Note that in some cases people will have to work together, regardless of their styles and preferences. When learners find themselves in a situation with such incompatibilities, the fact that each user has their own privacy preferences may result in conflicts between each user’s personalized settings. UTP should try to reconcile these preferences the best way possible; see Section 6.4 of this document.

Recommendation: Develop a social learning API for TLA to support the cohesive management of social privacy boundaries

One goal of TLA is to create a streamlined learning experience across various learning applications. As such, users should not have to re-engage in social network creation and management in each individual Learning Activity. Hence TLA should have a social learning API for the management of friends and other contacts across Activity Providers. This reduces social privacy boundary management practices, as it allows users to manage these boundaries across multiple learning applications at once. With a central API in place, UTP can easily be implemented to further support these privacy boundary management practices, e.g. through co-learner and group recommendation, recipient grouping, audience selection, and other boundary regulation practices.

3 Measuring privacy

The idea behind UTP is to model users’ privacy decisions. In Specification Vol. 1 (Section 1), we provided ample evidence that users vary extensively in their privacy decisions, due to both personal [118,119,351,352] and contextual [26,69,130,145,150,208,211,246,250,260,326,349] differences. In other words: no two people make exactly the same privacy decisions, and even for the same person the decision tends to depend on the context. Moreover, even when the decision is the same, the motivation behind it may differ per user and per context.

To accurately model users’ privacy decisions, one thus has to first “map out” how these decisions depend on personal and contextual variables, such as:

- **The data requested (What)**—Several researchers have found that people have different preferences for regarding the disclosure of different types of information [197,218,250,309], and engage in different protective behaviors to a different extent [361].
- **The user him/herself (Who)**—Research also shows that there exist distinct profiles of privacy behaviors among users [172,250,361].
- **The recipient of the information (To whom)**—The recipient seems to play an important role in users’ disclosure decisions, both in “commercial” and “social” privacy settings [150,168,173,208,261,349].
- **Other factors, usually system-specific**—In certain types of systems, privacy preferences may depend on other contextual factors [26,86,87,368].

This section discusses the existing research on these dependencies and argues that “purpose specificity” can be used as a guiding principle for mapping these dependencies. Given that these dependencies tend to differ per system, we argue that TLA implementers should conduct extensive studies to quantify these dependencies in the context of TLA. Section 4.1 discusses the technical matter of implementing these dependencies in UTP.

3.1 The data (What)

The data, or “what” is being requested, is the foremost factor in determining users’ privacy tendencies. In Specification Vol. 1 (Section 2) we outlined several types of input data that TLA-based systems are envisioned to collect. We also noted that these data may not all be equally sensitive, as they range from public knowledge (e.g. the user’s military rank) to psychological assessments and minute details about training performance. In this subsection, we provide evidence for the fact that users’ privacy concerns and behaviors regarding different types of information indeed vary substantially.

Privacy protection extends beyond disclosing information to other boundary-preserving behaviors [17]. For example, users can manage *relational* boundaries (e.g. friending and unfriending), *territorial* boundaries (e.g. deleting or contesting unwanted content posted by others), *network* boundaries (e.g. hiding or disclosing their contacts), and *interactional*

boundaries (e.g. blocking other users or hiding one's online status) [151,360]. We discuss research that shows that these behaviors also vary substantially per user.

Sensitivity and disclosure tendency vary by type of information

Several researchers have demonstrated a variability in the average perceived sensitivity of different types of information:

- Ackerman [3], Figure 1: people are most comfortable providing their favorite TV show, snack, and e-mail address; they are least comfortable providing their phone number, credit card number, and social security number.
- Metzger [233], Appendix A: sex, first name, education level, and age are considered least sensitive; telephone number, credit card number, banking information, and social security number are considered most sensitive.
- Metzger [232], Table 1: education level, first name, time spent online, and sex are considered least sensitive; social security number, telephone number, street address, and credit card number are considered most sensitive.
- Knijnenburg et al. [166], Table 12: age, gender, height, and parental status are considered least sensitive; birth control use, number of sexual partners, sports activity, and savings are considered most sensitive.
- Knijnenburg et al. [172], Table 12: people are most comfortable disclosing their gender and their favorite movie, food, and music to a retailer; they are least comfortable disclosing their phone number, income, and whether they use birth control.
- Wang et al. [347], Figure 1: people are most comfortable publicly disclosing their interests, groups, religion and links on their social network page; they are least comfortable disclosing their e-mail address, street address, and phone number.

Behaviorally, several researchers have shown that people are more likely to disclose certain types of information and less likely to disclose other types of information:

- Joinson et al. [146], Table 1: people are most likely to disclose their height, season in which they were born, and whether they are left- or right-handed; they are least likely to disclose their number of sexual partners, number of relationships, sexual orientation, and weight.
- Knijnenburg et al. [169], Table 2: people are most likely to disclose their gender, the amount they read, and their age to an app recommender system, and least likely to give a system access to their web browsing, e-mail messages, and credit card purchases. They also show that disclosure decreases as the amount of disclosed data accumulates.
- Kobsa et al. [190], Table 1: people are most likely to disclose their phone plan, relationship status, race, and field of work to an app recommender system, and least likely to give a system access to their microphone, e-mail messages, and credit card purchases.

- Wisniewski et al. [361], Figure 2: Facebook users are least likely to block apps, events, or people; they are most likely to set the default visibility of posts on their walls or in which they are tagged, and to restrict their chat availability.

Generally speaking, interests and basic demographics seem least sensitive, while financial information, contact information and details about a person's sex life are most sensitive. Interestingly, the sensitivity and disclosure tendencies vary by study (e.g. e-mail is among the most sensitive in some studies but among the least sensitive in others). Arguably, this depends on the context of the study, which determines the purpose for which the information is collected. This purpose specificity makes it difficult for TLA implementers to use existing research in gauging the relative sensitivity of data collected by TLA-based systems. Instead, they should conduct their own studies for this purpose.

Privacy behaviors are multi-dimensional

To support UTP, it would be useful to know to what extent users' disclosure tendencies regarding different types of information are interdependent: If they are decidedly not, then UTP must make separate predictions about users' disclosure tendency for each individual piece of information; if they are, then UTP can use users' known tendency to disclose certain items to predict unknown tendencies regarding other items.

Some studies treat the disclosure of each item as a separate decision and make no assumptions about correlations between these decisions [8,146]. Others treat disclosure behaviors as a unidimensional scale by summing individual disclosures into an overall measurement of disclosure tendency [142,232,234]. Again others group items into a number of distinct scales, based on the sensitivity of the items [164]. Finally, some studies distinguish different types of information [104,247,309], assuming that certain types of information are more closely related than others. In the latter category, some researchers have even been able to statistically validate such a multi-dimensional structure:

- Khalil and Connelly [160] suggest two dimensions of information disclosure in their context-aware telephony service: location/activity disclosures, and conversation/company disclosures.
- Buchanan et al. [43] uncover two dimensions of privacy behavior among 12 items: a general caution dimension and a technical protection dimension.
- In a study on community-based video surveillance, Koshimizu et al. [197] find seven factors of attitudes towards social and authoritative surveillance that relate to the responsibilities and concerns regarding the communal use of a surveillance system.
- Lusoli et al. [218] find four dimensions in users' preferences towards the collection of personal data by e-commerce sites: social information, biographical information, sensitive information, and security information. They also find six dimensions of privacy protection behaviors: reactive practices (e.g. spam- and spyware filters), proactive

practices (e.g. contacting websites about their privacy practices), withholding information, minimizing disclosure, avoiding the use of technology, and lying.

- In a study on disclosure behavior in an interpersonal privacy context, Olson et al. [250] find six different dimensions among 40 types of information, but the uncovered dimensions defy succinct description.
- Knijnenburg et al. [172] study dimensionality in three contexts. In a mobile app recommender system, they find two dimensions: disclosure of context data and demographics data. In a study on Facebook disclosure tendencies, they find four dimensions: Facebook activity, location, contact info, and life/interests. Finally, in a study on disclosure tendencies towards e-commerce retailers, they find four dimensions: health-related information, personal interests, job-related information, and contact information. Interestingly, while most correlations between dimensions are positive, they find that contact information is negatively correlated with the other dimensions.
- Wisniewski et al. [361] study the dimensionality of Facebook privacy protection behaviors. They find 11 dimensions among 28 behaviors: altering one's News Feed, moderating one's Timeline/Wall, reputation management, limiting visibility of information posted about the user by others, blocking people, blocking apps or event invitations, restricting chat availability, selective sharing, custom friend lists, withholding contact information, and withholding basic information.

These results are promising for the implementation of UTP, because they show that disclosure is more complex than a single tendency, but also not completely unstructured. The ostensive differences between the uncovered dimensional structures also show, though, that the underlying dimensionality of disclosure behaviors varies by context. This makes it difficult for TLA implementers to directly use existing research in gauging the dimensionality of data collected by TLA-based systems. Instead, they should conduct their own studies for this purpose.

***Recommendation: Study privacy behavior (and dimensionality)
in TLA-based systems to provide a basis for UTP***

No existing research on differences in (and the dimensionality of) privacy behaviors has been conducted in the context of personalized learning environments, so this research cannot be directly translated to TLA-based systems. We therefore strongly recommend that TLA implementers spend time researching users' privacy behaviors as a basis for input to the potential implementation of UTP.

As such, we envision that the “what” dimension can be implemented in UTP in four ways. The simplest method is to use content-agnostic algorithms that figure out the relative occurrence and structure of privacy behaviors on their own accord. These algorithms are likely to fall prey to the “cold start problem”, though, and may also be subject to overfitting (see Section 4.4).

Another method is to conduct user research on data collected during initial usage of the TLA-based system, using the methods presented in [172,361]. This allows one to carefully consider the dimensionality of the behaviors in situ, but it requires a delayed implementation of UTP.

A third method is to conduct scenario-based multi-factorial experiments to quantify the relative sensitivity and the dimensionality of various types data. This method does not use real behavior as a ground truth but can be conducted prior to implementation.

Finally, TLA implementers, with appropriate care, can attempt to translate existing findings to TLA-based systems. For example, the TLA processors are essentially app recommender systems, and thus may take note of the fact that users tend to disallow such systems access to their e-mails and credit card purchase. In such systems demographic data is regarded as a different type of data than context data. Similarly, a social learning component is in many ways similar to a social network, so developers of such a component can consider Wisniewski et al.'s [361] eleven dimensions in their design process (cf. [356]), and also take note of the fact that users tend to mostly engage in managing direct access (e.g., via chat) and posts about them by others (e.g., posts on their timeline, or posts in which they are tagged).

3.2 The user (Who)

In the previous subsection we demonstrated that the extent to which users disclose information (and, more generally, engage in privacy-related behaviors) differs substantially based on the type of information that is being requested (or, more generally, the type of behavior). In this section we provide evidence for the fact that these behaviors also vary substantially by user (the “who”). And just like the “what”, the “who” often shows some inherent structure, which can be captured in *disclosure profiles* or *privacy management profiles*.

Users differ in their level of privacy concern and behavior

Arguably the first attempt to measure between-person differences in privacy attitudes was the Equifax survey by Westin and Harris & Associates [351], who used a short questionnaire of three items across several surveys to classify people into three broad categories: privacy fundamentalists, pragmatists, and unconcerned [118,120]. This approach, while widely used even today, has received considerable critique [365], most importantly revolving the argument that a low-medium-high classification of privacy concerns is too simplistic to be useful in most cases [172].

Culnan expands upon the Equifax survey with two items from Smith et al. (cf. [300]) to measure users' loss of control and unauthorized secondary use of personal information [79]. Smith et al. [301], in turn, extend this scale to 15 items, creating the Concern For Information Privacy (CFIP) scale, which measures collection concerns, unauthorized access, fear of accidental errors, and secondary use. Finally, Malhotra et al. [223] further extended the CFIP scale and adapt it to an Internet environment, to create two scales: a 6-item scale for General Information Privacy

Concern (GIPC), and a 10-item Internet Users Information Privacy Concern (IUIPC) scale measuring collection, control, and awareness.

Behaviorally, many researchers have remarked that people differ substantially in how they deal with privacy [3,47,250]. Indeed, most of the works listed in Section 3.1 report sizable standard deviations in their estimates of disclosure tendencies and privacy behaviors.

Users can be clustered into a number of privacy profiles

One extension beyond the simplistic “levels of concern” approach is to create privacy clusters or “profiles”: using machine learning techniques, users can be sorted into groups that demonstrate similar behaviors. In some cases, this approach builds upon the idea of dimensions; people are clustered based on their values of the different dimensions. For example:

- Knijnenburg et al. [172], Study 1, cluster the users of their app recommender on the two dimensions of demographic data and context data. They find four clusters: three of users who score low, medium, or high on their likelihood to disclose either type of data, and a fourth cluster of users who are likely to disclose demographic data, but unlikely to disclose context data.
- Knijnenburg et al. [172], Study 2, cluster Facebook users on four dimensions (Facebook activity, location, contact info, and life/interests), and find five clusters: users with a tendency to disclose everything, users who disclose everything except contact info, users who disclose activity and life/interest only, users who disclose activity and location only, and users who are unwilling to disclose anything.
- Knijnenburg et al. [172], Study 3, cluster the e-commerce users on four dimensions (health, interests, work, and contact info). They find four clusters: users with a tendency to disclose everything except contact info, users with a tendency to disclose *only* contact info, users who disclose their interests and contact info, and users with a moderate level of disclosure on each dimension.
- Wisniewski et al. [361] cluster Facebook users on their 11 dimensions of Facebook privacy behaviors. They find six profiles: privacy maximizers, selective sharers, privacy balancers, time savers/consumers, self-censors, and privacy minimalists.

Alternatively, people could be clustered directly on their behaviors. For example, Bahirat et al. propose various clustering algorithms to sort users of an IoT privacy settings system into a number of categories [24]. Their optimal, 3-category solution has one cluster of users who want to prevent all IoT data collection; one cluster of users who are okay with the collection of certain types of data, but only if done by their own devices; and a third cluster of users who are okay with data collection by their own devices, friends’ devices, their employer or school’s devices, or devices of nearby businesses, as well as colleagues’ devices (but only if it is done for certain predetermined reasons).

The presented approaches to privacy profiles are nascent, and not without criticism. First of all, the uncovered profiles are static, while people’s privacy behaviors tend to change over time (cf. [313]). Moreover, not every person may fit perfectly into a single profile, and “fuzzy profiles” may create more accurate solutions [323]. Similarly, UTP can go beyond profiles altogether, and provide personalized predictions for each individual user. That said, profiles can still provide a useful input parameter for UTP.

Another problem is that clusters/profiles are completely data-driven, which means that they fall prey to the same “cold-start” and overfitting problems (see Section 4.4) as the dimensionality approaches discussed in Section 3.1. The remainder of this subsection circumvents this problem by analyzing traits and aspects that underlie the differences in users’ privacy behaviors.

Cultural differences have a significant impact on users’ privacy concerns and behaviors

At a macro-level, there are important cross-cultural differences in privacy behavior. For instance, Kim and Yun found that Koreans avoid information access by unwanted friends by temporarily closing their personal profiles rather than unfriending someone because of *Jeong*, a feeling that stresses relational interdependence [161]. Similarly, a recent study by Saway et al. [284] shows that Asian users (China, Japan, South Korea) are likely to exhibit less private behavior online and to be less concerned about online privacy. Users from France exhibited the most secure behavior, while Americans and users from the Emirates portrayed behaviors that were less secure than the French.

Mostly, though, differences in privacy behavior have been described in terms of differences in certain universal cultural dimensions [127,289]. The most prominent dimension is individualism/collectivism, which stresses the balance between a person’s autonomous needs and goals and those of the community. This research finds that users in individualistic cultures generally have more privacy concerns than those in collectivistic cultures, and are consequently less likely to disclose person information and more likely to adopt privacy management behaviors [61,212,235,237,269,310].

Notably, Li et al. [212] study information disclosure across eight different countries, and find that people from collectivistic countries (China, India) find it acceptable to disclose personal information to the government or employers, as they tend to “subjugate individual rights and goals for a sense of commitment to the group and of self-sacrifice for the common good.” In contrast, for people from individualistic countries (US, Germany) this is unacceptable, as they feel less obligated to disclose personal information if there is no specific demand. The authors indicate that these people are more willing to disclose for paid services, or for services with a pre-existing relationship to the service. In addition, users from individualistic countries find computers to be more relevant for a proper control of the disclosure of personal data, while in collectivistic countries mobile devices are more preferable and more likely to be used in data disclosure.

In the use of social media, people in collectivistic cultures are more likely to disclose personal information [269], be concerned with fake identities [347], limit friends to close ties [62], and have a higher level of intimate self-disclosure [62]. Users in individualistic cultures, on the other hand, are more likely to be privacy-concerned and are therefore more likely to adopt privacy protection behaviors [310,347], tend to have large networks with a variety of relationships [63,162], use social media for entertainment [162].

As TLA is employed in different cultural settings, it would be important to include cultural variations as part of UTP. As Li et al. [212] demonstrate, this will likely result in a significant improvement in the accuracy of privacy predictions.

Demographic differences have a significant impact on users' privacy concerns and behaviors

At a more individual level, research shows important demographic differences in privacy behavior. For example:

- A study by Hoy & Milne [129] found significant gender differences in young adults' privacy beliefs, reactions to behavioral advertising, and personal information-sharing and privacy protection behaviors on Facebook.
- Demographic variables were considered in an early study of privacy concern among job applicants [279]. Privacy concerns were measured regarding aspects of participants' family and background, personal history, interests and values, financial management, and social adjustment. The study found the level of concern regarding each type of concern differed significantly based on age, gender, education level, rural/urban background, and income level. The findings demonstrate demographics influence not only privacy attitudes overall, but also across specific domains.
- A study by Saway et al. [284] surprisingly shows no effect of gender and rural/urban background on security behavior. They do however find an effect of income: people with income level with at least 60\$K per year exhibit more security behaviors than people with a lower income.
- In a rare study on differences in users' reactions towards disclosure justifications, Knijnenburg and Kobsa [170] show a strong difference between males and females: While males with a low disclosure tendency are not persuaded by justifications, females with a low disclosure tendency can be convinced with an explanation. For males with a high disclosure tendency, a justification based on usefulness may be effective, while for females with a high disclosure tendency it is best to refrain from using any justifications.

Some of the effects of demographics may be explained by users' previous experience/comfort level with technology [314]. Sundar and Marathe compared the psychological appeal of personalization and customization. System-initiated personalization (SIP), is conducted by the system automatically using an individual's preferences and prior behavior. User-initiated customization (UIC), on the other hand, is done by users of systems, rather than the system

itself. Sundar and Marathe note that although both SIP and UIC ultimately provide users tailored content, user experience differs between the two. They suggest that with SIP, the outcome (i.e. tailored content) is key and the process does not matter. In contrast, with UIC, the affordance of user agency is most important in appraisals of a system. As such, the concept of “power users” explained why some users are satisfied with SIP and others prefer UIC. Power users—users who are more knowledgeable about computing and more comfortable managing system controls—had a greater appreciation for UIC, whereas non-power users, having less expertise and experience, had a more positive experience when SIP was implemented. This is in line with Knijnenburg et al. [166], who show that users with low domain knowledge have more trust in systems that are less user-controllable, while users with high domain knowledge have more trust in systems that are more user-controllable. The key implication of this study for UTP is the importance of identifying power users. If power users can be identified, the need to highlight privacy assurances and alleviate risk perceptions would be even greater for those users.

Demographic aspects are important to consider when employing UTP as part of TLA. Age and income may be particularly important, as these variables tend to correlate with military rank. Gender is of course also important, especially if one esteems not to create a gender-biased privacy recommender. While research on demographic differences in privacy concerns and behaviors is plentiful, it is at times contradictory. We therefore suggest that TLA implementers carefully investigate this aspect.

Personality differences have a significant impact on users’ privacy concerns and behaviors

Like demographics, personality has certain effects on users’ privacy concerns and behavior, but the effects seem rather inconsistent between studies:

- Studies have found that extraverts are more likely to use positive social media [71,278,282], while other studies have found no effect [149,196,253] or even that extraverts are less likely to disclose or spend time online [45].
- Results for agreeableness are equally divided: some find that it increases privacy concern [196,253], some find no effect [18,196], while others find that it decreases it [149]. Females who score low on agreeableness are less likely to upload pictures to Facebook [18].
- Those who are conscientious have higher privacy concerns [149,253] (although some studies find no effect [18,196], but also perceive a higher ease of use of social networks [278].
- Neuroticism increases privacy concerns via computer anxiety [196,253] (although some find no effect [18,149]), and has a negative effect on social media use [71] and posting photos [281].
- For openness to experience, some find a positive effect on concerns [279] via computer anxiety [196,253], but it also has a positive effect on social media use and disclosure [18,71].

The vast differences in effects of personality on privacy concerns and behavior may be due to different theoretical interpretations of privacy. On the one hand, in an early study of the personality correlates regarding privacy concern [279], Rosenbaum argues that personality represents “the type of image an individual predicts of himself in his interpersonal relationships” (p. 337). Therefore, those who are (or at least choose to present themselves as) more outgoing, cooperative, and forceful would perceive questions as less invasive.

Junglas et al., on the other hand, explain personality effects in terms of protection motivation theory: When faced with a threat, people assess the threat by estimating their vulnerability to the threat, as well as the seriousness of the threat. A coping appraisal then occurs in which a person determines how effective the response to the threat would be and the competence that would be required to execute that response. In this interpretation, people who are more outgoing, cooperative, and forceful would be more motivated to protect themselves [149].

The latter effect may also relate to self-confidence, which has been shown to have a large effect on users’ security behavior-intentions [284], even larger than users’ actual knowledge about computer-security. Self-confidence may increase disclosure if users feel confident about protecting themselves against privacy violations.

Personality may also influence users’ amenability to persuasive messages. For example, Orji et al. [252] found that tailoring persuasive games to gamer type increased the effectiveness of motivating healthy behavior. By analogy, personal characteristics of users may be important in presenting UTP recommendations in a persuasive manner.

The aforementioned research relates privacy concerns, social media use, and disclosure to general personality traits (e.g. the Big Five). Some research has shown that more communication-specific traits such as the “FYI communication style” are a more stable correlate of privacy-related behaviors. Users with a “high FYI” trait prefer to communicate passively, indirectly, and publicly, e.g. via status updates, while people who score low on FYI prefer to communicate actively, directly, and one-on-one, e.g. via text, voice or video-chat [254].

***Recommendation: Study differences in users’ privacy behavior
in TLA-based systems to provide a basis for UTP***

For User-Tailored Privacy to succeed, TLA implementers must carefully consider the “user-tailored” part. Hence, they should study and incorporate inter-user differences in privacy concerns and behaviors. Profiling is an interesting means to do this, because (unlike using demographics or personality) it requires little additional data beyond users’ actual disclosure behaviors. We have noted, though, that profiles fall prey to cold start and overfitting problems.

A possible solution would be to adopt profiles uncovered by existing research. Profiles are likely context-dependent, though, and thus hard to generalize to TLA directly. However, we can use

them as design personas. For example, in Wilkinson et al. [356] we developed such personas for TLA based on Wisniewski et al.'s [356] profiles. Table 1 shows how these personas can apply to TLA-based systems.

Table 1: TLA privacy personas, based on Wisniewski et al.'s profiles.

	Suggestion
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.

That said, UTP results are likely to be more precise if they are based on profiles that have been determined based on real TLA users. As mentioned in Section 27, this can be done by conducting user research on data collected during the initial usage of the TLA-based system, and then using the methods presented in [172,361]. This allows one to create profiles based on the behaviors in situ, but it requires a delayed implementation of UTP. Another method is to conduct scenario-based multi-factorial experiments to create profiles. This method does not use real behavior as a ground truth but can be conducted prior to implementation.

Beyond profiles, we recommend including cultural, demographic, and personality-based variations in privacy concerns and behavior as part of UTP, especially since TLA-based systems are envisioned to be deployed in environments that are diverse and multicultural.

3.3 The recipient (To whom)

A third important variable that influences users' privacy concerns and behaviors is the recipient of the information. In this section we provide evidence for the fact that users' privacy concerns and behaviors vary substantially depending on the recipient, guided by the principles of trust. And just like the "what" and the "who", recipients can be clustered into groups or "circles" which imposes some structure on the effect of recipient.

Users' privacy concerns and behaviors differ substantially by recipient

The fact that users' privacy concerns and behaviors differ substantially by recipient has been most prominently investigated in the field of social media. Several researchers have shown that users make inferences about the value and purpose of the information for a requester/recipient in their decision of what to disclose to whom [69,82,198,208,250,260,261]. Most social networks therefore allow users to share specific posts and profile items with specific recipients [221,373].

Privacy decisions are conceptually different when the recipient is a business or an application. Whereas decisions regarding people as recipients are usually governed by social conventions, decisions about what to disclose to applications are governed by information privacy concerns. Despite these different governing principles, users' privacy preferences tend to vary for different types of apps as well [158], and the most prominent platforms (i.e. iOS, Android, and Facebook) demand users to grant permissions to apps one by one [371]. Similarly, there is a substantial stream of research on granting selective permissions to different types of Web sites [20,35,75,193].

Recipients can be clustered into a number of groups or “circles”

To simplify disclosure decisions, Facebook and Google+ both allow users to categorize their contacts as a means to more efficiently determine what to share with whom ([150,349]). On both networks there are standard categories, and users can also create their own categories. Note though, that this feature is rarely used [82,312,349], because when people are prompted to categorize their friends into semantically meaningful categories, they often create categories that are inadequate for making privacy decisions [156].

In response, researchers have explored the idea of supporting users with “privacy setting wizards” [359], and with system-defined categories based on an analysis of the structure of their social network [96]. These kinds of systems could be adopted as part of UTP. Knijnenburg et al. suggest an alternative mechanism for segmenting recipients which is based on the statistical principles of convergent and discriminant validity [178].

Grouping by type has also been proposed for automating privacy decisions regarding different types of Web sites in the P3P protocol [20,35,75,193], but this protocol has lost its relevance in recent years. Interestingly, there has been no proposal to group (mobile) app privacy decisions by type of app.

Users' privacy concerns and behaviors towards various recipients are governed by trust

Users' privacy concerns and behaviors regarding various recipients—people or apps/businesses—are governed by the concept of trust [231]. In TLA, this concept can relate to the user's direct team versus other colleagues or known versus new learning applications.

Trust in learning applications is important in any learning environment, but especially in TLA-based systems that introduce personalized recommendations [33,378]. Wang and Benbasat [341] describe three types of trust that can be distinguished in this context: Competence, benevolence, and integrity. Users are capable of rating recommender systems on these dimensions [30], and users' trust in a personalization provider tends to significantly influence their disclosure tendencies [19,40,84,169,194,213,315,335,369,376].

In interpersonal situations, trust also influences users' tendency to answer requests for personal information [147], self-disclosure [50,269], and tendency to exert control [65]. Recipient grouping and trust have also be combined, noting that within each type of recipient, there may be some who are more trusted than others [270].

Trust in social settings can be categorized as cognition-based trust (reputation), personality-based trust (disposition), calculative trust (rationalization), institutional trust (normality and assurances), and knowledge-based trust (familiarity). Different types of trust seem required to facilitate different types of interactions, e.g. interpersonal communication, group communication and mass communication [59]. Likewise, trust seems to be influenced by the type of social media it occurs on: in closed social networks, people reported greater cognitive trust in user-generated content than in marketer generated content. In open social networks, people have more emotional trust in marketer generated content than in user-generated content [64].

Recommendation: Study users' trust of various TLA-based data recipients

Users of TLA-based systems will likely vary in their level of trust in the various recipients of their personal information, and UTP should take these variations into account. Research on recipients has been largely split between people as recipients (i.e. social media settings) and apps/businesses as recipients. Both of these lines of research are relevant for TLA, because as we mentioned in Specification Vol. 1 (Section 5), data in TLA can be shared with many different entities: e.g. applications, peer users, superiors, researchers.

In both lines of research, grouping/categorization has been proposed as a means to reduce the complexity of making recipient-specific privacy decisions. Interestingly, though, most systems leave the grouping up to the user themselves, and there is very little research on what groupings are most efficient for privacy decision-making (with [178] as a notable exception), let alone what types of recipients are considered generally most and least trustworthy. As such, TLA implementers should study this aspect themselves, by tracking users' privacy decisions in early implementations of TLA, or by running factorial scenario-based experiments on recipient categorization and trust.

3.4 Other factors (system specific, purpose specific)

Beyond “what”, “who”, and “to whom”, there are other contextual variables that can influence privacy concern and behavior. Rationally, such variables may have an influence on how useful or beneficial the revelation of information is. For example, TLA-based system users may find it less useful to share their location with fellow learners during their time off, since they are not available for collaborative learning during those times anyway. A similar effect was found by Xie et al. [368], who show that on weekdays, users are more likely to share their location with colleagues, while on weekends, they are more likely to share with family. Xie et al. argue that this

is due to users' likely availability to colleagues during work hours, versus family during their time off.

Context may also have an emotionally mediated effect: since privacy concern and behavior is based on a self-relevant evaluation of threat and trust, they may be influenced by how the user specifically feels at that specific time, location, etc.

Contextual variables seem most important for situations where information is continuously tracked (e.g. location-sharing services, heart rate monitors)—in these cases time and place are most relevant. Contextual variables can also apply to other situations, though. For example, TLA users' decision to share test results with peer learners may depend on their mood.

Time (of day, of week) influences users' privacy concerns and behaviors

Time as a contextual influencer of privacy concerns and behavior has specifically been investigated in the context of location-sharing. In TLA-based systems, time can play a role when it comes to continuous tracking of learning/training activity.

In the context of location-sharing, time-of-day is most influential on users' concerns and behaviors, but day-of-week also contributes to the variability in users' location-sharing preferences [26]. The former can be simplified to "business hours" versus evening/night [368], while the latter can be simplified to weekday versus weekend [26,368].

Researchers who have developed privacy-setting interfaces for location-sharing services have shown that time is the most prominent variable people use to specify rules about the disclosure of their location [316,359].

Location influences users' privacy concerns and behaviors

Similarly, location has also primarily been studied as a contextual variable in location-sharing systems, and these studies have shown it to be an important factor, especially in combination with time [26]. Again, in TLA-based systems location can play a role when it comes to continuous tracking of learning/training activity.

Like time, location can be simplified into a series of types of locations [317,368], with users making the most prominent distinctions between "home" and "work", especially when asked to share their location during business hours [261]. Among other locations, certain places are obviously more sensitive than others (e.g. hospital, place of worship, bar) [368]. Moreover, Toch et al. show that users' privacy concerns and behaviors regarding location-sharing are related to "entropy": the more uncommon the place is, the higher the concerns [326].

Other contextual factors may influence users' privacy concerns and behaviors as well, but little research has been done on this topic

The influence of time and location on privacy concerns and behaviors is fairly specific to systems that track the user continuously, such as location-tracking and activity-tracking systems. There may also be contextual factors that are specific to learning/training systems. One example is the user's performance: research shows that users are more likely to share things of which they are proud [177,249]. Hence, if a user performs well, they may be more willing to share their progress or performance.

TLA-enabled learning experiences are envisioned to be multi-device experiences [106] including smartphones, smart TVs, eBooks, smart watches, and a multitude of other devices. Especially when it comes to *displaying* learning recommendations, users' preferences may depend on the device they are using. As mentioned in Specification Vol. 1 (Section 3.2), users may want to restrict notifications on personal devices (as compared to work devices) and may be even more wary about devices that are with them even in private moments such as their smart phone or smart watch. We also mentioned that certain devices might be connected to a communal or publicly visible display (e.g. a smart TV), in which case users may want to restrict the notifications that may leak sensitive information.

Many contextual influences on privacy concerns and behavior are due to purpose specificity

Many of the described contextual influences on privacy—and many of the influences of data (what) and recipient (to whom) as well—are due to the user's inference of the *intent* behind the request for information. Nissenbaum alludes to this idea in her concept of *contextual integrity* [244,246], arguing that “contextual integrity is defined in terms of informational norms” (p. 140) which “render certain attributes appropriate or inappropriate in certain contexts, under certain conditions” (p. 143). These normative influences explain the idea of “purpose specificity” [246:140]. Indeed, many papers demonstrate how particular synergistic combinations of contextual variables result in higher levels of disclosure, due to this purpose specificity:

- Consolvo et al. find that people are more willing to disclose their location, and disclose it at a finer granularity, to people who are in the same city as them (i.e., for whom this information is arguably more useful) [69].
- Knijnenburg et al. show that users are more likely to disclose the type of information that matches the purpose of the Web site requesting it. Specifically, they find that people are more likely to disclose their personal interests to a blogging community, their job skills to a job search website, and their health-related information to a health insurer [173].

- Time, place, and recipient also show some synergy, with people being more comfortable sharing workplace and business-hours activities with colleagues, home and evening activities with family, and away and weekend activities with friends [368].
- Sleeper et al. find that Facebook users use various audience-limiting and reaching strategies to actively target their posts to an interest-based audience, where they try to select recipients that share an interest in the topic of the post [298].

Recommendation: Find the best contextual factors by studying purpose specificity

Beyond “what”, “who” and “to whom”, other contextual factors are likely going to influence privacy concerns and behaviors in TLA-based systems. Although TLA is a job-related technology, TLA-based environments are envisioned to track users continuously (even outside business hours), so time and location are likely going to influence users’ privacy concerns. Specifically, users will likely be more permissive regarding on-the-job tracking (at work, during business hours), but may be less permissive with tracking in their personal lives (at home, outside business hours).

Aside from this, TLA implementers should study other contextual factors that may influence users’ privacy decisions. But with contextual data abound, where should they focus their attention? Purpose specificity can provide a guiding principle here: TLA implementers can anticipate specific synergies that would lead to increased or decreased disclosure, and then inspect the variables involved to see if these synergies indeed hold. Examples of a “what” + “to whom” synergy would be: sharing my pedometer data with my personal trainer, sharing my study progress with my peers, and sharing my test outcomes with my supervisor. Such behaviors can be observed to see if they need to be incorporated in UTP.

4 Modeling privacy

In section 3, we explored how users’ privacy decisions depend on various contextual factors including the user, the data requested, and the recipient of the data. This section covers the technical aspects involved in the prediction/estimation of TLA users’ privacy behaviors and/or attitudes based on these factors. We will not cover specific algorithms (these essentially boil down to existing machine learning and recommendation algorithms), but specifically address aspects of the prediction problem that are unique to the domain of privacy prediction.

The task of modeling TLA users’ privacy behaviors and/or attitudes is not a trivial one. Importantly, the following aspects need to be considered:

- **The types of input used by the system**—Having established the important dependencies of users’ privacy decisions in section 3, the next step is to observe or elicit these decisions and dependencies in a TLA-based system, so that they can serve as input for the UTP algorithm.
- **The algorithm used to model users’ behaviors**—Most prominently, UTP can take either a collaborative filtering approach or a case-based reasoning approach, and each approach has its own strengths and weaknesses.
- **The target of the algorithm**—Once users’ decisions are modeled by the algorithm, the UTP approach can aspire to different targets: support the user’s current behavior, supplement their current behavior with complementary activities, or venture into new realms with the goal of potentially altering their behavioral patterns.

This section will also cover potential problems with the practice of modeling users’ privacy decisions. Specifically, we will address the potentially dynamic aspect of TLA users’ privacy decisions (i.e. their privacy preferences may change over time), the way the algorithm may balance the cost of over- versus under-disclosure, the potential tradeoff between privacy and other user and/or system goals, and the impact of traditional machine learning problems like overfitting and the cold start problem on UTP specifically.

Section 5 will conclude our triptych of Measure, Model, Adapt by discussing different ways in which UTP can implement adaptations in TLA-based systems.

4.1 Types of input

What types of input data allow UTP to accurately model users’ privacy decisions, including the dependencies that influence these decisions (as covered in section 3)? In this section, we cover various types of input data, so that TLA back-end developers can put the necessary functionality in place to facilitate their collection.

There are two ways to model users’ privacy decisions: direct observation of user behaviors, or inference from users’ attitudes. Likewise, there are two types of dependencies that can be collected to increase the precision of these user models: user traits and context data.

UTP can model users' decisions to disclose/withhold information, to allow/reject tracking, and to use/avoid privacy features

The goal of UTP is to support users' privacy decision-making practices, so these decisions themselves are a useful input for the user model. Users' privacy decisions often manifest themselves in observable behaviors.

The most rudimentary privacy decision behavior is users' decision to disclose or withhold information. Such behavior has been successfully used in user privacy modeling. For example:

- Fang and LeFevre [96] modeled Facebook users' willingness to share their profile information with each of their friends in order to automate the selection of their sharing settings.
- In a similar fashion, Ravichandran et al. [271] modeled users' contextualized location sharing decisions in an attempt to generate a set of personalized "sharing rules".
- Knijnenburg et al. [172] modeled users' disclosure of various items in three different datasets. They found that small group of distinct disclosure profiles could be extracted from the decisions in these datasets.
- Knijnenburg [166] observed whether users would disclose demographic information to a healthy-living recommender to model their disclosure tendency and adapt the order of demographic questions. If a user decided to skip a certain demographic data request, then subsequent requests would ask for less private information.
- Zhao et al. [375] modeled Facebook location check-ins to give privacy recommendations.

Another privacy decision behavior is users' decision to allow or reject a certain collection or sharing of data to occur automatically. Such behaviors occur in systems that track users' location, incoming email or messages, or other behavior over time. For example:

- Lee and Kobsa [209] asked study participants to decide whether to allow or reject certain IoT data collection scenarios, and Bahirat et al. [24] used this data to create privacy profiles for an IoT privacy-setting interface. In [210], Lee and Kobsa took this approach one step further by modeling users' reactions to live IoT data collection requests, given to them via Google Glass while walking on campus.
- Liu et al. [215] use users' Android app permissions to generate seven permission profiles.
- Friendship request decisions are a special example of allow/reject decisions in the realm of social networks, where it is a peer user (rather than a system) who is allowed/rejected to track the user's posts. For example, Dong et al. [87] model "disclosure behavior" through friendship on Twitter and Google Plus.
- Location-sharing policies are yet another social network-based example of an allow/reject decision. Xie et al. [368] model "disclosure behavior" by whether the user is willing to share their location to three different types of audience (family, friend, and colleague) in a location-sharing dataset. Similarly, Wilson et al. [359] modeled (hypothetical) location requests to generate location sharing "wizards".

- Taken in aggregate, location sharing decisions can be expressed as “sharing policies”. Several researchers have analyzed such policies to provide recommendations for new/better privacy rules [26,76,283,327].

Finally, some platforms (especially social networks) may offer users a plethora of features to control various aspects of their privacy. Facebook, for example, has functionality to allow users to 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) [151,360]. Wisniewski et al. [361] demonstrate how users’ engagement with these features can be captured in six privacy management profiles, and Wilkinson et al. demonstrate how to adapt Facebook’s interface to these profiles [356].

In TLA, the disclosure of personal information (e.g. during the onboarding procedure, or the installation of a learning application) can be used as a measure for their disclosure tendency. Their decision to allow or reject apps from tracking their training data, their location, or other behaviors also contributes to this measure. Note that this measure will likely be multi-dimensional and context-dependent: if users allow e.g. a personal trainer app to track their location, this does not necessarily mean that they agree with location tracking by other learning apps. Moreover, the social networking aspects of TLA will likely have features to manage one’s relational, territorial, network, and interactional boundaries. Tracking users’ engagement with these features will further enrich UTPs user model.

If combined with other data, implicit feedback on recommendations can be an important user modeling input

If UTP makes explicit privacy recommendations to the user, there is an opportunity to gather feedback on the quality of these recommendations. Recommender systems often thrive on explicit or implicit feedback on the presented recommendations.

Implicit feedback is easy to gather—it simply involves tracking whether the user follows the privacy recommendation or not. Such implicit data is inherently noisy, though. For example, if a TLA user ignores a recommended privacy setting, this does not necessarily mean that they disagree with the recommendation (e.g., they may simply not be paying attention to it). Moreover, implicit feedback has to deal with the “pigeon-holing” problem: it can only gather positive feedback on privacy settings that are recommended, so it is hard to steer UTP in a different direction by giving implicit feedback on the recommendations alone. As such, implicit feedback on recommendations should be combined with other input mechanisms (e.g., explicit feedback or “manual” privacy-setting behaviors).

In terms of explicit feedback, recommender systems have traditionally used 5-star rating scales (e.g. Amazon) and thumbs up/thumbs down (e.g. YouTube, Netflix). A scale with fewer options

allows for quicker rating [307], but captures less information [163]. Users prefer 5-star ratings, followed by 10-star ratings and thumbs up/thumbs down [108]. Alternative feedback mechanisms include critiquing, in which the user does not simply judge the current recommendation but also suggests changes to the recommendation [58,227].

While explicit feedback may be a useful technique for traditional recommender systems, it seems less suitable for UTP in TLA. The main goal of TLA is to give *learning* recommendations; privacy recommendations are a secondary functionality. Users are not likely to be very involved in this secondary recommendation process, and thus less likely to be motivated to give explicit feedback on the privacy recommendations.

Examples of feedback on privacy recommendations include:

- Kelley et al. [159] developed a mechanism for what they call “user-controllable policy learning”. It involves the incremental manipulation of location-sharing policies, where both the user and the system refine a common policy model. This feedback mechanism showed promising increases in policy accuracy.
- Patil et al. [264] created a system where participants could indicate their level of comfort with location disclosures that were controlled by a privacy profile (although this feedback was not used to create a user model).

Heuristic influences make behavior less precise than attitudes

Privacy decision behaviors are expected to occur frequently in TLA-based systems and are thus an abundant and easy to observe input for UTP. However, behavior is usually less precise than attitudinal data [185], and many researchers find a “gap” in predictability from attitudes to behaviors [224]. This gap also exists in the realm of privacy decisions [6,22], where it has been dubbed the “privacy paradox” [247]. This gap is further explained in Specification Vol. 1 (Section 1.1.).

Part of this lack of precision may be due to the fact that behaviors are often influenced by external factors [185]. In the area of privacy, decisions may for example be influenced by the anticipated benefits of disclosure as much as the associated privacy concerns [165,358]. The tradeoff between privacy and these other factors can be modeled—something that is discussed in Section 4.4.

Another part of this lack of precision is likely due to the fact that users’ privacy decisions are often heuristic [7,154,155]. Various of heuristics that influence privacy decisions are discussed in Specification Vol. 1 (Section 1.1.), but there are two problems that particularly influence the precision of privacy behaviors in the context of UTP: persuasion and reactance.

Recommender systems have a persuasive power [72,110], and Knijnenburg and Jin [177] showed that this is also true for privacy recommendations. Consequently, online services can manipulate user's disclosure behavior, e.g., by providing privacy recommendations that systematically encourage over-disclose or under-disclose compared to the user's preferences. Even privacy recommenders that remain in the background by adaptively changing default settings have a persuasive effect by virtue of the default effect [168,200].

Conversely, if provided recommendations are too far from the user's actual preferences, users may disregard or even counter such recommendations [16,72,99]. The strength of this reactance effect depends on the perceived expertise of the source of the recommendation, as well as the valence of the recommendation (positive or negative). Reactance has not been studied in the specific area of privacy recommendations; this would be an interesting area of future research.

If privacy behaviors are not reflective of users' preferences, then how can we ascertain that our predictions are accurate? One step towards this goal is to investigate the psychological reasons behind heuristic decision practices [85] which can help create decision environments where they are less likely to occur. For example:

- Knijnenburg et al. [173] developed a new form auto-completion tool that eliminates default effects in self-disclosure.
- Wu et al. [367] demonstrate that the order of information disclosure requests influences the variability of users' disclosures which in turn influences the accuracy of a privacy-adaptive recommender system. Particularly, they find the best predictive accuracy when alternating sensitive and non-sensitive questions.
- Patil et al. [264] show that giving users immediate feedback on a system's location-sharing decisions increases users' discomfort with the decision, which leads to feelings of oversharing. Providing feedback immediately after system-enacted disclosures may thus create a heightened sense of disagreement with the decision. It is therefore advisable to delay disclosure feedback, allowing for a reasonable 'cooling off' period.

An alternative solution is to attempt to distinguish heuristic from "rational" decisions. This can be done by measuring indicators of accuracy (e.g. the time it took the user to make the decision) or covariates of accuracy (e.g. the Elaboration Likelihood Model suggests that motivation and self-efficacy are correlated with more elaborate decisions; see Specification Vol. 1, Section 1.2). UTP could then decide to consider only behavior that meets a certain threshold, or use 'propensity score matching' [280] to factor out the heuristic effect.

Attitudes/preferences take effort to elicit but may be required to create an accurate user model for "privacy novices"

Aside from behavior, UTP can use users' attitudes and/or preferences as input. Attitudes are typically more stable than behaviors [185], but require explicit elicitation. Like with explicit feedback on recommendations, explicit elicitation may be unrealistic for UTP in TLA.

Moreover, due to the privacy paradox [247], there may exist a difference between users' attitudes and their behaviors. Hence, if recommendations are based on users' attitudes, they may in some cases want to behave in opposition to the recommendations. In this case, the persuasive effect of privacy recommendations [177] may bring their behavior closer to their attitudes, which can be a good thing in the long run. However, the "gap" can also lead to reactance.

Another complication in using attitudes or preferences as input for UTP is the *principle of compatibility* [224]: attitudes are more predictive of behavior if they are compatible with respect to action, object, and context. Therefore, given that UTP adapts to who, what, to whom, and context (see Section 3), the measurement of attitudes and preferences should be contextualized using these same parameters as well. Having to elicit contextualized preferences may however make the elicitation of attitudes/preferences prohibitively cumbersome.

The following works have used attitudes and/or preferences as input for user modeling in the area of privacy:

- Foundational work by Westin et al. [118,119,351,352] clustered participants into three groups according to their answers to four privacy-related questions: privacy fundamentalists, pragmatic majority, and marginally concerned. This categorization has been extensively used in privacy research, but recent work has shown that it has important shortcomings [365].
- Ackerman et al. [3] and Spiekermann et al. [309] use the outcomes of a privacy concern questionnaire to cluster participants into different groups. Ackerman et al. find three clusters aligned with Westin's work, while in Spiekermann et al. the pragmatic majority cluster is split into "profile averse" and "identity concerned" individuals. Interestingly, Knijnenburg et al. [172] find a similar distinction in behavioral data.
- Bahirat et al. [24] find that attitude-based clustering can be used to create profiles for a privacy-setting interface for public IoT devices.
- In a study on location-sharing, Knijnenburg and Jin [177] asked users to evaluate the activity performed at the location, and used this as an input to decide which sharing options to recommend.
- In their study on privacy profiles, Knijnenburg et al. [172] show that for their Android app recommendation dataset profile membership can be predicted by users' general privacy concerns and collection concerns, and for their Facebook dataset membership can be predicted by trust in Facebook and need for consent.
- Li et al. [212] find that users' perceptions of social privacy threats and the importance of notice and control rank among the top features for predicting disclosure behavior.

Attitude/preference-based input has been extensively studied in the context of multi-attribute utility theory (MAUT) [32], which has been used as a mechanism for recommender systems [121]. Usability studies of preference-elicitation methods suggest that domain novices tend to be bad at turning their preferences into system-level attributes [181,183,184]. In the realm of privacy recommendations, this would mean that users who are not privacy experts are bad at

turning their preferences into system-based privacy settings (a conjecture which the ample research on the heuristic nature of privacy decisions seems to support [7,154,155]). Novices are much better at expressing their preferences at a higher conceptual level [186], which would mean that they could potentially be asked to judge the overall balance between the benefits of disclosure and their privacy (which UTP would then have to turn into specific settings).

We expect that in the context of TLA, some users will be privacy experts (e.g. officers trained in cyber defense/offense), while many others will be novices. It is thus advised to provide TLA users with a combination of input mechanisms where privacy preference-based input is considered. Future research should test the viability of this approach.

User traits can kick start UTP's user model and are essential in a TLA environment with a wide variety of user types

Users' privacy preferences could also be derived from their traits, such as their culture, demographics, job title, or personality. Such data is envisioned to be readily available in TLA, which means that it can be used to instantly create (a first approximation of) a user privacy model, thereby overcoming the cold-start problem (see Section 4.4).

Given the envisioned wide variety of TLA users, trait data can be an important user-modeling aspect, because it makes the UTP user models user-dependent. For example, it allows the model of a Colonel to be different from the model of an Admiral. The following works have used user traits as input for user modeling in the area of privacy:

- Dong et al. show that the user's "follow tendency" (the ratio of following to followers) is an important aspect in predicting Twitter and Google+ users' following behavior.
- In their study on privacy profiles, Knijnenburg et al. [172] show that for their Android app recommendation dataset profile membership can be predicted by users' mobile internet usage; for their Facebook dataset membership can be predicted by users' age and gender; and for their online retailer dataset membership can be predicted by users' age.
- Knijnenburg and Kobsa [170] demonstrate that users' preferred justification message depends on their gender and overall disclosure tendency.
- In a study on cross-cultural privacy predictions conducted in 8 different countries, Li et al. [212] show that privacy predictions can be significantly improved by taking country-level cultural variables, especially country-level measures of individualism.

Future research should investigate relevant user trait inputs to UTP in the context of the TLA.

*Context is an essential user-modeling input for UTP
in pervasive TLA-based systems*

As mentioned in Section 3.4, TLA users' privacy decisions are likely to be heavily context-dependent. Many contextual variables (the current time, device, location, etc.) are envisioned to be readily available in TLA, which makes them suitable to use as input for UTP user models. It is important to automatically include context in a user model, as the absence of contextual information would likely increase the number of times users have to interact with their privacy settings [257].

Given the envisioned pervasiveness of TLA-based applications—on the job, just in time, multi-device [101]—context data can be an important user-modeling aspect for UTP user models in TLA. For example, Specification Vol. 1 (Section 3.2) discussed important differences in the use of personalized learning notifications on different devices: notifications should be reduced outside work hours on personal devices like smartphones and smartwatches, while they should be scrubbed of sensitive information on (semi-)public displays like smart TVs. In this example, the context variables “device modality” and “time” are considered to be crucial for an effective, privacy-aware implementation of notifications in TLA. The organizational context (division, current mission) can also be crucial in terms of rules that may be in place to keep certain training data confidential.

The following works have used context as input for user modeling in the area of privacy:

- Bahirat et al. [24] demonstrate that users' decision to allow/reject public IoT tracking scenarios is crucially dependent on the recipient, the data type, the purpose of collection, and the persistence of the collection. These contextual variables alone can predict users' disclosure decisions with 73.1% accuracy.
- Dong et al. [87] measured the trustworthiness of a new follower by the ratio of followers to following, and their appropriateness by the overlap in followers and following with the user. They show that both of these contextual variables are instrumental in predicting whether the user will follow the new follower back.
- Dong et al. [87] also show that the user's tendency to share their location depends on the trustworthiness of the recipient (i.e., their average tendency to share with that recipient) and the sensitivity of the location (i.e. their average tendency to share that particular type of location).
- Benisch et al. [26] show that location and time are important predictors of location sharing. There are important differences between weekdays and weekends as well.
- Xie et al. [368] show that location, time, companion, and emotion have a high impact on users' decision to share their location. There are some appropriateness patterns here: users are more likely to share their location with family when they are visiting somewhere with other family, with friends when they are with other friends, and with colleagues when they are with other colleagues. TLA could use these patterns to predict users' context-relevant privacy settings.

Contextual variables can introduce a large amount of sparsity to a user model, or result in overfitting [368]. Luckily, various mechanisms exist that allow a recommender system to integrate contextual variables without significantly reducing the amount of data available for each prediction context [14,368]. Another way to prevent overfitting and sparsity is to have a psychological theory behind the measurement of certain contextual variables [87]. An example of this is Toch et al.’s [326] realization that entropy is the most important aspect of location sensitivity, and Li et al.’s [212] finding that country-level cultural variables are a better predictor than country itself. Future research should investigate the relevance and structure of contextual inputs to UTP in the context of the TLA.

Recommendation: Unobtrusively gather user-modeling input from as many different sources as possible but make sure to respect the user’s privacy

This subsection has introduced many examples of successful privacy decision models. A TLA-based implementation of UTP should use a combination of privacy decision behavior, attitudes, and feedback on UTP recommendations as input for its user models. Moreover, UTP should make use of available user traits and context as additional predictors for these models.

This input data should be gathered unobtrusively, though; beyond the recommendations themselves (see Section 5), users should be bothered as little as possible with questions pertaining to their privacy, because privacy is not the main task and goal of the TLA user.

It is important to note that gathering user modeling input for UTP can in itself turn into a privacy problem, especially if the user declines to have all TLA-based applications access a certain data. In that case, UTP should probably also decide not to track that variable, even if this means the UTP user model will be less precise. For example, if a TLA user is based at an undisclosed location, UTP will likely switch off location-sharing or tracking for all TLA-based applications. In that case, UTP itself should also switch off its own location-tracking practices, until the user turns them back on (or until UTP can derive with some certainty from other data that the user is no longer at an undisclosed location).

Another way to reduce privacy-related data storage is to only store recent data. Given that users’ privacy preferences are dynamic, this practice actually ensures that UTP’s user privacy models reflect potential changes in user preferences that happen over time.

4.2 Algorithms

This section covers the task of algorithmically transforming the input data (Section 4.1) into privacy recommendations. From a technical perspective, the algorithms that can calculate privacy recommendations are no different from the standard algorithms discussed in the recommender systems literature. Therefore, this section does not cover the existing research on specific algorithmic approaches and efficient implementations; for these topics TLA back-end developers can consult existing reference works (e.g. [139,276]).

In this section, we primarily cover existing applications of algorithms to predict privacy preferences. Among existing approaches, a main distinction can be made between methods that rely on other users' behaviors (so-called "collaborative filtering" methods), versus methods that rely on the target user's behaviors only (so-called "case-based reasoning" methods).

Collaborative filtering leverages other users' privacy settings to provide recommendations but can "leak" sensitive information about settings

Collaborative filtering leverages other users' privacy behaviors to help predict the current user's privacy preferences. An example is the nearest neighbor approach, where the target user's behaviors are matched with other users' behaviors in an attempt to find the users who are most similar to the target user. Once a set of nearest neighbor users has been found, any unknown preferences of the target user can be predicted using the preferences of these neighbors. This approach is called "user-based collaborative filtering", as opposed to "item-based collaborative filtering", which applies the same approach to the items instead of the users.

In a networked setting, this approach can be augmented to what are called "trust-enhancing" algorithms, by leveraging the network ties between users [333]. In this case, the algorithm recommends settings based on the target user's friends, or people they know or are close to. Such algorithms may at time be less accurate due to the decreased sample size (only "neighbors" with ties are considered); however, the trust aspect of the system would be enhanced, as privacy recommendations could be presented with an effective justification (that is potentially adaptive; see Section 5.2), i.e., telling the user the recommendation was based on the behavior of a known friend [176].

More advanced versions of the collaborative filtering approach use dimensional reduction to overcome the issue of sparsity, and to increase the efficiency and accuracy of the predictions. Matrix factorization is currently the state-of-the-art approach, with deep learning algorithms still lacking behind, but gaining quickly in popularity and efficiency [195,338].

Existing works in User-Tailored Privacy that use collaborative filtering include the following:

- Ismail et al. [136] incorporate user-based collaborative filtering into the mobile app privacy domain and use it for predicting the suitability of various permission sets. A similar approach is taken by Liu et al. [216].
- Toch et al. [327] propose a method for suggesting configurable privacy policy defaults for location-sharing by analyzing existing policies using user-collaborative policy analysis to cluster similar users. They also let users personalize their policy profile using a User Interaction Model. A similar approach is taken by Sadeh et al. [283].
- Xie et al. [368] develop PPRec, which uses a hybrid of user-based and item-based algorithms to provide context-aware privacy recommendations for check-in-based location-sharing services.

In section 4.1, we noted that gathering user modeling input for UTP can in itself turn into a privacy problem. Similarly, privacy recommendations that are based on collaborative filtering can “leak” information about users’ privacy preferences, and thus create security violations. For example, if a hacker has access to TLA users’ privacy settings, they may be able to derive from these settings what kind of information users find most sensitive. Zhao et al. [374] propose a system that treats users’ privacy recommendations themselves as sensitive information, and in response, they build a differentially private privacy recommender using standard data obfuscation techniques.

Case-based reasoning can be used to generate smart default settings or profiles, but are usually “static” in their predictions

An alternative approach to privacy recommendation is “case-based reasoning”, which applies contextualized rules to decide on the best outcome in a given situation [303]. The rules could be based on common sense (e.g. recommendations could be pre-defined for various types of applications) or based on (past) data of other users (e.g. past user data could be used to establish a “privacy score” for each type of app, which then informs future recommendations). Regardless, one benefit of case-based recommendation is that the system does not require “live” user data (which reduces the chance of “privacy leaking”) and can easily be implemented on the client-side (which voids the need for user data to be shared with the recommender) [38].

While most case-based approaches to privacy recommendation are context-dependent, not all of them are *personalized*, i.e., many of them will produce the same recommendation in the same situation, regardless of the user. Personalized versions are usually profile-based, where the specific set of rules to apply to the target user’s behavior depends on the profile that was assigned to them. Profiles turn the user modeling from a multidimensional tracking problem into a simpler classification problem [172], and offer personalization without requiring a central server to calculate the recommendations.

Existing works in User-Tailored Privacy that use a version of case-based reasoning include the following:

- Pallapa et al. [257] proposed context-aware approaches to privacy preservation in wireless and mobile pervasive environments. One of their solutions leverages the history of interaction between users to determine the level of privacy required in new situations.
- Watson et al. [350] train a classifier to create a rule-based recommender for Facebook privacy settings. They personalize the classifier through segmentation, but this approach shows no improvement over the non-personalized classifier.
- Bahirat et al. [24] employ tree learning to create a default policy for an IoT privacy-setting interface. Their non-personalized policy attains 73.1% accuracy. Clustering users to create three privacy profiles improves the accuracy to 81.5%.
- Li et al. [212] create a classifier for privacy behaviors based on demographic, contextual, attitudinal, and cultural input data. They demonstrate that these factors all have an

impact on prediction accuracy (which reaches 76.8%), but accuracy is good with just a single type of input data as well (74.5–76.2%), indicating that such approaches can work even with limited data. Tree-learning and logistic regression classifiers reach the highest accuracy on their dataset; nearest neighbor algorithms perform decidedly worse.

- Ravichandran et al. [271] employ using k-means clustering on users’ contextualized location sharing decisions to come up with default privacy policies. Liu et al. [216] use a similar approach to generate Android app permission policies.

A downside of case-based privacy recommendations is that its rules are static: unless the algorithm behind the rules is re-trained, the rules will not change even if users’ behavior evolves. The same holds for the profiles in personalized versions of case-based privacy recommenders: users are typically not re-classified into a different profile, and if profile assignment is done by the user, the Control Paradox (see Specification Vol. 1, Section 6.2) predicts that they are not likely to do so themselves either. Case-based recommendation approaches should thus be combined with a manual settings feature to allow for changes over time [24].

Recommendation: *when it comes to privacy recommendations, explainable and user-controllable algorithms are better than high prediction accuracy*

While we advise TLA back-end developers tasked with implementing UTP to leverage advances in recommender systems, we warn them against going overboard on the implementation of an algorithm. Research shows that simple algorithms are preferred for privacy recommendation, because they are more easily explainable and user-controllable [76,159,339,348].

User control can be provided through preference input or explicit feedback (see Section 4.1), but more complex, conversational approaches have been proposed in the realm of recommending privacy rules (cf. [76]). Explanations have been studied extensively in the recommender systems field [102,107,325], but no work to date has focused on explaining privacy recommendations. This is unfortunate, because research has found that users like explanations [123], and that they increase users’ understanding of the recommendation process [107,337], their trust in the quality of the recommendations, and their perception of competence and benevolence of the system [73,97,341]. Future research on explanations of privacy recommendations could not only inform the implementation of UTP in TLA, but also advance the state-of-the-art in privacy recommendation research.

4.3 Target

Once the user’s privacy behavior or attitude is known, the question remains how UTP should adapt to this behavior/attitude. Traditional recommender systems primarily attempt to find items that are the closest match to the user’s preferences, but more recently researchers have begun to question the validity of this approach. For one thing, research shows that user preferences are fleeting, constructed on the fly, and vulnerable to distorting influences, rather

than well-defined, fixed, and invariant [138] (see also Section 1.1 of Specification Vol. 1). Consequently, researchers have advised to increase the diversity of recommendations [36,357] and more generally, to provide recommendations with the goal of encouraging exploration, self-actualization [182], and the fulfillment of longer-term goals [92].

A similar dilemma has to be resolved regarding User-Tailored Privacy for TLA. On the one hand, the system could attempt to match users' existing behaviors and attitudes, thereby alleviating some of the hassle of privacy management. On the other hand, the system could facilitate users' behaviors by automating or supporting synergistic auxiliary behaviors, or even attempt to go beyond users' current behavioral patterns and help users explore new ways to manage their privacy [180].

Alleviate the burden of privacy management by matching TLA users' current behavior/attitudes

UTP can recommend and/or automate actions that match one-to-one the user's current behavior. This is a safe adaptation practice that relieves some of the burden of making privacy decisions. For example, if a TLA user has hidden their recent training outcomes for a particular class from their colleagues, the system can automatically also hide these outcomes from a new colleague.

Recent research has shown, however, that suggesting or highlighting users' existing practices can be regarded as redundant or a nuisance [240]. Fully automating such practices may be a better solution, as it significantly reduces the burden of having to engage in the task oneself. However, the same research also finds that users sometimes consider their own privacy practices superior to proposed algorithmic solutions, especially when it comes to privacy behaviors they engage in frequently. In such situations, they may prefer to manually engage the privacy management behavior, rather than having it automated.

Solidify users' privacy management practices by recommending privacy behaviors that dovetail with TLA users' current behavior/attitude patterns

UTP can also recommend and/or automate actions/settings that go beyond users' current privacy practices but are selected to support the existing practices. Such "privacy tips" (see Section 6.2) may solidify the user's current privacy management practices. For example, if a TLA user has hidden recent training outcomes from their colleagues, the system can recommend to also hide these outcomes from a social learning application that could inadvertently "leak" some of these outcomes to the user's colleagues.

This approach to user-tailored privacy requires some knowledge about activities that may be considered synergistic. Work on the dimensionality of disclosure behavior [172], privacy profiles [361], and recipient groups [178] provide a solution here: if a user frequently engages in privacy

behaviors that belong to a certain dimension or profile (or that is geared towards recipients of a certain group), the system could recommend other behaviors that belong to the same dimension or profile (or recommend applying the behavior to other recipients in the same group).

Recent research has shown that recommendations regarding behaviors the user does not frequently engage in are better made in manner that is not fully automated, e.g. by highlighting or suggesting the recommended behavior to the user. Suggestions have the added benefit of giving the system the opportunity to explain the recommendation. This would allow the system to explain how the recommended behavior dovetails with users' existing privacy management practices [240].

Move beyond current behavior/attitude patterns

Finally, UTP can recommend and/or automate actions/settings that move beyond the practices that the user is currently familiar with. This would effectively teach the user new privacy behaviors. For example, if a TLA user has hidden all training outcomes from their colleagues and has never shared anything selectively, the system could recommend sharing certain training outcomes selectively with certain colleagues only. This would teach the user a new practice (namely "selective sharing").

Research by Wisniewski et al. shows that on social networks, users' privacy profile is significantly related to their level of privacy knowledge [361]. Novices tend to employ a veritable grab bag of privacy management behaviors or tend to shy away from contributing any information to the network (and instead focus on consumption). Experts, on the other hand, tend to be privacy maximizers or selective sharers. They argue that users' privacy profile can be a function of their knowledge as much as their preference and propose that education can help elevate novice users' privacy practices.

A similar approach can be taken for TLA: especially when using explicit suggestions for recommendations, the system has an opportunity to teach users about privacy practices that are currently not known to them (see Section 6.2). Recent research shows that users will likely appreciate such a personalized educational approach [240].

Recommendation: Combine different recommendation targets in a usable manner

Assuming that the TLA environment will have a diverse offering of privacy functionalities for end users to manage their privacy (something we argued for in Specification Vol. 1), those end users will likely end up with a wide variety of privacy management practices, using certain privacy features very regularly, and other much less frequently. One can also assume that most TLA users will not know all of the available privacy features, but instead familiarize themselves with a subset of features that seems to best suit their personal privacy preferences.

As such, certain frequent privacy behaviors that are part of the the user’s main privacy management strategy can be fully automated, so as to completely alleviate the burden associated with these behaviors. Namara et al. suggest that such automation is most appropriate for privacy behaviors that have low-impact consequences, lest users are worried that the system may cause problems if it incorrectly predicts their decisions [240].

UTP can also recommend auxiliary behaviors, but rather than automating these behaviors, TLA users should be “nudged” towards these behaviors using subtle design interventions or suggestions. Namara et al. show that this approach helps users to solidify their privacy management practices without taking too much control [240].

Finally, UTP has the opportunity to teach the TLA user about privacy features (see Section 6.2). Tailored education focuses these efforts on privacy practices the user is currently unaware of, which ascertains that users are not overwhelmed by lengthy privacy tutorials [361].

4.4 Problems

The practice of modeling users’ privacy decisions is not without potential problems. We have already discussed the potential negative influence of heuristic behavior and the privacy paradox on the accuracy of user privacy models [247], as well as the problematic “positive feedback loop” that may be caused by the persuasive effect of privacy recommendations [177]. We also mentioned the fact that user models should respect the dynamic nature of users’ privacy preferences which may change with context as well as over time.

Here, we will discuss the balance between over- versus under-disclosure, the potential tradeoff between privacy and other user goals (e.g. the benefits of personalized learning) and/or system goals (e.g. mission-tactical goals), and the impact of traditional machine learning problems like overfitting and the cold start problem.

Study the relative cost of over-disclosure versus under-disclosure, and build this cost into the UTP algorithms

Prediction algorithms usually try to predict users’ behaviors as accurately as possible and will penalize a solution with many “false positives” and “false negatives”. Typical metrics of prediction accuracy such as F1 and AUC treat these two types of errors as equally bad, but this may not be desirable. Indeed, in the privacy domain false positives and false negatives translate into over-disclosures (recommend disclosure when the user in fact does not want to disclose) and under-disclosures (recommend non-disclosure when the user in fact wants to disclose), and one could argue whether these two types of mistakes are equally problematic. “Cost-sensitive” learning and meta-learning algorithms exist that allow one to specify the relative importance of each type of mistake [93], but the relative weight of the mistakes (known as the cost function) is a decision the developer will need to make based on their best judgment.

Most privacy researchers would argue that over-disclosure is worse than under-disclosure [262,346], but how much worse is unclear. Indeed, research shows that under-disclosure occurs quite frequently and may thus be more problematic in aggregate [264]. Moreover, the cost of over-disclosure may be more prominent for TLA environments that deal with classified information and qualifications, while under-disclosure may be more problematic in TLA environments that promote social learning.

There is not much research on this topic in the realm of privacy, with the exception of Benisch et al. [26], who do not attempt to determine the cost function, but instead investigate the impact of changing the relative cost of over- versus under-disclosure from 1 to 100 on the output of their privacy policy recommender. At the very least, TLA-based systems should take a similar approach, and investigate the impact of such cost functions. A better solution would be to actively study the relative impact of over- versus under-disclosure in various contexts, and subsequently build the resulting cost function into the deployed UTP algorithms.

Trade-off privacy with other user goals, as well as the goals of the system, the institution, and other users

Another complication is the fact that privacy decisions are rarely made in isolation, but usually considered as a trade-off with other goals. Indeed, research shows that users are typically willing to give up some privacy in return for personalization benefits [219,319]. To more carefully model whether the user will disclose a piece of information, UTP should model this underlying trade-off between the costs of disclosure and the benefits for personalization [166]. This trade-off is the foundation behind the privacy calculus [205,206] (see Section 1.2). Therefore, the incorporation of other goals into the privacy prediction algorithm essentially turns the privacy calculus into a prescriptive model of disclosure behavior [180] (see Section 1.3).

This approach has been attempted before [26,124,166], most explicitly in a recent study on a demographics-based recommender system [166]. In this study, the order in which demographic questions are presented to the user is determined on the fly, based on a trade-off between the sensitivity of the information and the potential benefit of disclosing it to the recommender. The former was determined in a pre-study, while the latter is calculated on the fly, based on a prediction of how much the underlying recommendation user model changes if the information is disclosed. The trade-off is based on a sensitivity threshold: below this threshold, items are ordered by decreasing benefit, while above the threshold they are ordered by increasing sensitivity. In some tested versions of the system, the threshold itself is determined on the fly, using an estimation of users' disclosure tendency based on their previous disclosure behaviors.

Trade-off based user privacy models bring along a number of additional questions regarding their design and implementation. For example, one would have to decide whether the trade-off between privacy and benefits should be a compensatory (linear) trade-off or a non-compensatory (threshold-based) trade-off. Moreover, it requires estimates of both risk and benefit, which can be either derived from objective system parameters, or subjective user

experiences [180]. Solove’s taxonomy of risks and benefits also suggests that both risks and benefits come in multiple forms, and that different users may have different perceptions and preferences regarding these different types of risks and benefits [304].

Taking this trade-off-based user privacy model a step further, one could incorporate the goals of other TLA users (e.g. in case the sharing decision involves team performance), the institution (e.g. in case the user is institutionally prohibited from sharing certain information), or the learning system itself (e.g. in case the system would derive particular benefits from sharing certain information). Note that these external considerations could result in a situation where the UTP recommendation deviates substantially from the user’s own sharing preferences. To avoid ethical dilemmas or widespread reactance to the privacy recommendations in such situations, UTP should carefully explain the considerations behind the recommendation, so that the user can be informed about the decision (see Sections 4.2 and 5.2). From an ethical perspective this is especially important when recommenders take factors other than the user’s own preferences into account.

Prevent overfitting by couching the recommendation logic in psychological principles

In calculating users’ contextualized privacy preferences, one has to deal with the problem of overfitting: the more granular the contextual preferences get, the less data the predictions will be based on. Finding a solution that can provide robust predictions at a reasonable level of granularity is an important aspect of UTP. For example, Xie et al. present a context-aware privacy prediction algorithm that combines input at various levels of granularity to produce recommendations that are both robust and highly context-specific [368]. Their solution uses a distance function to incorporate information from previous decisions that were contextually similar but not identical.

Beyond granularity, the availability of large amounts of contextual information can also present an overfitting threat: The more information is incorporated in the model, the more accurate its predictions are likely to become. Extensive cross-validation and the judicious use of feature selection mechanisms can prevent overfitting, but these efforts crucially rely on deep knowledge of the problem domain. In that regard, Dong et al. argue that user privacy models should be couching in extensive psychological principles to ensure their robustness [87]. Their work ensures that the predictions of their privacy prediction model reflect underlying psychological principles, which they argue helps to reduce overfitting.

Prevent the cold start problem by asking new users to select a default profile or to input a specified set of preferences

As mentioned in Section 4.1, privacy prediction needs a certain amount of input before it can determine the user's privacy preferences. Without such input, it is impossible to create an accurate user privacy model. In the user modeling community, this problem is known as the "cold start" problem [288], which occurs both at the item side (i.e., not enough data for this item to recommend this item to any user) and the user side (i.e., not enough data for this user to make any recommendations to them). Especially on the user side, this problem is hard to solve.

In order to overcome the cold start problem, TLA-based systems can probe users' privacy preferences by asking them a small set of general questions, for example based on their comfort with respect to sharing data with other certain apps, TLA users, managers etc. This can then be used to make a first-order approximation of their sharing tendencies. The following works take this type of approach:

- Liu et al. [215] developed a profile-based personalized privacy assistant that elicits a small set of preferences pertaining to whether or not the user felt comfortable granting some permissions to apps from certain categories. Based on these answers, assistant was able to identify a privacy profile that matches the user's preferences, and to recommend a number of permission settings changes to the user based on this profile.
- Similarly, Lin et al. [214] generate privacy profiles for app privacy settings, taking into consideration purpose information and users' self-reported willingness to potentially grant access, elicited in a scenario-based online study.
- Wilson et al. [359] describe a small set of general questions to help users specify policies for sharing their location with their social connections in Locaccino, an application that allows users to share their current location with their friends subject to user-controllable privacy rules.
- Fang and LeFevre [96] evaluate a Facebook privacy wizard that categorizes friends based on an small initial set of labeled friends. They demonstrate that the wizard can create such categories with high accuracy (90%) after users manually label only 25 friends.
- Kolter and Pernul [193] developed a user preference generator for 12 Internet service types based on the input to a configuration wizard. After the user uses the wizard, the application switches to a privacy cockpit that provides clear overview of the configured preferences.

One problem with this approach is that users may initially lack the experience required to express their privacy preferences. The initial preference elicitation phase may thus be influenced by the language and ambiguous expressions during this process. This means that it is of vital importance in this phase to assist users with suitable user-interfaces with clear information visualizations that help them comprehend all the privacy-related options.

Another way to help users in the initial elicitation process that overcomes the cold start problem is to simplify this process through “privacy profiles” that are generated based on previously collected data. For example, Bahirat et al. [24] leverage a dataset from Lee and Kobsa [209] to generate a set of privacy profiles that span the preferences of a wide variety of users. Each profile contains a myriad of default settings for various public IoT scenarios. By selecting among the available profiles (their optimal solution has three of them), users select a personal initial default setting that requires little adjustment. While this approach does not provide dynamically updating recommendations, it provides a certain amount of personalization without requiring any knowledge about the user.

Recommendation: Implement UTP using a layered and gracefully degrading approach

Not all privacy-related situations in TLA require a user-tailored approach. It is important to remember that the majority of privacy problems should be prevented by design (see Specification Vol. 1) rather than solved by UTP. Moreover, implementing UTP will be a complex endeavor that requires a solid infrastructure for privacy-related data collection, user privacy modeling, and recommendation. As such we recommend that UTP should be implemented incrementally, starting with simple “smart profile”-based approaches, and moving to more complex user privacy modeling solutions later in the development process.

Combining simple and complex user modeling approaches within the same system allows for “graceful degradation” of the user modeling approach. For example, a collaborative filtering recommender will not work when too little user data is available, or when the user is offline. In such cases, the system can fall back on a profile-based approach, or even just recommend the settings of the average user.

Finally, to ascertain the quality and validity of the UTP module, it should be evaluated using the “layered evaluation” approach [259], which allows input, processing, and output procedures to be evaluated separately.

5 Adapting privacy

While most existing work on user-tailored privacy covers the “modeling privacy” aspect (Section 4), it is of utmost importance to also cover various ways in which TLA-based systems can leverage these user privacy models to provide user-tailored privacy decision support. This section therefore discusses how UTP can implement adaptations throughout the TLA-based system that help users make better privacy decisions. Particularly we discuss:

- **Adaptations involving the privacy **settings** or information requests themselves**—These adaptations alleviate the burden of privacy decision-making, either through fully automated adaptive defaults or adaptive “nudges” in the form of highlights or suggestions. For example, UTP may automatically decide to share a TLA user’s learning outcome with some but not all of their teammates.
- **Adaptations involving the **justification** for certain settings or information requests**—In their simplest form, adaptive justifications can inform users about the reasons behind a recommendation, or act as a “nudge” that gives users a rationale for engaging in a privacy-related behavior. More complex forms of justifications can educate users about the risks and benefits involved in a privacy decision. For example, UTP may inform a TLA user that sharing their training data with a TLA processor may increase the effectiveness of upcoming training modules. It can offer the user to learn more about how TLA processors use detailed training data to personalize training recommendations.
- **Interface adaptations**—These adaptations restructure the user interface of the system to make certain privacy actions easier to accomplish. For example, UTP may notice that a certain TLA user often shares learning outcomes selectively with a subset of possible recipients, and subsequently make this action easier to accomplish by giving it a more prominent place in the “training data-sharing” section of the TLA user interface.
- **Privacy-adaptive **personalization****—These adaptations influence the types of personalization a system can engage in based on the collected user data, thereby preventing potential unwanted inferences to be made. For example, if a TLA user is diagnosed with PTSD due to a past mission, UTP can instruct TLA processors to refrain from using information about that mission as a basis for its personalized training recommendation.

This section will conclude with a discussion of the tradeoff between fully automated privacy adaptations (which are more convenient, but less engaging, and can be dangerous when they are wrong) and keeping the user in the loop on recommendations (which allows users to become more informed about privacy, but can be overwhelming, and could potentially have a negative effect if the privacy adaptations are overly persuasive or authoritative).

5.1 Adapt the setting

The most commonly studied application of UTP is “adaptive privacy settings”. This application extends work on default settings as nudges [144,200] to the idea of adaptive defaults [302],

which take into account the fact that the optimal default setting in privacy crucially depends on the user and the context of the decision. Their adaptive nature prevents these nudges from threatening the users' autonomy (see Section 1.4).

Adaptive defaults can be applied to any privacy settings that can be represented as a checkbox, radio button, or multiple-choice question. This could include app privacy settings (e.g. "May this financial health awareness app have access to your bank account?") [169,190] and social sharing settings (e.g. a "sharing matrix" indicating which learning outcomes will be shared with which colleague) [168,171]. This section covers various ways in which adaptive default can be implemented in such situations.

Automatically apply settings to alleviate TLA users from frequent privacy behaviors, but avoid automating decisions with far-reaching consequences

Arguably the most straightforward way to effect adaptive privacy settings is to automatically apply them by default. If privacy settings are largely in line with users' preferences, it will be easier for users to choose the right settings [85]. In fact, in the ideal case where the system makes no errors in deciding on the default setting, the user does not have to take any action at all. As such, adaptive defaults are the least intrusive way to implement adaptive privacy settings, which makes them a very effective method for relieving users' decision burden.

Defaults also provide an implicit normative cue, e.g., a default value communicates what the system thinks the TLA user should do [41,228,292], and they create an 'endowment effect' where people are less willing to pay for what they perceive to be a gain in privacy than what they would demand if the same decision were framed as a loss [9,330]. Because of this, automatically applying settings by default will likely nudge users in the direction of that default [324].

While a large number of previous works on privacy prediction assume some sort of automatic application of the predicted privacy policy, very few works test whether users appreciate—and go along with—such adaptive defaults. Studying the potential for user-tailored privacy on Facebook, Namara et al. find that users appreciate adaptive defaults for features that they use frequently, as this may relieve them of a considerable amount of burden in these cases [240]. Users are skeptical about the accuracy of privacy recommendations, though, and therefore dislike the idea of automating privacy features that may have far-reaching social consequences, such as blocking unwanted friends or applications.

Moreover, Knijnenburg and Kobsa compare a static "smart default" setting (non-personalized but based on the average disclosure of past users) against a "shared-by-default" setting and a "private-by-default" setting in a mockup of a social network system [168]. They show that users are much more likely to follow the smart default than either of the other two defaults, and that users with high privacy concerns share more information in this situation, without increasing their perceived over-disclosure threat. These findings suggest that an adaptive default setting

can strike a balance between threat and disclosure that is absent from the typical “shared-by-default” and “private-by-default” options.

Highlight suggested settings to reduce TLA users’ cognitive burden in a subtle but useful manner

Another way to adapt settings is to highlight the setting that UTP identifies as most suitable. This method still requires some user involvement, which makes it a bit more burdensome, but also allows users to more readily correct any possible mistakes. Moreover, if UTP is uncertain about the best setting, it can highlight more than one.

Highlighted suggestions do not reduce the physical burden of decision-making, but they do reduce the cognitive burden, because in most cases users can go for the highlighted setting without much thinking. The cognitive burden can be further reduced by deemphasizing or hiding the settings that are not recommended, rather than highlighting the ones that are. The idea of adaptiveness through emphasis and deemphasis overlaps with the idea of interface adaptations (see Section 5.3).

Again, very few works test whether users appreciate and go along with “adaptive highlights”. Namara et al. find that users appreciate this adaptation method, but not for adaptations that are deemed “obvious” [240]. Participants in their study remark that they like the awareness-increasing ability of adaptive highlights.

Moreover, Knijnenburg and Jin test hiding and highlighting in a check-in based location sharing system [177]. They find that users are likely to follow highlighted recommendations, and even more so with hiding. In fact, they demonstrate that highlighting and hiding have a persuasive effect: users are more likely to follow the recommendations than what would be expected based on the prediction accuracy of the underlying algorithm. Finally, they find that users perceive a significant increase in privacy support when the system displays only a short list of recommended options (and hides the remaining options).

Suggest privacy settings/rules to keep TLA users involved in their privacy decisions, but avoid making suggestions with awkward social consequences

Another mechanism for effecting adaptive settings is to “suggest” settings to the user. Such suggestions can for instance come as textual advice from a “privacy agent” (e.g. Facebook’s privacy dinosaur [251]). Suggestions require more user involvement but are less risky than the other adaptation mechanisms. Moreover, their textual nature allows for a great amount of flexibility, in that their persuasiveness can be fine-tuned, and they can include both a suggestion and a justification (see Section 5.2).

Namara et al. test adaptive suggestions in a Facebook context using the privacy dinosaur as a vehicle for displaying suggestions [240]. They find that users appreciate suggestions for infrequent privacy behaviors and argue that suggestions can leverage the opportunity to teach users about privacy. At the same time, though, they find that the explicit nature of suggestions makes them less suitable for recommending privacy behaviors that may be socially awkward, such as suggesting the user to delete a friend's post.

Suggestions can be aggregated into larger chunks of advice. This includes “rules” for sharing or “profiles”, which are essentially “generators” of individual settings. A few papers test how users react to such privacy rule suggestions:

- Liu et al. demonstrate the effectiveness and usability of their mobile privacy policy recommendations [215]. Specifically, they show that users are likely to accept their suggestions, that the suggestions help them converge more quickly on a final privacy setting, and that users felt comfortable with the recommendations and deemed them helpful.
- Wilson et al. show that privacy profiles significantly increase sharing without reducing users' comfort.

Future work should study whether human-like characters improve or hamper users' acceptance of privacy suggestions.

Manipulate the order of sequentially presented settings and information requests to prioritize the disclosure of certain types of information

When settings or information requests are presented sequentially, e.g. in a conversational setting, UTP can also change the order in which settings or information requests are presented to the user. The order of sensitive versus less sensitive requests has an impact on disclosure. For example, Acquisti et al. find that disclosure rates are lower when asking less sensitive questions first [9], and Knijnenburg and Kobsa find that sharing rates in social networks are higher when users are asked to share with weaker ties first [168]. Adaptive request orders can thus regulate the permissiveness of users' privacy settings or the amount of information they disclose.

Adaptive request orders can also regulate what specific items users disclose: previous work shows that users are more likely to answer requests for personal information that are presented first, at the expense of requests made later in the process [169]. This is useful when the amount of information users provide is arbitrary, e.g., a TLA processor that uses any information it is given to provide learning recommendations. In these situations, the system can prioritize requests that are useful to the system and that users are predicted to be more likely to answer [166].

Knijnenburg conducted a comprehensive evaluation of adaptive request orders in a demographics-based health recommender system [166]. In this system, users answer questions

that vary in sensitivity as well as in the extent to which they are useful to the recommender. The UTP algorithm in this system makes a tradeoff by prioritizing questions that are useful to the recommender but that don't surpass a certain sensitivity threshold. This threshold can either be static or adapted to users' privacy concerns on the fly (as indicated by their answering behavior).

Trading usefulness and sensitivity did indeed improve the users' experience. In particular, using a static tradeoff with a high sensitivity-to-usefulness threshold resulted in higher levels of trust and user satisfaction among participants with domain expertise or low privacy concerns. Moreover, the adaptive request orders resulted in better recommendations because users answered more questions when the algorithm avoided overly sensitive questions.

Recommendation: A hybrid adaptive privacy-setting procedure can increase users' acceptance of and comfort with the adaptations

Research shows that different types of users react differently to automatic default settings. Brown et al. [41] demonstrate that expertise and social intelligence (known as "market metacognition") are important aspects in this regard. Particularly, if users are privacy novices and unaware of the intentions of the various actors involved in their privacy decisions, they are more likely to adhere to the suggested default setting (persuasion). However, users who are privacy experts and aware of the intentions of the various actors involved in their privacy decisions tend to be *less* likely to adhere to the suggested default setting (reactance). For those with either expertise or market metacognition, the effect depends on the value of the default: a low default is persuasive, while a high default causes reactance. Similar effects have been found in privacy, where the effect of default orders and settings depend on users' overall privacy concerns [166,168,170].

In cases where defaults cause reactance, highlights and suggestions may be better able to help users with their privacy decisions. Highlights and suggestions are also more appropriate when the UTP system is less confident about the optimal setting or decision. As such, the mechanisms presented in this section can be combined in various interesting ways. An example of this would be a UTP system that:

- automatically discloses any items to a TLA learning provider that it knows the user is very likely going to be comfortable disclosing;
- automatically withholds any items to a TLA learning provider that it knows the user is very likely going to be uncomfortable disclosing;
- makes suggestions for items and learning providers it is not certain about;
- presents these suggestions in decreasing order of certainty.

For TLA users with a certain level of privacy expertise and market metacognition (which are potentially observable from their extent of making manual privacy decisions), the UTP system can increase its automation thresholds, especially for automatic disclosures.

This combination of various adaptive privacy-setting techniques arguably increases users' acceptance of and comfort with the adaptations. Future work should test this hybrid procedure in a live setting.

5.2 Adapt the justification

Users can also be supported with justifications. A justification is usually accompanied by a recommended setting or action and provides a succinct reason to engage or not engage in the recommended privacy-related behavior. Different types of justifications are covered in Specification Vol. 1, Section 6.3. Here we specifically address the idea of tailoring the justification to the user.

Tailoring the timing, framing, and content of justifications makes them personalized and relatable. This arguably increases the trust and perception of support between the TLA-based system and its' users, enabling them to make better privacy decisions regarding the information used by various learning and training applications [169].

Make justifications context-relevant

Good justifications incorporate meaningful representations of information rather than raw data [140]. Therefore, one way to tailor justifications is to make them context-relevant.

Contextualized justifications are important in TLA-based systems: UTP is employed to tailor the user interface and privacy setting based on the user's privacy preference, so the justification for these user-tailored adjustments should be tailored to the context too.

Indeed, Kobsa and Teltzrow found that websites that had contextual explanations were viewed more favorably by users [187]. As a result, users rated those websites' privacy practices and perceived benefit significantly higher than websites that did not offer any contextual explanations. They also made considerably more purchases on such websites. Similarly, Tsai et al. showed that providing real-time feedback to users of a location-sharing system about when and who queried their location tends to make them more comfortable about sharing location information, and increases their chance of continued use of the system [329].

Likewise, Wijesekera et al. conclude that current privacy notice approaches lack contextual integrity that would help users understand the purpose of a particular permission request [354]. Although users consider some permission requests as inappropriate, the user's preferences towards appropriateness may change over the time. Thus, quantifying the dynamics of user privacy preferences is an important factor to be considered.

Time justifications carefully

Context-relevant justifications are likely more effective if shown when users are actually in the process of making a commitment/decision regarding the use of the app, rather than “post-hoc”. In this regard, Egelman et al. argue that justifications are most impactful early in the process, when users haven’t made up their minds yet [90]. However, research also shows that giving the justifications after the fact may actually result in more accurate feedback [264]. The optimal timing for justifications should thus be studied in future work.

Regardless, real-time contextualized feedback could result in information overload, and may be regarded as intrusive due to the constant interruptions they cause. The frequency and timing of privacy-related justifications is therefore very important, as users should not be overburdened by justifications, and the justifications should not cause an “ironic transparency effect” where the justifications could reduce users’ overall trust and satisfaction due to being reminded of privacy [169]. As such, justifications should only happen in situations where they may have a short-term (e.g. impact the user’s decision) or long-term (e.g. increase the user’s privacy knowledge) impact. For example, Felt et al. suggest to avoid overloading users with unnecessary low-risk privacy warnings in Web browsers [98].

Jedrzejczyk et al. [140] show that users’ willingness to receive a feedback notification depends on several contextual factors such as time, location, and the importance of the information. They find that the user’s current task is an important factor in determining whether giving a notice is appropriate. In other words, if the TLA user is in the middle of a training and unavailable to read the justification, UTP may decide to apply the adaptations without justification, or postpone the justification (or even the adaptation itself) until later. If the user is inadvertently interrupted anyway, a “snooze function” allows them to dismiss the justification without discarding it altogether [140].

Tailor the type of justification (explanation, usefulness, or social norm) to the user’s personal characteristics

Justifications come in various shapes and sizes. Developers therefore need to carefully consider how to construct such justifications and consider how users are likely to interpret them. The following justifications have been tested in privacy-related works:

- Some researchers justify information requests by providing a reason for requesting the information [69,169]. This can be particularly beneficial in personalized systems like TLA, where seemingly unimportant and unrelated information about the user can have an unexpected beneficial effect on the accuracy of learning recommendations. In their overview of justification methods, Knijnenburg and Kobsa found overall negative results,

but explanations were the least problematic justification method, as they did not significantly reduce users' trust, satisfaction or disclosure [169].

- Other researchers justify information requests by highlighting the benefits of disclosure [169,187,341]. A study by Kobsa and Teltzrow [187] showed that users were about 8.3% more likely to disclose information when they knew the benefits of disclosing the information. Knijnenburg and Kobsa found that while benefits-based justifications are considered helpful, they decrease users' trust, satisfaction, and disclosure [169].
- Others justify information requests by appealing to the social norm [8,31,169,263]. Among these, Acquisti et al. find that users were 27% more likely to do this when they learned that many others decided to disclose the same information [8]. But Besmer et al. [31], who use social navigation to help users make better privacy and security decisions using community knowledge and expertise, find that social cues have barely any effect on users' Facebook privacy settings: only the small subset of users who take the time to customize their settings may be influenced by strong negative social cues. Likewise, Patil et al.[263] use aggregate information about the privacy preferences of a users' social circle to help them make informed choices about their own privacy preferences. However, they rate social navigation cues as a secondary effect. Finally, Knijnenburg and Kobsa find that users do not consider social navigation cues to be helpful, even though they are the only studied justification type that influences disclosure [169]. Social navigation also reduces trust and satisfaction in their study.

Not everyone is equally amenable to justifications, and people differ in the types of arguments that may persuade them. In a study of privacy seals, Rifon et al. [277] show that while seals did not seem to ease privacy concerns for *all* users, users with low privacy self-efficacy and users who rarely interact with the site experienced more trust and had a greater intention to disclose personal information when the site displayed a privacy seal. Tailoring privacy seals to the user may thus provide a way to influence users who are most amenable to them, without bothering others. Likewise, Felt et al. [98] suggest customizing privacy warnings in web browsers (which are a type of justification) according to users' personal privacy concern.

Knijnenburg and Kobsa are to our knowledge the only ones who have explored this possibility in some detail [170]. They demonstrate the potential for personalized justifications by taking the results from their overview of justification methods [169], in which they found that all justifications had an overall negative effect on user satisfaction and disclosure, and re-analyzing the data using a personalized approach. Their analysis shows that tailoring the justification method to users' gender and disclosure tendency can increase the effectiveness of the justification. Particularly:

- For females with low disclosure tendency they find that it is best to ask demographics questions first with an "explanation" justification.
- For males with low disclosure tendency, it is best to ask demographics questions first with no justification.
- For females with high disclosure tendency, it is best to ask context questions first with no justification.

- For males with high disclosure tendency, it is best to ask demographics questions first with a “usefulness” justification.

These strategies provide the best balance between satisfaction and disclosure; different strategies are optimal when the goal is to improve either satisfaction or disclosure exclusively.

Pending a comprehensive review of this technique, the framing of justifications can be leveraged by UTP to adaptively influence disclosure

Research has shown that the framing of a privacy decision can significantly influence the level of disclosure, with negative framing leading to significantly lower levels of disclosure than positive framing [144,200]. Justifications can also be used to frame a privacy decision. In a study on mobile app store privacy indicators, Choe et al. found that positively framed icons are more likely to nudge people away from privacy invasive apps [63]. Interestingly, this is counter to the traditional framing effect, in which a negative frame reduces disclosure due to loss aversion.

Pending more research on framing in privacy justifications, this technique can potentially be used in user-tailored privacy to increase or decrease sharing as desired. Note, though, that framing does not always work for everyone: Hardisty et al. [114] show that framing of a carbon fee as either a “tax” or an “offset” has an influence on people who identify as Republicans (who generally have a more negative view of taxes), but not on people who identify as Democrats. Therefore, framing should be applied in a way that fits the user’s context and frame of reference.

Recommendation: Conduct more comprehensive research into adaptive justifications and their delivery

Justifications have the potential to nudge users into a certain direction, so an adaptive application of this approach in privacy decision-making can be a valuable tool for UTP. Moreover, justifications can give users insight into the reasoning behind the UTP recommendations, keep them in the loop on their privacy decisions, and even teach them valuable, context-relevant lessons about privacy.

In theory, the pervasive use of privacy justifications can give users a feeling of support, and establish trust in the system [341]. System developers may be able to capitalize on this trust relationship by anthropomorphizing the source of the justifications, like with Facebook’s privacy dinosaur [251]. Note, though, that findings outside the area of privacy suggest that anthropomorphism may hurt explanations of interface adaptations [174].

As we have shown throughout this subsection, though, very little research on the value of justifications exists, and research on user-tailored justifications is even less prevalent. More research should be conducted into the use of justifications for various UTP-related purposes.

Arguably, this research would be the key to turning UTP from a user-support tool into a tool that increases users' awareness and teaches them valuable lessons about privacy (see Section 6).

5.3 Adapt the interface

User-tailored settings and justifications are most suitable when managing privacy involves explicit decisions, particularly those that can be represented as a checkbox, radio button, or multiple-choice question. However, most modern information systems allow 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 [126,201,238]. For example, in the case of Facebook, existing research demonstrates 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) [360].

These privacy behaviors extend beyond the sharing matrix—they are enabled in Facebook's interface by a variety of designed privacy features. However, these privacy features are often difficult to access, and create an unwieldy “labyrinth” of privacy functionality that users find difficult to use [70,222]. Luckily, research on users' utilization of these various Facebook privacy behaviors has demonstrated that users substantially differ in the extent to which they use each behavior [361–364]. Consequently, UTP can be used to tailor the design of the interface itself to the user, emphasizing features the current user is expected to use most often, and de-emphasizing the features they only seldom use. Recent work labels this approach “User-Tailored Privacy by Design” (UTPbD) and explores how it can be applied directly using user profiles mined from the same system, or indirectly (“extrapolated”) based on “personas” abstracted from research on a different system [356]. Below we describe both approaches in some detail.

Apply UTPbD directly, by measuring privacy profiles and then tailoring the user interface to these profiles

Both UTPbD approaches start with a clustering of users' privacy behaviors into various “profiles” [361]. In a direct application of UTPbD, the privacy controls of the system for which the profiles were developed are tailored in a way that changes their salience depending on the profile of the current user. Research on information disclosure shows that the salience [100,128,131,173,371] 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.

Wilkinson et al. present designs based on this UTPbD approach for Facebook, based on Wisniewski et al.'s profiles. Below are some example interface adaptations they propose.

Selective Sharers want to limit the audience with whom they share information. Their tailored design makes it easier to assign friends to custom friend lists (Figure 2) and increases the prominence of the selective sharing options that Facebook provides when submitting a new post (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.

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 4). 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 5).




Figure 4: Default visibility for Contact and Basic Info is set to “only me”.




Figure 5: 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).

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, their tailored design makes it easier to go offline on chat at any time by means of a toggle button (Figure 6). It also makes it easier to alter their News Feed by deleting stories or hiding posts (Figure 7).




Figure 6: 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.




Figure 7: A more prominent design for Timeline moderation. Users can easily delete or hide posts on their Timeline.

Privacy Balancers display moderate levels of privacy management. Designing for these users is particularly hard; the proposed solution is to make certain key privacy features more prominent. Particularly, their tailored design increases the accessibility of the settings for restricting their chat availability, timeline moderation, altering post on their News Feed, and blocking apps, events, and people (Figure 8).



Figure 8: A more prominent design for blocking apps, events, and people, displayed directly in the notifications.

Privacy Maximizers display the widest variety of privacy behaviors, i.e. they utilize almost all the available privacy features. Therefore, increasing the accessibility of all aforementioned privacy activities. Combination of these emphasized features is likely going to result in a considerably more cluttered interface. However, Privacy Maximizers are likely to have such strong privacy concerns that they prefer this cluttered interface over the standard Facebook interface (cf. [344,345]).

Finally, 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.

Apply extrapolated UTPbD, by leveraging profiles from other systems as “personas”, and designing tailored interfaces for these personas

Aside from a direct application, UTPbD also allows for an extrapolated application, where the user profiles identified in one system are used as “personas” to develop privacy design guidelines for a yet-to-be-implemented system that 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 [10,11]. 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.

To explore extrapolated UTPbD, Wilkinson et al. [356] use the same six profiles uncovered by Wisniewski et al. [361] to develop user-tailored design guidelines for TLA. Users of TLA-based 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, there are several ways in TLA-based systems can adapt their privacy interface to these different types of users.

In a network of training applications, Selective Sharers may be selective regarding the applications that they are willing to use. Selective Sharers would also be more restrictive regarding the social aspects of TLA. Specifically, they would be likely to carefully manage who is within their network, what training outcomes are posted publicly, and who gets to see the information that is collected or generated by the training system or that other people share about them. This means that for Selective Sharers TLA-based systems should increase the prominence of 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.

Self-Censors tend to manage privacy by withholding information, so TLA-based systems should allow these users to limit the extent to which the system collects and tracks information about their skills, interests, training schedule and achievements are shared with the system. The system should also emphasize its browsing functionality, since it can expose Self-Censors to relevant training modules without the need for extensive tracking. If Self-Censors refuse to share their learning outcomes with the system, this could prevent them from getting credit for these learning activities, 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 coarse level (e.g. 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. 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.

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 and indirect communication. They should thus have the ability to opt out of social connectivity features such as chat or status updates if the TLA implementation has such functionality. Similarly, TLA-based systems should allow Time 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 and allowing them to switch off push notifications and emails sent out by the system.

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. Privacy Balancers should get the same functionality as Time Savers, 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.

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.

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.

Recommendation: Implement extrapolated UTPbD at design-time, then collect data for a direct implementation of UTPbD at deployment-time

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 UTPbD [272,273]. However, the TLA specifications have yet to be implemented in an actual learning ecosystem. Therefore, the direct approach for creating user profiles based on users' past privacy behavior within the system is not feasible at this stage in the development of TLA.

Therefore, at this point the only feasible approach is to extrapolate previously developed profiles to personas that can inform UTPbD for TLA. While this approach takes an untested theoretical jump, 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.

5.4 Adapt the personalization

A final venue for UTP to affect users' privacy decisions is by modifying the outcome of TLA's data collection practices: the learning recommendations. This UTP approach acknowledges that users may differ in their attitudes towards data *collection* versus data *use*. For example, a TLA user may not mind the system knowing about a certain qualification (e.g. a language skill) but may prefer the system to refrain from using this skill as input to training recommendations (e.g. if the user no longer wishes to work in the region where this language is spoken).

This section surveys differences in attitudes between data collection and use and discusses a user modeling framework that can adjust itself to users' privacy practices.

Use UTP to prevent erroneous inferences and other types of unwanted personalization

The conceptual difference between users' attitudes towards the collection of personal information and its use was first carefully examined by Smith et al. in their development of the Concern For Information Privacy (CFIP) scale. Their factor analysis shows that there is indeed a difference in users' "concern that extensive amounts of personally identifiable data are being collected and stored in databases" (Collection), their "concern that protections against deliberate and accidental errors in personal data are inadequate" (Errors), their "concern that data about individuals are readily available to people not properly authorized to view or work with this data" (Improper Access), and their "concern that information is collected from individuals for one purpose but is used for another, secondary purpose (internally or after disclosure to an external party) without authorization from the individuals" (Unauthorized Secondary Use). Across three different samples, they find that the levels of the latter three concerns are higher than the concern about Collection.

In TLA-based systems, personalized learning recommendations may not always be regarded as the primary purpose of the system, and it is not unlikely that some personalized functionalities will be introduced at a later time [137]. Given users' higher level of concerns about secondary use of personal information, privacy experts argue that such secondary use of the information should be explicitly communicated to the users, otherwise they may be surprised to find out about it, and feel that their privacy is violated [79,321].

Likewise, while users are likely to value the personalized learning recommendations provided by TLA-based systems, they will get annoyed if the system makes an incorrect prediction or inference about their goals and preferences (see Specification Vol. 1, Section 2.3). Possibly worse yet are unwanted or creepy correct predictions [295,322], which cause users to engage in some kind of "reputation management" when using personalized systems. For example, in interviews regarding the data collected by a mobile app recommender, Knijnenburg et al. [169] found that users would occasionally decide not to disclose a certain piece of information because "it doesn't accurately represent me as a person". To mitigate these problems, researchers suggest that users should have the opportunity to scrutinize [152] and correct [103] potential mistakes in the system's predictions.

Importantly, though, the suggested mitigations are classical examples of notice and choice, which will likely leave the user overwhelmed if implemented in a system as complex as TLA (see Section 1.3). Hence, we should Implement UTP to make privacy-aware personalization more manageable.

Recommendation: Integrate UTP into a dynamic privacy-enabling user modeling framework

Wang and Kobsa’s dynamic privacy-enabling user modeling framework (Figure 9) can act as a basis for user-tailored privacy-aware personalization [342]. This framework comprises a User Modeling Server (UMS) that is similar to TLA’s “Processor” entity: it stores user characteristics and behavior, integrates external user-related information, applies user modeling methods to derive additional assumptions about the user, allow multiple external user adaptive applications (cf. User Facing Apps created by Activity Providers) to retrieve user information from the server in parallel. The adaptive functionality provided by the UMS is codified in a number of User Modeling Components (UMCs), which represent classes of inferences that can be made. Importantly, the server allows for each user to have their own privacy constraints on these UMCs, e.g. related to the use of certain data, or the ability to make certain inferences. At the beginning of the interaction with a user, the Selector verifies for every UMC whether it may operate under the user’s privacy constraints.

A fundamental question surrounding the real-world use of such UMSs, is how to determine the privacy constraints for each user. Wang and Kobsa primarily consider prevailing privacy laws and regulations based on the country of residence of the user [342,343]. Our suggestion is to support a much more granular level of personalized constraints based on UTP. Users’ privacy preferences as modeled by UTP (see Section 4) can serve as a basis to determine whether certain types of inferences or data use should be allowed, and thus steer the selection of UMCs in a more dynamic fashion.




Figure 9: Wang and Kobsa's Dynamic Privacy-Enabling User Modeling Framework [342]

5.5 Problems with adaptations

The adaptations discussed in this section span a wide variety of “degrees of automation”, as outlined by Sheridan and Verplank [293], ranging from “the computer offers no assistance” to “the computer decides on everything and acts autonomously”, with several levels in between. In light of these various degrees of automation, we present two important problems in this section, as well as a potential solution.

The first problem is the trade-off between *burden* and *control*: Automation (and subsequently higher degrees of automation) reduces the *burden* of having to engage privacy-related behaviors oneself, but it comes at the cost of relinquishing some *control*. The first problem is thus one of finding the optimal degree of automation that balances this trade-off for TLA users.

The second problem involves the persuasive effect of the UTP-based adaptations (see e.g. [177,359]), which arguably increase with the degree of automation. While the goal of UTP is to nudge TLA users to employ the privacy behaviors that match their observed or stated preferences, one needs to be careful not to turn UTP into a privacy “dark pattern” [37].

Finding the optimal balance between burden and control is a difficult trial-and-error process

While it is clear that the degree of automation determines the balance between (the reduction in) user burden and control, existing literature provides little guidance regarding optimal degree of automation that balances burden and control. We argue for UTP based on the fact that the vast complexity of TLA-based systems leaves their users ill-equipped to take full control over their privacy, regardless of the usability of the proposed control mechanisms. As such, we advocate for UTP as a means to reduce user burden. On the other hand, though, we acknowledge that UTP cannot function without informed input, and we strive for TLA users to be informed about and have control over their privacy decisions—at least regarding aspects of their privacy that are truly important to them. Hence, we do not want to fully relinquish control.

Bokhove et al. evaluate 15 technical and behavioral privacy control strategies, and show that granular control and usability are fundamentally at odds with each other [35]: most strategies that provide granular control also cause a significant amount of annoyance, and easy-to-use privacy control methods usually fail to provide granular control.

Patil et al. demonstrate the difficulty of finding the optimal level of control in a location-sharing study [73]. They find that immediate feedback without the opportunity to control actions increases user discomfort with location sharing and leads to feelings of oversharing. However, they also find that providing feedback immediately after system-enacted disclosures may create a heightened sense of disagreement with the decision. They therefore advise that the optimal level of control is delayed disclosure feedback, allowing for a reasonable ‘cooling off’ period before the user gives feedback on the automated location-sharing decision.

Hence, we argue that the degree of UTP automation implemented in TLA-based systems should fundamentally be considered a matter of burden and control. TLA stakeholders should recognize the trade-off between these two values, explicitly state their opinion regarding its optimal balance, and implement adaptation methods accordingly. Ideally, these adaptation methods should be evaluated after implementation along the dimensions of burden and control. A comprehensive field trial may point out whether the implemented UTP adaptation methods take over too much control, or rather leave too much of the burden on the user.

Finding the optimal level of persuasion is a difficult trial-and-error process

Research shows that users are prone to agree with a recommender’s predicted ratings [72] and to follow a recommender’s advice [77,110]. Specifically in the field of privacy, Knijnenburg and Jin show that users in their study were more likely to follow the privacy recommendations than what would be expected based on the precision of their recommender as determined in their offline evaluation [177]. Similarly, Wilson et al. demonstrate that their personalized privacy

wizard influences users to share significantly more without a substantial difference in comfort [359].

Higher degrees of automation, and for instance more strongly worded justifications, likely result in even more persuasive recommendations. We argue for UTP based on the fact that we want to make it easier for users to engage in the privacy behaviors that fit their preferences. As such, we welcome the persuasive nature of UTP. On the other hand, though, we acknowledge that UTP should not become too paternalistic, and give users the opportunity to deviate from the recommendations, even if they are informed by their own previous preferences and behaviors.

An overly persuasive recommender is bad for three reasons: First of all, Knijnenburg et al. demonstrate that online services can use persuasive recommendations manipulate user’s disclosure behavior [177]. For example, they can provide over-disclosing privacy recommendations as a means to increase their disclosure. Secondly, overly persuasive privacy recommendations can be perceived as untrustworthy [169], which may cause reactance (see Sections 4.1 and 5.1)—behavior which is opposite to the effect we want to achieve. Finally, the persuasive nature of recommender systems creates what Lanier calls a “positive feedback loop” [202]: rather than going through the trouble of developing their own privacy preferences, simply follow whatever UTP recommends them to do.

Hence, we argue that the UTP automation methods implemented in TLA-based systems should also be evaluated along the dimension of persuasion. A comprehensive field trial may point out whether the implemented UTP adaptation methods are too persuasive, or rather fail to convince users to pursue the suggested privacy-related behavior. The type and wording of the implemented methods can be tweaked according the results of this field trial.

Recommendation: Use different privacy adaptation methods for different situations

The optimal privacy adaptation method remains an open question [180], and a comprehensive field trial can shed light on the best approach to implement the privacy recommendations provided by UTP. Beyond this, we argue that the optimal adaptation method may depend on the *goals* of UTP, be it automation, awareness, guidance or education (see Section 6). For example, according to Sheridan and Verplank [293], lower degrees of automation are most suitable in situations where control is desirable, and a certain amount of burden is justifiable, e.g. for the purpose of educating users about a privacy feature.

Moreover, Namara et al. demonstrate that the optimal adaptation method may differ from feature to feature, and crucially depend on the user’s prior engagement with each feature [240]. They show that when users are unfamiliar with a privacy feature, they prefer adaptations in terms of explicit suggestions, mainly because this allows for the adaptive behavior to be explained. This additional layer of control outweighs the burden of having to follow up on the suggestion. On the other hand, when users use a feature frequently, they prefer a fully

automated approach, suggesting that they are willing to give up some control in return for the significant reduction in the burden that this approach offers them.

6 Goals

This document presents UTP as a means to provide TLA users with personalized privacy decision-support. It covers the operation of UTP as a “measure, model, adapt” framework, indicating how it can give users advice that is tailored to their privacy preferences. So far, though, we have not covered *why* UTP should provide such support—at least not beyond the rudimentary notion that privacy decisions are difficult and existing solutions for support are lacking.

The foremost goal of UTP is to support the user in their privacy decision-making, but in providing such support one must acknowledge that users themselves may have conflicting goals, and it will be up to UTP developers to decide which of these goals to prioritize. Alternatively, UTP can support users by taking on a “teaching role” and giving them the tools, they need to decide for themselves on how to meet their privacy goals.

To further complicate things, UTP can take the privacy requirements of the requester/recipient of the information, other users, and the organization as a whole into account as well. This is particularly difficult when those goals conflict with the user’s goals. In that case, UTP needs to reconcile various requirements in a way that does not betray the trust of the user.

6.1 Support the user

The primary goal of UTP is to support the user. This section covers practical and conceptual obstacles that can make it difficult for UTP to attain this goal:

- TLA users may not trust UTP to make privacy decisions for them
- TLA users’ privacy preferences may be the result of conflicting goals
- TLA users may want control over their data without spending too much effort on privacy decision-making

While these problems are essentially impossible to overcome, this subsection discusses ways to mitigate them.

*Trust is a prerequisite for supporting the user
and may depend on the entity operating UTP*

For UTP to effectively support the user, it needs to be obvious to the user themselves that this is the goal. Users’ acceptance of default settings or privacy recommendations is dependent on their “market metacognition” [41] (see Section 5.1): users will only follow such advice if they believe in the benevolence of the recommending authority [229]. If the user does not trust UTP to act in their best interest, they will likely show reactance to the recommendation and thereby thwart the goal of supporting the user.

Monetizing users' personal information is a common profit strategy of commercial Web applications. As such, users may believe that the goal of a typical learning application is to collect as much information about its users as possible, which may conflict with the user's own goals of keeping (some of) their information private. Hence, the user may not trust such learning applications to make privacy recommendations with their best interests in mind. Due to this apparent conflict of interest, it is thus better for UTP to be implemented by a higher-level entity that is independent from the learning applications [166]. The TLA architecture has several entities that operate independently and at a higher level than the learning applications, for example, the "TLA processors" that provide meta-adaptations [101]. The TLA processors would thus be an appropriate entity for operating UTP and providing privacy recommendations.

Another option is to build UTP as a separate entity, potentially putting it under the control of a trusted party. Training department managers could be this trusted party, as it is their job to provide the best training experience to the people in their department.

*User-supporting UTP must reconcile users' various,
potentially conflicting goals*

UTP manages users' disclosure behavior in a way that meets their privacy preferences. It is important to note that those preferences arise from various, potentially conflicting goals.

The user's primary goal in using TLA is to get a personalized learning experience through meta-, macro-, and micro-adaptations that tailor the learning activities to their needs and preferences. From this perspective, collecting more user data generally equals a better, more personalized experience [311]. Therefore, UTP can best support this goal by asking users to provide as much information as they are comfortable with. This is not a trivial task, because users' comfort cannot simply be inferred from their disclosure behavior. For example, using the "door-in-the-face" technique (i.e., asking the most sensitive question first when asking multiple disclosure questions) has been demonstrated to increase disclosure [8] but may severely increase privacy concerns among users [168]. Given that privacy recommendations are persuasive [177], UTP must be careful not to assume that users' agreement with recommended disclosures indicates that they are comfortable with these requests.

Moreover, a secondary goal in using TLA is to remain private. As argued in Specification Vol. 1 (Section 2.2) that effective learning can be inhibited by continuous observation and tracking. From this perspective, collecting *less* user data generally equals a more suitable learning environment. UTP can attain this goal of privacy by providing a meaningful experience with the least amount of data, i.e., by only asking for data that are actually useful to the requesting application.

These perspectives are endpoints on a spectrum, which depends on what kind of personalization is considered "meaningful", and what kind of tracking users consider to be "comfortable". UTP can help determine what the user finds useful and what they are concerned about, but the

tradeoff between concerns and usefulness is essentially a designer question. Personalizing this tradeoff can circumvent some of the ethical considerations of making this tradeoff as a developer [302], but even then, a certain threshold needs to be set. For example, if UTP predicts the probability of a user disclosing a piece of information to be at a certain percentage, then what percentage should serve as the threshold above which it is acceptable to recommend the user to disclose this piece of information? And at what percentage can the disclosure occur automatically? The answer to these questions depends on how the system prioritizes the user's goals.

Recommendation: Use the concept of stewardship to provide an ideal balance between effort and control

A related question concerns the amount of control UTP gives to the user. Using different recommendation targets (see Section 4.3) and adaptation methods (see Section 5.1), UTP comprises a broad spectrum of approaches ranging from full automation to lightly assisted manual control.

On the “full automation” side, UTP could aim to exactly mimic users' own privacy behaviors and apply them automatically. This approach relieves the user from the cognitive and physical effort of taking those actions. There are a few problems with this approach. For one, if UTP takes over all interactions with TLA-related privacy settings, users never get the opportunity to indicate their preferences. Worse even, if privacy settings are fully automated, TLA users could become disinterested and ill-equipped to set such settings manually, should the need arise. This could result in an erosion of user autonomy [182] that could ultimately result in severe privacy violations.

On the “manual control” side of the spectrum, UTP could provide privacy suggestions, and/or teach the user how to deal with privacy settings by themselves (see Section 6.2). This approach allows users to “stay in the loop” regarding their privacy settings, allowing them to double-check each privacy-related decision. There are a few problems with this approach as well, though. First of all, it is difficult to motivate users to actually take control over their privacy settings [68], especially when they are numerous and complex. As such, it is not unlikely that users will avoid the cognitive effort of making privacy decisions and approve any suggested actions without any critical evaluation. This creates the same situation as “full automation”, but with a false sense of control (which can in itself erode privacy [39]) and without reduced physical effort.

A possible middle ground can be found in the concept of “stewardship”, where UTP makes most of the individual privacy decisions, but the user controls the general rules by which these decisions are made—for example, the general prioritization between privacy and benefits as described above. UTP can follow these rules, raise exceptions where appropriate to ask for user feedback, and incorporate such feedback as a means to update the rules.

This “stewardship” approach minimizes the required physical effort and distills the required cognitive effort to a more conceptual expression of the user’s general goals. Should users still feel ill-equipped to express their privacy preferences at this level, then they can decide to share their “stewardship” with a trusted third party, such as a training department manager (see Specification Vol. 1, Section 4.2).

6.2 Teach the user

Most existing works on UTP essentially (partially) *replace* the user’s privacy decision-making practices by taking (partial) control over the user’s privacy settings. Alternatively, UTP can actively *support* users’ privacy decision-making practices by helping them to decide for themselves how to meet their privacy goals. This subsection discusses several approaches to “privacy education”.

Personalized privacy tips can help solidify users’ privacy-related behaviors

An opportunity for privacy training exists whenever UTP’s recommendations go beyond what the user currently does. One way to give TLA users personalized “tips” about privacy decision-making is by highlighting inconsistencies in their current behavior. This requires a somewhat prescriptive approach, where privacy-related behaviors are deemed “consistent” or “inconsistent” based on logical interdependencies between privacy settings [179]. For example, if a TLA user indicates that they don’t want to share a certain learning outcome with Bob, but they are sharing that learning outcome in a public forum that Bob also frequents, the latter privacy behavior logically contradicts the former.

Such logical inconsistencies can be detected through formal logical analysis, although this requires that the consequences of privacy actions can be exhaustively represented in a machine-readable format. Alternatively, the detection of such inconsistencies can be “crowdsourced” either by following the average settings of a user’s nearest neighbors (assuming that such aggregated privacy behaviors are more consistent than each individual’s behaviors) or by following the advice of “privacy experts” (assuming that such experts exists and can be detected).

A more generalized approach to “privacy tips” is discussed in Section 4.3: The system can highlight privacy-related actions that complement users’ current practices. This requires knowledge about synergistic privacy actions, which fortunately is a topic previous user-tailored privacy research has paid a considerable amount of attention to [172,178,361].

Privacy education through self-actualization gives users more confidence in their overall privacy strategy and supports the evolution of this strategy

Section 4.3 also discusses the opportunity to move beyond users' current privacy practices by teaching them *new* privacy practices that they are currently unaware of. This approach to privacy education is equally amenable to the user-tailored approach: users should not be taught practices they already engage in nor practices they have decided they do not want to engage in [361].

Like "privacy tips", the goal of this educational practice is to actively involve users in the privacy decision-making process. However, whereas privacy tips make users aware of synergistic privacy practices with the goal of fixing the flaws in users' current behaviors, making users aware of alternative privacy practices allows the user to discover new ways of thinking about privacy and adopt these new mechanisms if they want. This supports the more long-term goal of giving users more confidence in their overall privacy strategy and the evolution of this strategy. We therefore call this approach "privacy education through self-actualization". Recent research shows that users will likely appreciate this personalized educational approach [240].

Recommendation: *When employing personalized privacy education, make sure to actively involve the user and carefully explain the recommendations*

As privacy education typically involves recommending actions outside the user's current purview, such recommendation should not be made automatically but rather be based on active user involvement. Indeed, Namara et al. [240] find that the best approach for recommendations that go beyond the users' current actions is to "nudge" users towards these behaviors using subtle design interventions or suggestions (see Sections 4.3 and 5.1). This ascertains that users are actively involved in the recommendation process. As noted in Section 5.3, changes to the user interface can also ascertain that the user retain active control over their privacy decision-making.

Another way to support self-actualization is to explain how the privacy suggestions provided by UTP have come about. This could give the user insight into their own privacy practices and teach them how to evolve beyond their current practices. The field of recommender systems has shown that explanations increase users' understanding of the recommendation process [107,337], thereby increasing the opportunity for self-actualization. Moreover, while explanations are useful to support any privacy recommendation (see Section 4.2), it is particularly important to explain suggestions that go beyond the user's current practices, as the relevance such suggestions may not be readily apparent to the user. Namara et al.'s work demonstrates that users appreciate explanations of privacy adaptations, especially when they are not intimately familiar with the adapted privacy functionality [240].

Going beyond explanations of the privacy recommendations, UTP could also give users a better insight into their own privacy practices, e.g. by visualizing these practices, how they evolved over time, and how they compare to other, similar users. While currently untested, such practices have been proposed in the field of recommender systems as novel interaction mechanisms to increase self-actualization [182].

6.3 Help the recipient

While the primary goal of UTP is to support and or educate the user, this does not necessarily mean that it has to ignore the interests of the other parties involved in the information exchange that happens when users use a TLA-based learning application (i.e., the eventual “recipients” of the collected information). Particularly, UTP can also help the learning activity providers, researchers, supervisors, and the organization itself attain their goals.

This subsection outlines the goals of these other actors and discusses how TLA can support them. Note that the goals of these actors may at times conflict with the user’s goals, which creates an ethical dilemma. Section 6.4 explains how to deal with this dilemma.

*UTP can be used to improve personalization, research,
and personnel-related decision-making*

The goals of other parties involved in a deployed TLA environment are generally best served if they can collect as much data about the user as possible. For a learning activity provider, user data can be used to improve its curriculum (e.g., if users are disproportionally failing a certain learning unit then that unit can be overhauled) and to increase its personalization power (i.e., it is a common practice in the field of recommender systems to periodically or even continuously re-train the algorithms based on collected user data [274]). UTP can try to convince users to make their learner runtime activity available to the application (or even make such data available automatically, without asking the user at all), but Specification Vol. 1 (Section 2.2) notes that this data can be particularly sensitive.

TLA users may also benefit from each other’s disclosures. This benefit can be derived directly, as it allows users to compare their learning progress and performance with their peers (social learning experiences are covered in more detail in Specification Vol. 1, Section 5.2). Users can also indirectly derive benefit from others’ disclosure through the provision of better-informed or even socially traceable recommendations (see above). UTP can thus help users attain social learning benefits by recommending them to share data with each other.

Education researchers can use collected user data to discover new insights about learning strategies. For example, by analyzing the data of hundreds of teams in a collaborative training exercise, researchers can determine the most effective collaboration strategies, and incorporate these strategies into new training modules. Recent scandals with Facebook [134,265] and Dropbox [88] suggest that researchers should not use such data without consent from the user.

Hence, UTP can help attain this goal by convincing users to share the data that pertain to the studied phenomenon with the researchers.

TLA users' supervisors may want to take as much of their data into account as possible to make promotion and mission planning decisions (see also Specification Vol. 1, Section 5.3). Military deployments increasingly consist of small, synergistic teams of individuals with complementary skills and an affinity for the location or the situation in which they are deployed. Similarly, requirements for promotion have become increasingly based on evident mastery of skills. Note, though, that a soldier's skills are increasingly captured in "micro-credentials" and thus harder to represent in traditional degrees or certificates [94]. As such, the goals of supervisors and strategists would be supported if TLA would make more detailed data about skills and affinities available to them.

UTP can be used to meet organizational constraints

The organizational entity behind the TLA user may have rules outlining that certain data collection practices are required or rather prohibited. In these cases, UTP will have to adhere to these rules, regardless of the user's (or even the recipient's) preferences.

On the "required" side of this equation, a user may want to keep certain training data private, but the entity paying for the training activity may require it to be disclosed—at least to said entity. Such organizational constraints can prevent the free-riding behavior of users who want to reap the benefits of TLA-based training without paying the price of having their data collected by their employer. That said, organizational data collection requirements should stay within legal limits: the collection of data in off-duty situations should never be required; certain activities (such as going to the restroom) are private even when the user is on duty, and care must be taken that collected data is not at such detailed level that protected private information (e.g., religious preferences, sexual orientation) can be derived from the data.

On the "prohibited" side, it is not always possible for e.g. activity providers to collect detailed data about the user due to organizational constraints. For example, in a military context one must be very cautious about potential leaks of mission-critical information through training data. An activity provider may thus have to provide strong security assurances before it can be allowed to collect any training data, and researchers may need security clearance before they can inspect data for research purposes.

In both of these situations, (tailored) justifications can help explain why these organizational constraints are in place. This can mitigate the negative impact of organizationally mandated or prohibited data collection practices.

Recommendation: To maintain or increase users' long-term data disclosure, focus on increasing their trust in the recipients of the data

As mentioned at the beginning of this subsection, the goals of other parties involved in a deployed TLA environment are generally best served if they can collect as much data about the user as possible. Note, however, that getting more data does not always mean increasing data collection efforts. In fact, if TLA users feel pressured to provide data they may provide fake or lower-quality data [52]. So even when the goal is to collect as much data as possible, it serves the UTP system to be mindful of the user's privacy preferences.

Another reason not to forcefully maximize data collection is that the recipient's longer-term data collection goals may be best served building a trust-based relationship with the user. Building user trust is usually in the best interest of the recipient, because as Specification Vol. 1 (Section 1.1) points out, trust is the most important antecedent of disclosure. If UTP practices increase users' concern about a certain recipient (or the TLA-based system as a whole), their trust is likely to go down. Hence, UTP should collect data in a manner that does not violate users' trust.

An important aspect in maintaining users' trust is to ostensibly inform users about how the UTP-based recommendations reconcile the goals of the various stakeholders involved. This practice is further discussed in Section 6.4.

6.4 Reconciling different goals

The previous subsections have highlighted the fact that UTP may have to serve multiple stakeholders. In complex organizational information systems like learning management systems based on TLA, UTP may not only be required to provide privacy support to the user; it may also be required to consider the goals of the learning activity providers, TLA researchers, supervisors, other users, and the organization as a whole. In such multi-stakeholder environments, UTP will have to learn what everyone wants, and then decide on how to adapt.

The goals of these stakeholders may be in conflict, so it will be necessary for UTP to make a tradeoff between these various goals. Moreover, as outlined in the previous subsections, each entity may have multiple, internally conflicting goals. This means there are also internal tradeoffs that need to be made. This subsection addresses technical mechanisms and ethical principles for making such multi-stakeholder tradeoffs.

The field of group recommender systems and group decision-making can inspire multi-stakeholder decision strategies for UTP

The field of group recommender systems has studied various ways to integrate the preferences of multiple stakeholders in the recommendation process. At a rudimentary level, a

recommender system can simply take the average of users' preferences. More sophisticated methods include voting schemes and methods that avoid options that are strongly disliked by at least one of the stakeholders [226]. The latter methods have particular merit in UTP, where avoiding decisions that go squarely against some stakeholders' preferences is arguably more important than finding the best possible aggregate option.

Research into how people themselves make decisions in a group shows that they mainly use simple strategies [225]. Moreover, just like how individual preferences are often constructed on the fly [32], so are group preferences [81], which suggests that stakeholders should be able to discuss their preferences with each other [242]—a process that may be considered too involved when it comes to privacy settings.

Ethically speaking, UTP should put the end-user first and apply the concept of reciprocity to reconcile conflicting goals among multiple stakeholders

Beyond these technical approaches, UTP should consider ethical principles in its privacy decision-making process. One ethical principle for reconciling the goals of multiple stakeholders is to always put the end-user first. This principle makes sense both from a Kantian perspective (since users' data is the means of TLA user modeling, serving the user's needs should also be its primary goal), as well as from a utilitarian perspective (since in the end, the success of TLA is dependent on the participation of its users, which requires that they are comfortable using TLA-based systems).

This ethical principle can be practically implemented in various ways, but one useful formulation of a universal rule could be: "UTP can take others' goals into consideration only to the extent that they do not unreasonably conflict with the user's goals". Note that this rule requires a clear understanding of users' goals and preferences, which is a particularly vexing problem in the field of privacy, where users may have internally conflicted goals and may not even understand their own preferences.

This ethical principle does not give sufficient recourse in cases of organizational constraints (which often cannot be traded off against users' preferences) or conflicts between users (who may have equal stake in the management of a piece of co-created data, e.g. in the case of a team exercise). In these cases, UTP should appeal to notions of fairness [91] and reciprocity [256]: certain privacy-related practices may involve a give-and-take between multiple users ("you get to share X if I get to share Y"), and certain practices may be disadvantageous to the individual user but benefit the group. Making users aware of these dynamics is arguably a better strategy than keeping them tacit.

Recommendation: Allow users to manage UTP's multi-stakeholder privacy decision-making practices through a reintroduction of transparency and control

This subsection has discussed practical and ethical considerations regarding the reconciliation of multi-stakeholder privacy preferences in UTP. From both perspectives, it seems advantageous to honestly inform users about the optimization strategy of UTP, i.e., to explain to them how different conflicting internal and external goals are taken into consideration. If necessary, the user can subsequently be given the option to control the multi-stakeholder privacy decision process; i.e., they can adjust their preferences and/or the used tradeoff policies. This strategy highlights reciprocity-based decisions and also allows for a discussion/negotiation of preferences.

This idea brings back the notion of transparency and control to the privacy decision-making process, but in this case transparency and control operate at a higher level than their traditional application: Users are not asked to make individual decisions, but to reflect upon UTP's policies for making privacy at a meta-level. This principle is not only useful for multi-stakeholder privacy decision-making, but for privacy decision-making in general: UTP can manage the user's privacy settings, and users can manage the principles by which UTP operates.

7 Conclusion

In this document we introduced the concept of User-Tailored Privacy (UTP) as a means to support the privacy management practices of the users of TLA-based systems. As UTP relies on user modeling to provide personalized privacy decision support, this document provides recommendations regarding the Modeling Factors involved in personalizing privacy support for TLA. These recommendations will allow ADL and other TLA performers to implement UTP in TLA. The recommendations in the current version of this document are tentative; the text will be updated (Volume 2.1) based on the outcomes of the User-Tailored Privacy Summit. Moreover, the specifics of the various modeling factors will be decided upon after intensive discussion with ADL and other TLA performers during the development of final version of this document.

In the meanwhile, we **made a case for UTP** by highlighting the shortcomings of technical solutions, privacy by design, notice and choice, and privacy nudging. We recommend that TLA performers integrate technical privacy-preserving solutions but complement them with user-centric that support users' privacy management practices. Many of these user-centric solutions can follow the concept of Privacy by Design (see Specification Vol. 1), but where this concept cannot resolve privacy problems they should employ the personalized approach afforded by UTP.

In this document, we **define** UTP as a “measure, model, adapt” framework, and recommend that TLA performers implement this in TLA's learning applications and social learning practices. Particularly. A TLA privacy API should be developed for this purpose.

In **measuring privacy**, TLA performers should acknowledge the plurality and multi-dimensionality of users' decision-making practices. They should also note the variability of users' privacy practices can often be captured by a concise set of “privacy profiles” and that data recipients can often similarly be grouped into a number of groups or “circles”.

In **modeling privacy**, we particularly note that matching the users' current privacy practices may not always be the best modeling strategy; in certain cases, UTP should recommend complementary practices, while in other cases UTP can completely move beyond users' current practices. TLA performers should carefully balance these various approaches. Moreover, since privacy modeling may not always be successful, TLA performers should build UTP as a layered and gracefully degrading modeling component.

In **adapting privacy**, UTP can personalize the privacy settings of a TLA-based application, the justification it gives for requesting certain information, its privacy-setting interface, and its learning recommendation practices. TLA performers should carefully balance proactive and conservative adaptation strategies in order to give users reduce users' burden but at the same time give them sufficient control and reduce undue persuasion.

Finally, TLA performers should carefully consider the various **goals** that UTP can support. They should acknowledge that UTP must reconcile users' various, potentially conflicting goals, and

they should balance the goal of replacing the users' privacy decision-making practices with the longer-term goal of teaching them about privacy. Moreover, TLA performers should consider that UTP's support can help other stakeholders in the privacy decision-making process as well. Regarding this, they should carefully consider how to reconcile the potentially conflicting goals of these various stakeholders.

References

1. Michael Aagaard. 2013. How Privacy Policy Affects Sign-Ups – Surprising Data From 4 A/B Tests. *ContentVerve.com*. Retrieved May 28, 2013 from <http://contentverve.com/sign-up-privacy-policy-tests/>
2. Accenture. 2015. U.S. Consumers Want More Personalized Retail Experience and Control Over Personal Information, Accenture Survey Shows. Retrieved March 17, 2015 from <http://newsroom.accenture.com/news/us-consumers-want-more-personalized-retail-experience-and-control-over-personal-information-accenture-survey-shows.htm>
3. Mark S Ackerman, Lorrie Faith Cranor, and Joseph Reagle. 1999. Privacy in e-commerce: examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM conference on electronic commerce (EC '99)*, 1–8. <https://doi.org/10.1145/336992.336995>
4. Alessandro Acquisti. 2004. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, 21–29. <https://doi.org/10.1145/988772.988777>
5. Alessandro Acquisti. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security and Privacy* 7: 82–85. <https://doi.org/10.1109/MSP.2009.163>
6. Alessandro Acquisti and Jens Grossklags. 2004. Privacy Attitudes and Privacy Behavior. In *Economics of Information Security*, L. Camp and Stephen Lewis (eds.). Springer US, 165–178.
7. Alessandro Acquisti and Jens Grossklags. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3, 1: 26–33. <https://doi.org/10.1109/MSP.2005.22>
8. 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>
9. Alessandro Acquisti, Leslie K. John, and George Loewenstein. 2013. What is privacy worth? *The Journal of Legal Studies* 42, 2: 249–274. <https://doi.org/10.1086/671754>
10. 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.
11. 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>
12. Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (SOUPS '13)*, 9:1–9:11. <https://doi.org/10.1145/2501604.2501613>
13. William F Adkinson, Jeffrey A Eisenach, and Thomas M Lenard. 2002. *Privacy Online: A Report on the Information Practices and Policies of Commercial Web Sites*. Privacy & Freedom Foundation.
14. Gediminas Adomavicius and Alexander Tuzhilin. 2011. Context-Aware Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira and Paul B. Kantor (eds.). Springer US, Boston, MA, 217–253.
15. Zeynep Ahmet and Kaisa Väänänen-Vainio Mattila. 2012. Mobile Service Distribution from the End-user Perspective: The Survey Study on Recommendation Practices. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems (CHI EA '12)*, 573–588. <https://doi.org/10.1145/2212776.2212831>
16. Muhammad Aljukhadar, Valerie Trifts, and Sylvain Senecal. 2017. Consumer self-construal and trust as determinants of the reactance to a recommender advice. *Psychology & Marketing* 34, 7: 708–719. <https://doi.org/10.1002/mar.21017>
17. I. Altman. 1975. *The environment and social behavior*. Brooks/Cole Pub. Co., Monterey, CA.
18. Yair Amichai-Hamburger and Gideon Vinitzky. 2010. Social network use and personality. *Computers in Human Behavior* 26, 6: 1289–1295. <https://doi.org/10.1016/j.chb.2010.03.018>
19. Eduardo B Andrade, Velitchka Kaltcheva, and Barton Weitz. 2002. Self-Disclosure on the Web: The Impact of Privacy Policy, Reward, and Company Reputation. In *Advances in Consumer Research*, Susan M Broniarczyk and Kent Nakamoto (eds.). Association for Consumer Research, Valdosta, GA, 350–353.
20. Claudio Agostino Ardagna, Sabrina Capitani di Vimercati, and Pierangela Samarati. 2011. Privacy Models and Languages: Access Control and Data Handling Policies. In *Digital Privacy*, Jan Camenisch, Ronald Leenes and

- Dieter Sommer (eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 309–329. Retrieved September 27, 2011 from DOI 10.1007/978-3-642-19050-6_11
21. Robert M Arlein, Ben Jai, Markus Jakobsson, Fabian Monrose, and Michael K Reiter. 2000. Privacy-Preserving Global Customization. In *2nd ACM Conference on Electronic Commerce*, 176–184.
 22. Susan Athey, Christian Catalini, and Catherine Tucker. 2017. *The Digital Privacy Paradox: Small Money, Small Costs, Small Talk*. National Bureau of Economic Research. <https://doi.org/10.3386/w23488>
 23. Naveen Farag Awad and M. S. Krishnan. 2006. The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled Online for Personalization. *MIS Quarterly* 30, 1: 13–28.
 24. Paritosh Bahirat, Yangyang He, Abhilash Menon, and Bart Knijnenburg. 2018. A Data-Driven Approach to Developing IoT Privacy-Setting Interfaces. In *23rd International Conference on Intelligent User Interfaces (IUI '18)*, 165–176. <https://doi.org/10.1145/3172944.3172982>
 25. R. Balebako, P. G. Leon, J. Mugan, A. Acquisti, L. F. Cranor, and N. Sadeh. 2011. Nudging users towards privacy on mobile devices. In *CHI 2011 workshop on Persuasion, Influence, Nudge and Coercion Through Mobile Devices*, 23–26. Retrieved December 24, 2012 from <http://www.andrew.cmu.edu/user/jmugan/Publications/chiworkshop.pdf>
 26. Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal Ubiquitous Computing* 15, 7: 679–694. <https://doi.org/10.1007/s00779-010-0346-0>
 27. Bettina Berendt, Oliver Günther, and Sarah Spiekermann. 2005. Privacy in E-Commerce: Stated Preferences vs. Actual Behavior. *Communications of the ACM* 48, 4: 101–106. <https://doi.org/10.1145/1053291.1053295>
 28. Mike Bergmann. 2009. Testing Privacy Awareness. In *The Future of Identity in the Information Society*, Vashek Matyáš, Simone Fischer-Hübner, Daniel Cvrček and Petr Švenda (eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 237–253.
 29. Shlomo Berkovsky, Yaniv Eytani, Tsvi Kuflik, and Francesco Ricci. 2006. Hierarchical Neighborhood Topology for Privacy Enhanced Collaborative Filtering. In *Proceedings of PEP06, CHI 2006 Workshop on Privacy-Enhanced Personalization*, 6–13.
 30. Shlomo Berkovsky, Ronnie Taib, and Dan Conway. 2017. How to Recommend?: User Trust Factors in Movie Recommender Systems. In *Proceedings of the 22Nd International Conference on Intelligent User Interfaces (IUI '17)*, 287–300. <https://doi.org/10.1145/3025171.3025209>
 31. 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>
 32. James R. Bettman, Mary Frances Luce, and John W. Payne. 1998. Constructive Consumer Choice Processes. *Journal of Consumer Research* 25, 3: 187–217. <https://doi.org/10.1086/209535>
 33. Anol Bhattacharjee. 2002. Individual Trust in Online Firms: Scale Development and Initial Test. *Journal of Management Information Systems* 19, 1: 211–241. <https://doi.org/10.1080/07421222.2002.11045715>
 34. Kemal Bicakci, Nart Bedin Atalay, and Hakan Ezgi Kiziloz. 2011. Johnny in internet café: user study and exploration of password autocomplete in web browsers. In *Proceedings of the 7th ACM workshop on Digital identity management*, 33–42. <https://doi.org/10.1145/2046642.2046652>
 35. Wouter Bokhove, Bob Hulsebosch, Bas Van Schoonhoven, Maya Sappelli, and Kees Wouters. 2012. User Privacy in Applications for Well-being and Well-working. In *AMBIENT 2012, The Second International Conference on Ambient Computing, Applications, Services and Technologies*, 53–59. Retrieved December 24, 2012 from http://www.thinkmind.org/index.php?view=article&articleid=ambient_2012_3_10_70014
 36. Dirk Bollen, Bart P. Knijnenburg, Martijn C. Willemsen, and Mark Graus. 2010. Understanding choice overload in recommender systems. In *Proceedings of the fourth ACM conference on Recommender systems*, 63–70. <https://doi.org/10.1145/1864708.1864724>
 37. Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns. *Proceedings on Privacy Enhancing Technologies* 2016, 4: 237–254. <https://doi.org/10.1515/popets-2016-0038>
 38. Keith Bradley, Rachael Rafter, and Barry Smyth. 2000. Case-Based User Profiling for Content Personalisation. In *Adaptive Hypermedia and Adaptive Web-Based Systems (Lecture Notes in Computer Science)*, 62–72. https://doi.org/10.1007/3-540-44595-1_7

39. Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. 2013. Misplaced Confidences: Privacy and the Control Paradox. *Social Psychological and Personality Science* 4, 3: 340–347. <https://doi.org/10.1177/1948550612455931>
40. Pamela Briggs, Brad Simpson, and Antonella De Angeli. 2004. Personalisation and Trust: A Reciprocal Relationship? In *Designing Personalized User Experiences in eCommerce*, Clare-Marie Karat, Jan O. Blom and John Karat (eds.). Kluwer Academic Publishers, Dordrecht, Netherlands, 39–55.
41. Christina L. Brown and Aradhna Krishna. 2004. The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice. *Journal of Consumer Research* 31, 3: 529–539. <https://doi.org/10.1086/425087>
42. Michael J Brzozowski and Daniel M Romero. 2011. Who Should I Follow? Recommending People in Directed Social Networks. In *Fifth International AAAI Conference on Weblogs and Social Media*. Retrieved December 30, 2011 from <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewPaper/2867>
43. Tom Buchanan, Carina Paine, Adam N. Joinson, and Ulf-Dietrich Reips. 2007. Development of measures of online privacy concern and protection for use on the Internet. *Journal of the American Society for Information Science and Technology* 58, 2: 157–165. <https://doi.org/10.1002/asi.20459>
44. Burcu Bulgurcu. 2012. Understanding the information privacy-related perceptions and behaviors of an online social network user. University of British Columbia, Vancouver, BC.
45. Moira Burke, Robert Kraut, and Cameron Marlow. 2011. Social capital on facebook: differentiating uses and users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, 571–580. <https://doi.org/10.1145/1978942.1979023>
46. L Bustos. 2012. Best Practice Gone Bad: 4 Shocking A/B Tests. *GetElastic*. Retrieved January 2, 2013 from <http://www.getelastic.com/best-practice-gone-bad-4-shocking-ab-tests/>
47. Ralf Caers, Tim De Feyter, Marijke De Couck, Talia Stough, Claudia Vigna, and Cind Du Bois. 2013. Facebook: A literature review. *New Media & Society* 15, 6: 982–1002. <https://doi.org/10.1177/1461444813488061>
48. Joseph A. Calandrino, Ann Kilzer, Arvind Narayanan, Edward W. Felten, and Vitaly Shmatikov. 2011. You Might Also Like: Privacy Risks of Collaborative Filtering. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, 231–246. <https://doi.org/10.1109/SP.2011.40>
49. John Canny. 2002. Collaborative Filtering with Privacy. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, 45–57. <https://doi.org/10.1109/SECPR.2002.1004361>
50. Jinwei Cao, Kamile Asli Basoglu, Hong Sheng, and Paul Benjamin Lowry. 2015. A systematic review of social networks research in information systems: Building a foundation for exciting future research. *Communications of the Association for Information Systems* 36: 727–758.
51. Jinwei Cao and Andrea Everard. 2007. Influence of Culture on Attitude Towards Instant Messaging: Balance Between Awareness and Privacy. In *Human-Computer Interaction. Interaction Platforms and Techniques*, Julie Jacko (ed.). Springer Berlin / Heidelberg, 236–240. Retrieved November 25, 2012 from <http://www.springerlink.com/content/576224625811r819/abstract/>
52. Avner Caspi and Paul Gorsky. 2006. Online Deception: Prevalence, Motivation, and Emotion. *CyberPsychology & Behavior* 9: 54–59. <https://doi.org/10.1089/cpb.2006.9.54>
53. Lillian N Cassel and Ursula Wolz. 2001. Client Side Personalization. In *DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries*, 8–12. Retrieved from <http://www.ercim.eu/publication/ws-proceedings/DelNoe02/CasselWolz.pdf>
54. Ann Cavoukian. 2009. *Privacy by Design*. Information and Privacy Commissioner of Ontario, Canada. Retrieved from <http://www.privacybydesign.ca/content/uploads/2010/03/PrivacybyDesignBook.pdf>
55. 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*.
56. R. K Chellappa and R. G Sin. 2005. Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management* 6, 2: 181–202. <https://doi.org/10.1007/s10799-005-5879-y>
57. Jilin Chen, Werner Geyer, Casey Dugan, Michael Muller, and Ido Guy. 2009. Make new friends, but keep the old: recommending people on social networking sites. In *Proceedings of the 27th international conference on Human factors in computing systems*, 201–210. <https://doi.org/10.1145/1518701.1518735>
58. Li Chen and Pearl Pu. 2012. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction* 22, 1–2: 125–150. <https://doi.org/10.1007/s11257-011-9108-6>

59. Xusen Cheng, Shixuan Fu, and Gert-Jan de Vreede. 2017. Understanding trust influencing factors in social media communication: A qualitative study. *International Journal of Information Management* 37, 2: 25–35. <https://doi.org/10.1016/j.ijinfomgt.2016.11.009>
60. Daegon Cho, Soodong Kim, and A. Acquisti. 2012. Empirical analysis of online anonymity and user behaviors: the impact of real name policy. In *2012 45th Hawaii International Conference on System Science*, 3041–3050. <https://doi.org/10.1109/HICSS.2012.241>
61. Hichang Cho, Milagros Rivera-Sánchez, and Sun Sun Lim. 2009. A multinational study on online privacy: global concerns and local responses. *New Media & Society* 11, 3: 395–416. <https://doi.org/10.1177/1461444808101618>
62. Seong Eun Cho and Han Woo Park. 2013. A qualitative analysis of cross-cultural new media research: SNS use in Asia and the West. *Quality & Quantity* 47, 4: 2319–2330. <https://doi.org/10.1007/s11135-011-9658-z>
63. Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. 2013. Nudging people away from privacy-invasive mobile apps through visual framing. In *IFIP Conference on Human-Computer Interaction*, 74–91. https://doi.org/10.1007/978-3-642-40477-1_5
64. Boreum Choi and Inseong Lee. 2017. Trust in open versus closed social media: The relative influence of user- and marketer-generated content in social network services on customer trust. *Telematics and Informatics* 34, 5: 550–559. <https://doi.org/10.1016/j.tele.2016.11.005>
65. Emily Christofides, Amy Muise, and Serge Desmarais. 2009. Information Disclosure and Control on Facebook: Are They Two Sides of the Same Coin or Two Different Processes? *CyberPsychology & Behavior* 12, 3: 341–345. <https://doi.org/10.1089/cpb.2008.0226>
66. Richard Cissé and Sahin Albayrak. 2007. An Agent-based Approach for Privacy-preserving Recommender Systems. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '07)*, 182:1–182:8. <https://doi.org/10.1145/1329125.1329345>
67. Sophie Cockcroft and Saphira Rekker. 2015. The relationship between culture and information privacy policy. *Electronic Markets*: 1–18. <https://doi.org/10.1007/s12525-015-0195-9>
68. 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
69. Sunny Consolvo, Ian Smith, Terra Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. 2005. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 81–90. <https://doi.org/10.1145/1054972.1054985>
70. Consumer Reports. 2012. Facebook & your privacy: Who sees the data you share on the biggest social network? *Consumer Reports*. Retrieved from <http://www.consumerreports.org/cro/magazine/2012/06/facebook-your-privacy>
71. Teresa Correa, Amber Willard Hinsley, and Homero Gil de Zúñiga. 2010. Who interacts on the Web?: The intersection of users’ personality and social media use. *Computers in Human Behavior* 26, 2: 247–253. <https://doi.org/10.1016/j.chb.2009.09.003>
72. Dan Cosley, Shyong K. Lam, Istvan Albert, Joseph A. Konstan, and John Riedl. 2003. Is Seeing Believing?: How Recommender System Interfaces Affect Users’ Opinions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*, 585–592. <https://doi.org/10.1145/642611.642713>
73. Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction* 18, 5: 455–496. <https://doi.org/10.1007/s11257-008-9051-3>
74. Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006. User Interfaces for Privacy Agents. *ACM Transactions on Human-Computer Interactions*. <https://doi.org/10.1145/1165734.1165735>
75. Lorrie Cranor, Marc Langheinrich, and Massimo Marchiori. 2002. *A P3P Preference Exchange Language 1.0 (APPEL1.0)*.
76. J. Cranshaw, J. Mugan, and N. Sadeh. 2011. User-Controllable Learning of Location Privacy Policies with Gaussian Mixture Models. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AI '11)*, 1146–1152. Retrieved February 5, 2013 from <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/viewPDFInterstitial/3785/4052>

77. Paolo Cremonesi, Franca Garzotto, and Roberto Turrin. 2012. Investigating the Persuasion Potential of Recommender Systems from a Quality Perspective: An Empirical Study. *ACM Transactions on Interactive Intelligent Systems* 2, 2: 11:1–11:41. <https://doi.org/10.1145/2209310.2209314>
78. Leon Cremonini and Lorenzo Valeri. 2003. *Benchmarking Security and Trust in Europe and the US*. RAND Europe.
79. Mary J. Culnan. 1993. “How Did They Get My Name?”: An Exploratory Investigation of Consumer Attitudes toward Secondary Information Use. *MIS Quarterly* 17, 3: 341–363. <https://doi.org/10.2307/249775>
80. Christoffer Davidsson and Simon Moritz. 2011. Utilizing Implicit Feedback and Context to Recommend Mobile Applications from First Use. In *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation* (CaRR ’11), 19–22. <https://doi.org/10.1145/1961634.1961639>
81. Amra Delic, Julia Neidhardt, Thuy Ngoc Nguyen, Francesco Ricci, Laurens Rook, Hannes Werthner, and Markus Zanker. 2016. Observing Group Decision Making Processes. In *Proceedings of the 10th ACM Conference on Recommender Systems* (RecSys ’16), 147–150. <https://doi.org/10.1145/2959100.2959168>
82. André Deuker. 2012. Friend-to-Friend Privacy Protection on Social Networking Sites: A Grounded Theory Study. In *AMCIS 2012 Proceedings*. Retrieved from <http://aisel.aisnet.org/amcis2012/proceedings/SocialIssues/5>
83. Tamara Dinev, Massimo Bellotto, Paul Hart, Vincenzo Russo, Ilaria Serra, and Christian Colautti. 2006. Privacy calculus model in e-commerce - a study of Italy and the United States. *European Journal of Information Systems* 15, 4: 389–402. <http://dx.doi.org.janus.lib.tue.nl/10.1057/palgrave.ejis.3000590>
84. Tamara Dinev and Paul Hart. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research* 17, 1: 61–80. <https://doi.org/10.1287/isre.1060.0080>
85. Isaac Dinner, Eric J. Johnson, Daniel G. Goldstein, and Kaiya Liu. 2011. Partitioning Default Effects: Why People Choose Not to Choose. *Journal of Experimental Psychology: Applied* 17, 4: 332–341. Retrieved June 26, 2017 from <https://www.learntechlib.org/p/64714/>
86. Cailing Dong, Hongxia Jin, and Bart P. Knijnenburg. 2015. Predicting Privacy Behavior on Online Social Networks. In *Ninth International AAAI Conference on Web and Social Media*, 91–100. Retrieved May 22, 2015 from <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10554>
87. Cailing Dong, Hongxia Jin, and Bart P. Knijnenburg. 2016. PPM: A Privacy Prediction Model for Online Social Networks. In *Proceedings of The International Conference on Social Informatics* (SocInfo ’16), 400–420. https://doi.org/10.1007/978-3-319-47874-6_28
88. Emily Dreyfuss. 2018. Was It Ethical for Dropbox to Share Customer Data with Scientists? *Wired*. Retrieved July 29, 2018 from <https://www.wired.com/story/dropbox-sharing-data-study-ethics/>
89. Cynthia Dwork and Moni Naor. 2008. On the Difficulties of Disclosure Prevention in Statistical Databases or The Case for Differential Privacy. *Journal of Privacy and Confidentiality* 2, 1. Retrieved from <http://repository.cmu.edu/jpc/vol2/iss1/8>
90. S. Egelman, J. Tsai, L. F. Cranor, and A. Acquisti. 2009. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proceedings of the 27th international conference on Human factors in computing systems*, 319–328.
91. Michael D. Ekstrand, Rezvan Joshaghani, and Hoda Mehrpouyan. 2018. Privacy for All: Ensuring Fair and Equitable Privacy Protections. In *Conference on Fairness, Accountability and Transparency*, 35–47. Retrieved July 30, 2018 from <http://proceedings.mlr.press/v81/ekstrand18a.html>
92. Michael D. Ekstrand and Martijn C. Willemsen. 2016. Behaviorism is Not Enough: Better Recommendations Through Listening to Users. In *Proceedings of the 10th ACM Conference on Recommender Systems* (RecSys ’16), 221–224. <https://doi.org/10.1145/2959100.2959179>
93. Charles Elkan. 2001. The Foundations of Cost-sensitive Learning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2* (IJCAI’01), 973–978. <https://doi.org/10.29012/jpc.v2i1.585>
94. Larry E. Ellis, Sandra G. Nunn, and John T. Avella. 2016. Digital Badges and Micro-credentials: Historical Overview, Motivational Aspects, Issues, and Challenges. In *Foundation of Digital Badges and Micro-Credentials*. Springer, Cham, 3–21. https://doi.org/10.1007/978-3-319-15425-1_1
95. EU. 2012. *Proposal for a directive of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data by competent authorities for the purposes of prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties,*

- and the free movement of such data*. Retrieved from <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:52012PC0010:en:NOT>
96. 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>
 97. Alexander Felfernig. 2007. Knowledge-Based Recommender Technologies for Marketing and Sales. *International Journal of Pattern Recognition and Artificial Intelligence* 21, 2: 333–354. <https://doi.org/10.1142/S0218001407005417>
 98. Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android Permissions: User Attention, Comprehension, and Behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS '12)*, 3:1–3:14. <https://doi.org/10.1145/2335356.2335360>
 99. Gavan J. Fitzsimons and Donald R. Lehmann. 2004. Reactance to Recommendations: When Unsolicited Advice Yields Contrary Responses. *Marketing Science* 23, 1: 82–94. <https://doi.org/10.1287/mksc.1030.0033>
 100. B Fogg. 2003. *Persuasive technology : using computers to change what we think and do*. Morgan Kaufmann Publishers, Amsterdam.
 101. 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*.
 102. Gerhard Friedrich and Markus Zanker. 2011. A Taxonomy for Generating Explanations in Recommender Systems. *AI Magazine* 32, 3: 90–98. <https://doi.org/10.1609/aimag.v32i3.2365>
 103. FTC. 2000. *Privacy Online: Fair Information Practices in the Electronic Marketplace. A Report to Congress*. Federal Trade Commission. Retrieved from www.ftc.gov/reports/privacy2000/privacy2000.pdf
 104. Evelien van de Garde-Perik, Panos Markopoulos, Boris de Ruyter, Berry Eggen, and Wijnand Ijsselstein. 2008. Investigating Privacy Attitudes and Behavior in Relation to Personalization. *Social Science Computer Review* 26, 1: 20–43. <https://doi.org/10.1177/0894439307307682>
 105. James Gardner. 2012. 12 Surprising A/B Test Results to Stop You Making Assumptions. *Unbounce*. Retrieved February 16, 2014 from <http://unbounce.com/a-b-testing/shocking-results/>
 106. Amit Garg. 2015. Multi-device learning. *Training & Development* 42, 4: 10. Retrieved February 5, 2017 from <http://search.informit.com.au/documentSummary;dn=412351784359697;res=IELBUS>
 107. Fatih Gedikli, Dietmar Jannach, and Mouzhi Ge. 2014. How should I explain? A comparison of different explanation types for recommender systems. *International Journal of Human-Computer Studies* 72, 4: 367–382. <https://doi.org/10.1016/j.ijhcs.2013.12.007>
 108. Cristina Gena, Roberto Brogi, Federica Cena, and Fabiana Vernerio. 2011. The Impact of Rating Scales on User's Rating Behavior. In *User Modeling, Adaption and Personalization*, Joseph A. Konstan, Ricardo Conejo, José L. Marzo and Nuria Oliver (eds.). Springer Berlin Heidelberg, 123–134.
 109. Andrea Girardello and Florian Michahelles. 2010. AppAware: Which Mobile Applications Are Hot? In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '10)*, 431–434. <https://doi.org/10.1145/1851600.1851698>
 110. Ulrike Gretzel and Daniel R. Fesenmaier. 2006. Persuasion in Recommender Systems. *International Journal of Electronic Commerce* 11, 2: 81–100. <https://doi.org/10.2753/JEC1086-4415110204>
 111. Ido Guy, Inbal Ronen, and Eric Wilcox. 2009. Do you know?: recommending people to invite into your social network. In *Proceedings of the 14th international conference on Intelligent user interfaces (IUI '09)*, 77–86. <https://doi.org/10.1145/1502650.1502664>
 112. John Hagel and Jeffrey F. Rayport. 1999. The Coming Battle for Customer Information. In *Creating value in the network economy*. Harvard Business School Press, Boston, MA, 159–171. Retrieved May 26, 2014 from <http://hbr.org/1997/01/the-coming-battle-for-customer-information/ar/3>
 113. Il-Horn Hann, Kai-Lung Hui, Sang-Yong Lee, and Ivan Png. 2007. Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach. *Journal of Management Information Systems* 24, 2: 13–42. <https://doi.org/10.2753/MIS0742-1222240202>
 114. David J. Hardisty, Eric J. Johnson, and Elke U. Weber. 2010. A Dirty Word or a Dirty World?: Attribute Framing, Political Affiliation, and Query Theory. *Psychological Science* 21, 1: 86–92. <https://doi.org/10.1177/0956797609355572>

115. Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *27th USENIX Security Symposium (USENIX Security 18)*. Retrieved from <https://www.usenix.org/conference/usenixsecurity18/presentation/harkous>
116. Hamza Harkous, Rameez Rahman, and Karl Aberer. 2016. Data-Driven Privacy Indicators. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Retrieved from <https://www.usenix.org/conference/soups2016/workshop-program/wpi/presentation/harkous>
117. Harris. 2001. *Privacy Notices Research: Final Results*. Harris Interactive, Inc. Retrieved from <http://www.bbbonline.org/UnderstandingPrivacy/library/datasum.pdf>
118. 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>
119. 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.
120. Michael M. Harris, Greet Van Hove, and Filip Lievens. 2003. Privacy and Attitudes Towards Internet-Based Selection Systems: A Cross-Cultural Comparison. *International Journal of Selection and Assessment* 11, 2–3: 230–236. <https://doi.org/10.1111/1468-2389.00246>
121. Gerald Häubl and Valerie Trifts. 2000. Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids. *Marketing Science* 19, 1: 4–21. Retrieved March 2, 2014 from <http://www.jstor.org/stable/193256>
122. Benjamin Heitmann, James G Kim, Alexandre Passant, Conor Hayes, and Hong-Gee Kim. 2010. An architecture for privacy-enabled user profile portability on the web of data. In *HetRec’10 Proceedings of the 1st International Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec ’10)*, 16–23. <https://doi.org/10.1145/1869446.1869449>
123. Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. 2000. Explaining collaborative filtering recommendations. In *Proc. of the 2000 ACM conference on Computer supported cooperative work*, 241–250. <https://doi.org/10.1145/358916.358995>
124. Ron Hirschrung, Eran Toch, Frank Bolton, and Oded Maimon. 2016. A methodology for estimating the value of privacy in information disclosure systems. *Computers in Human Behavior* 61: 443–453. <https://doi.org/10.1016/j.chb.2016.03.033>
125. Shuk Ying Ho and Kar Tam. 2006. Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes. *MIS Quarterly* 30, 4: 865–890.
126. 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
127. Geert Hofstede. 1980. Culture and Organizations. *International Studies of Management & Organization* 10, 4: 15–41. Retrieved May 25, 2015 from <http://www.jstor.org/stable/40396875>
128. 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>
129. Mariea Grubbs Hoy and George Milne. 2010. Gender Differences in Privacy-Related Measures for Young Adult Facebook Users. *Journal of Interactive Advertising* 10, 2: 28–45. <https://doi.org/10.1080/15252019.2010.10722168>
130. Chiung-wen Hsu. 2006. Privacy concerns, privacy practices and web site categories: Toward a situational paradigm. *Online Information Review* 30, 5: 569–586. <https://doi.org/10.1108/14684520610706433>
131. 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>
132. Kai-Lung Hui, Bernard C. Y. Tan, and Chyan-Yee Goh. 2006. Online information disclosure: Motivators and measurements. *ACM Transactions on Internet Technology* 6, 4: 415–441. <https://doi.org/10.1145/1183463.1183467>
133. Kai-Lung Hui, Hock Hai Teo, and Sang-Yong Tom Lee. 2007. The Value of Privacy Assurance: An Exploratory Field Experiment. *MIS Quarterly* 31, 1: 19–33.
134. David Hunter and Nicholas Evans. 2016. Facebook emotional contagion experiment controversy. *Research Ethics* 12, 1: 2–3. <https://doi.org/10.1177/1747016115626341>

135. Internet Society. 2012. *Global Internet User Survey 2012*. Internet Society. Retrieved from <https://www.internetsociety.org/internet/global-internet-user-survey-2012>
136. Qatrunnada Ismail, Tousif Ahmed, Apu Kapadia, and Michael K. Reiter. 2015. Crowdsourced Exploration of Security Configurations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 467–476. <https://doi.org/10.1145/2702123.2702370>
137. Johnson Iyilade. 2013. Enforcing Privacy in Secondary User Information Sharing and Usage. In *User Modeling, Adaptation, and Personalization*, Sandra Carberry, Stephan Weibelzahl, Alessandro Micarelli and Giovanni Semeraro (eds.). Springer Berlin Heidelberg, 396–400.
138. Anthony Jameson, Martijn C. Willemsen, Alexander Felfernig, Marco de Gemmis, Pasquale Lops, Giovanni Semeraro, and Li Chen. 2015. Human Decision Making and Recommender Systems. In *Recommender Systems Handbook*. 611–648.
139. Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. 2010. *Recommender Systems: An Introduction*. Cambridge University Press, New York.
140. Lukasz Jedrzejczyk, Blaine A. Price, Arosha K. Bandara, and Bashar Nuseibeh. 2010. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, 14:1-14:12. <https://doi.org/10.1145/1837110.1837129>
141. Carlos Jensen, Colin Potts, and Christian Jensen. 2005. Privacy Practices of Internet Users: Self-Reports versus Observed Behavior. *International Journal of Human-Computer Studies* 63, 1–2: 203–227. <https://doi.org/10.1016/j.ijhcs.2005.04.019>
142. Leslie K. John, Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of consumer research* 37, 5: 858–873. <https://doi.org/10.1086/656423>
143. David W. Johnson. 1994. *Learning Together and Alone. Cooperative, Competitive, and Individualistic Learning. Fourth Edition*. Allyn and Bacon, 160 Gould Street, Needham Heights, MA 02194. Retrieved February 1, 2017 from <https://eric.ed.gov/?id=ED369778>
144. 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>
145. Maritza Johnson, Serge Egelman, and Steven M. Bellovin. 2012. Facebook and privacy: it's complicated. In *Proceedings of the 8th Symposium on Usable Privacy and Security*, 9:1-9:15. <https://doi.org/10.1145/2335356.2335369>
146. Adam N. Joinson, Carina Paine, Tom Buchanan, and Ulf-Dietrich Reips. 2008. Measuring self-disclosure online: Blurring and non-response to sensitive items in web-based surveys. *Computers in Human Behavior* 24, 5: 2158–2171. <https://doi.org/10.1016/j.chb.2007.10.005>
147. Adam N. Joinson, Ulf-Dietrich Reips, Tom Buchanan, and Carina B. Paine Schofield. 2010. Privacy, Trust, and Self-Disclosure Online. *Human-Computer Interaction* 25, 1: 1–24. <https://doi.org/10.1080/07370020903586662>
148. Ari Juels. 2001. Targeted Advertising ... and Privacy Too. In *Topics in Cryptology — CT-RSA 2001*, David Naccache (ed.). Springer, Berlin/Heidelberg, 408–424.
149. Iris A. Junglas, Norman A. Johnson, and Christiane Spitzmüller. 2008. Personality traits and concern for privacy: an empirical study in the context of location-based services. *European Journal of Information Systems* 17, 4: 387–402. <https://doi.org/10.1057/ejis.2008.29>
150. Sanjay Kairam, Mike Brzozowski, David Huffaker, and Ed Chi. 2012. Talking in circles: Selective Sharing in Google+. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1065–1074. <https://doi.org/10.1145/2207676.2208552>
151. Pamela Karr-Wisniewski, David C. Wilson, and Heather Richter-Lipford. 2011. A New Social Order: Mechanisms for Social Network Site Boundary Regulation. In *AMCIS 2011 Proceedings*, Paper 101. Retrieved from http://aisel.aisnet.org/amcis2011_submissions/101
152. Judy Kay and Bob Kummerfeld. 2013. Creating personalized systems that people can scrutinize and control: Drivers, principles and experience. *ACM Transactions on Interactive Intelligent Systems* 2, 4: 24:1–24:42. <https://doi.org/10.1145/2395123.2395129>
153. Anthony Kaye. 1992. Learning Together Apart. In *Collaborative Learning Through Computer Conferencing*, Anthony R. Kaye (ed.). Springer Berlin Heidelberg, 1–24. https://doi.org/10.1007/978-3-642-77684-7_1

154. Flavius Kehr, Tobias Kowatsch, Daniel Wentzel, and Elgar Fleisch. 2015. Thinking Styles and Privacy Decisions: Need for Cognition, Faith into Intuition, and the Privacy Calculus. In *12th International Conference on Wirtschaftsinformatik*.
155. Flavius Kehr, Daniel Wentzel, and Peter Mayer. 2013. Rethinking the Privacy Calculus: On the Role of Dispositional Factors and Affect. In *ICIS 2013 Proceedings*.
156. Patrick Gage Kelley, Robin Brewer, Yael Mayer, Lorrie Faith Cranor, and Norman Sadeh. 2011. An Investigation into Facebook Friend Grouping. In *INTERACT*, Pedro Campos, Nicholas Graham, Joaquim Jorge, Nuno Nunes, Philippe Palanque and Marco Winckler (eds.). Springer Heidelberg, Lisbon, Portugal, 216–233.
157. Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. 2010. Standardizing privacy notices: an online study of the nutrition label approach. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI '10)*, 1573–1582. <https://doi.org/10.1145/1753326.1753561>
158. Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy As Part of the App Decision-making Process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 3393–3402. <https://doi.org/10.1145/2470654.2466466>
159. Patrick Gage Kelley, Paul Hanks Drielsma, Norman Sadeh, and Lorrie Faith Cranor. 2008. User-controllable Learning of Security and Privacy Policies. In *Proceedings of the 1st ACM Workshop on Workshop on AISEC (AISEC '08)*, 11–18. <https://doi.org/10.1145/1456377.1456380>
160. Ashraf Khalil and Kay Connelly. 2006. Context-aware telephony: privacy preferences and sharing patterns. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, 469–478. <https://doi.org/10.1145/1180875.1180947>
161. Kyung-Hee Kim and Haejin Yun. 2007. Cying for me, Cying for us: Relational dialectics in a Korean social network site. *Journal of Computer-Mediated Communication* 13, 1: 298–318. <https://doi.org/10.1111/j.1083-6101.2007.00397.x>
162. Yoojung Kim, Dongyoung Sohn, and Sejung Marina Choi. 2011. Cultural difference in motivations for using social network sites: A comparative study of American and Korean college students. *Computers in Human Behavior* 27, 1: 365–372. <https://doi.org/10.1016/j.chb.2010.08.015>
163. Daniel Kluver, Tien T. Nguyen, Michael Ekstrand, Shilad Sen, and John Riedl. 2012. How Many Bits Per Rating? In *Proceedings of the Sixth ACM Conference on Recommender Systems (RecSys '12)*, 99–106. <https://doi.org/10.1145/2365952.2365974>
164. Herschel Knapp and Stuart A. Kirk. 2003. Using pencil and paper, Internet and touch-tone phones for self-administered surveys: does methodology matter? *Computers in Human Behavior* 19, 1: 117–134. [https://doi.org/10.1016/S0747-5632\(02\)00008-0](https://doi.org/10.1016/S0747-5632(02)00008-0)
165. Lionel Knight. 2010. Social experiment:online privacy vs. personalization paradox. Retrieved November 1, 2010 from http://www.upshot.net/wp-content/uploads/2010/08/Social_Experiment.pdf
166. B. P. Knijnenburg. 2015. A user-tailored approach to privacy decision support. University of California, Irvine, Irvine, CA. Retrieved from <http://search.proquest.com/docview/1725139739/abstract>
167. B. P. Knijnenburg. 2017. Privacy? I Can't Even! Making a Case for User-Tailored Privacy. *IEEE Security Privacy* 15, 4: 62–67. <https://doi.org/10.1109/MSP.2017.3151331>
168. 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*.
169. 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>
170. B. P. Knijnenburg and Alfred Kobsa. 2013. Helping users with information disclosure decisions: potential for adaptation. In *Proceedings of the 2013 ACM international conference on Intelligent User Interfaces*, 407–416. <https://doi.org/10.1145/2449396.2449448>
171. B. P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Preference-based location sharing: are more privacy options really better? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 2667–2676. <https://doi.org/10.1145/2470654.2481369>
172. 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>

173. 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*.
174. Bart. P. Knijnenburg. 2009. Adaptive advice: adapting a recommender system for energy-saving behaviors to personal differences in decision-making. Eindhoven University of Technology, Eindhoven, Netherlands. Retrieved from <http://alexandria.tue.nl/extra2/afstversl/tm/Knijnenburg%202009.pdf>
175. Bart P Knijnenburg. 2013. Simplifying Privacy Decisions: Towards Interactive and Adaptive Solutions. In *Proceedings of the Recsys 2013 Workshop on Human Decision Making in Recommender Systems*, 40–41.
176. Bart P. Knijnenburg, Svetlin Bostandjiev, John O'Donovan, and Alfred Kobsa. 2012. Inspectability and control in social recommenders. In *Proceedings of the sixth ACM conference on Recommender systems (RecSys '12)*, 43–50. <https://doi.org/10.1145/2365952.2365966>
177. Bart P Knijnenburg and Hongxia Jin. 2013. The Persuasive Effect of Privacy Recommendations. In *Twelfth Annual Workshop on HCI Research in MIS (SigHCI '13)*. Retrieved from <http://aisel.aisnet.org/sighci2013/16>
178. Bart P. Knijnenburg, A. Kobsa, and Hongxia Jin. 2014. Segmenting the Recipients of Personal Information.
179. Bart P. Knijnenburg, Alfred Kobsa, and Martijn C. Willemsen. 2016. *Taking Control of Household IoT Device Privacy - A White Paper for the Sociotechnical Cybersecurity Workshop*. The Computing Community Consortium, La Jolla, CA. Retrieved from <https://www.usabart.nl/portfolio/paper-IoT2016.html>
180. Bart P. Knijnenburg, Elaine M. Raybourn, David Cherry, Daricia Wilkinson, Saadhika Sivakumar, and Henry Sloan. 2017. Death to the Privacy Calculus? In *Proceedings of the 2017 Networked Privacy Workshop at CSCW*. Retrieved March 8, 2017 from <https://papers.ssrn.com/abstract=2923806>
181. Bart P. Knijnenburg, Niels J.M. Reijmer, and Martijn C. Willemsen. 2011. Each to his own: how different users call for different interaction methods in recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, 141–148. <https://doi.org/10.1145/2043932.2043960>
182. Bart P. Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender Systems for Self-Actualization. In *Proceedings of the 10th ACM Conference on Recommender Systems*, 11–14. <https://doi.org/10.1145/2959100.2959189>
183. Bart P. Knijnenburg and Martijn C. Willemsen. 2009. Understanding the Effect of Adaptive Preference Elicitation Methods on User Satisfaction of a Recommender System. In *Proceedings of the Third ACM Conference on Recommender Systems (RecSys '09)*, 381–384. <https://doi.org/10.1145/1639714.1639793>
184. Bart P. Knijnenburg and Martijn C. Willemsen. 2010. The effect of preference elicitation methods on the user experience of a recommender system. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, 3457–3462. <https://doi.org/10.1145/1753846.1754001>
185. Bart P. Knijnenburg and Martijn C. Willemsen. 2015. Evaluating Recommender Systems with User Experiments. In *Recommender Systems Handbook* (2nd ed.), Francesco Ricci, Lior Rokach and Bracha Shapira (eds.). Springer US, New York, NY, 309–352. https://doi.org/10.1007/978-1-4899-7637-6_9
186. Bart P. Knijnenburg, Martijn C. Willemsen, and R. Broeders. 2014. Smart Sustainability through System Satisfaction: Tailored Preference Elicitation for Energy-saving Recommenders. In *AMCIS 2014 proceedings*.
187. A. Kobsa and M. Teltzrow. 2005. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. In *Privacy Enhancing Technologies: Revised Selected Papers of the 4th International Workshop: PET '04*, David Martin and Andrei Serjantov (eds.). Springer Berlin Heidelberg, 329–343.
188. Alfred Kobsa. 2001. Tailoring Privacy to Users' Needs 1. In *User Modeling 2001*, 301–313. https://doi.org/10.1007/3-540-44566-8_52
189. Alfred Kobsa. 2007. Privacy-Enhanced Personalization. *Communications of the ACM* 50, 8: 24–33.
190. Alfred Kobsa, Hichang Cho, and Bart P. Knijnenburg. 2016. The Effect of Personalization Provider Characteristics on Privacy Attitudes and Behaviors: An Elaboration Likelihood Model Approach. *Journal of the Association for Information Science and Technology*. <https://doi.org/10.1002/asi.23629>
191. Alfred Kobsa, Bart P. Knijnenburg, and Benjamin Livshits. 2014. Let's Do It at My Place Instead?: Attitudinal and Behavioral Study of Privacy in Client-side Personalization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, 81–90. <https://doi.org/10.1145/2556288.2557102>
192. Alfred Kobsa and Jörg Schreck. 2003. Privacy through pseudonymity in user-adaptive systems. *ACM Transactions on Internet Technology* 3, 2: 149–183. <https://doi.org/10.1145/767193.767196>
193. Jan Kolter and Günther Pernul. 2009. Generating User-Understandable Privacy Preferences. In *Conf. on Availability, Reliability and Security*, 299–306. <https://doi.org/10.1109/ARES.2009.89>

194. S. Y.X Komiak and I. Benbasat. 2006. The effects of personalization and familiarity on trust and adoption of recommendation agents. *Mis Quarterly* 30, 4: 941–960.
195. Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8: 30–37. <https://doi.org/10.1109/MC.2009.263>
196. Melinda L. Korzaan and Katherine T. Boswell. 2008. The Influence of Personality Traits and Information Privacy Concerns on Behavioral Intentions. *Journal of Computer Information Systems* 48, 4: 15–24. <https://doi.org/10.1080/08874417.2008.11646031>
197. Takashi Koshimizu, Tomoji Toriyama, and Noboru Babaguchi. 2006. Factors on the sense of privacy in video surveillance. In *Proceedings of the 3rd ACM workshop on Continuous archival and retrieval of personal experiences* (CARPE '06), 35–44. <https://doi.org/10.1145/1178657.1178665>
198. Hanna Krasnova, Thomas Hildebrand, and Oliver Guenther. 2009. Investigating the Value of Privacy in Online Social Networks: Conjoint Analysis. In *ICIS 2009 Proceedings*. Retrieved from <http://aisel.aisnet.org/icis2009/173>
199. Yee-Lin Lai and Kai-Lung Hui. 2004. Opting-in or opting-out on the Internet: Does it Really Matter? In *ICIS 2004: Twenty-Fifth International Conference on Information Systems*, 781–792. Retrieved from <http://aisel.aisnet.org/icis2004/63>
200. 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 CPR Conference on Computer Personnel Research: Forty Four Years of Computer Personnel Research: Achievements, Challenges & the Future* (SIGMIS CPR '06), 253–263. <https://doi.org/10.1145/1125170.1125230>
201. 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
202. Jaron Lanier. 2010. *You Are Not a Gadget: A Manifesto*. Thorndike Press, Waterville, Me.
203. Elsy Lao and Alfred Kobsa. 2005. *Privacy Attitudes of Internet Users in the U.S. and Europe*. University of California, Irvine.
204. Robert Larose and Nora J Rifon. 2007. Promoting i-Safety: Effects of Privacy Warnings and Privacy Seals on Risk Assessment and Online Privacy Behavior. *Journal of Consumer Affairs* 41, 1: 127–149. <https://doi.org/10.1111/j.1745-6606.2006.00071.x>
205. Robert S Laufer, Harold M Proshansky, and Maxine Wolfe. 1973. Some Analytic Dimensions of Privacy. In *Proceedings of the lund conference on architectural psychology*, 353–372.
206. Robert S. Laufer and Maxine Wolfe. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of social issues* 33, 3: 22–42. <https://doi.org/10.1111/j.1540-4560.1977.tb01880.x>
207. 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>
208. Scott Lederer, Jennifer Mankoff, and Anind K. Dey. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 724–725. <https://doi.org/10.1145/765891.765952>
209. H. Lee and A. Kobsa. 2016. Understanding user privacy in Internet of Things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, 407–412. <https://doi.org/10.1109/WF-IoT.2016.7845392>
210. H. Lee and A. Kobsa. 2017. Privacy preference modeling and prediction in a simulated campuswide IoT environment. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 276–285. <https://doi.org/10.1109/PERCOM.2017.7917874>
211. H. Li, R. Sarathy, and H. Xu. 2010. Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems* 51, 1: 62–71.
212. Yao Li, Alfred Kobsa, Bart P Knijnenburg, and MH Carolyn Nguyen. 2017. Cross-Cultural Privacy Prediction. *Proceedings on Privacy Enhancing Technologies* 2: 93–112.
213. Yuan Li. 2014. The impact of disposition to privacy, website reputation and website familiarity on information privacy concerns. *Decision Support Systems* 57: 343–354. <https://doi.org/10.1016/j.dss.2013.09.018>
214. Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I. Hong. 2014. Modeling Users' Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings. In *Symposium on Usable Privacy and Security* (SOUPS '14).

215. Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhiemedi, 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 '16)*, 27–41.
216. 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>
217. Yabing Liu, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing Facebook privacy settings: user expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, 61–70. <https://doi.org/10.1145/2068816.2068823>
218. Wainer Lusoli, Margherita Bacigalupo, Francisco Lupiáñez-Villanueva, Norberto Andrade, Shara Monteleone, and Ioannis Maghiros. 2012. *Pan-European Survey of Practices, Attitudes and Policy Preferences as Regards Personal Identity Data Management*. Social Science Research Network, Rochester, NY. Retrieved February 26, 2013 from <http://papers.ssrn.com/abstract=2086579>
219. Kevin Mabley. 2000. *Privacy vs. Personalization: Part III*. Cyber Dialogue, Inc. Retrieved from <http://www.egov.vic.gov.au/pdfs/wp-2000-privacy3.pdf>
220. Ashwin Machanavajjhala, Aleksandra Korolova, and Atish Das Sarma. 2011. Personalized Social Recommendations: Accurate or Private. *Proceedings of the VLDB Endowment* 4, 7: 440–450. <https://doi.org/10.14778/1988776.1988780>
221. Mary Madden. 2012. *Privacy management on social media sites*. Pew Internet & American Life Project, Pew Research Center, Washington, DC. Retrieved from <http://www.pewinternet.org/2012/02/24/privacy-management-on-social-media-sites/>
222. M. Madejski, M. Johnson, and S.M. Bellovin. 2012. A study of privacy settings errors in an online social network. In *Fourth International Workshop on Security and Social Networking (SECSOC '12)*, 340–345. <https://doi.org/10.1109/PerComW.2012.6197507>
223. Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research* 15, 4: 336–355. <https://doi.org/10.1287/isre.1040.0032>
224. Antony S.R. Manstead. 1996. Attitudes and Behaviour. In *Applied Social Psychology*. SAGE Publications Ltd, London, 3–29. <https://doi.org/10.4135/9781446250556>
225. Judith Masthoff. 2004. Group Modeling: Selecting a Sequence of Television Items to Suit a Group of Viewers. *User Modeling and User-Adapted Interaction* 14, 1: 37–85. <https://doi.org/10.1023/B:USER.0000010138.79319.fd>
226. Judith Masthoff. 2015. Group Recommender Systems: Aggregation, Satisfaction and Group Attributes. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach and Bracha Shapira (eds.). Springer US, Boston, MA, 743–776. https://doi.org/10.1007/978-1-4899-7637-6_22
227. Lorraine McGinty and Barry Smyth. 2006. Adaptive Selection: An Analysis of Critiquing and Preference-Based Feedback in Conversational Recommender Systems. *International Journal of Electronic Commerce* 11, 2: 35–57. <https://doi.org/10.2753/JEC1086-4415110202>
228. 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>
229. D. Harrison McKnight, Vivek Choudhury, and Charles Kacmar. 2002. Developing and Validating Trust Measures for e-Commerce: An Integrative Typology. *Information Systems Research* 13, 3: 334–359. <https://doi.org/10.1287/isre.13.3.334.81>
230. Frank McSherry and Ilya Mironov. 2009. Differentially Private Recommender Systems: Building Privacy into the Net. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '09)*, 627–636. <https://doi.org/10.1145/1557019.1557090>
231. Gustavo S. Mesch. 2012. Is online trust and trust in social institutions associated with online disclosure of identifiable information online? *Computers in Human Behavior* 28, 4: 1471–1477. <https://doi.org/10.1016/j.chb.2012.03.010>
232. M. J Metzger. 2007. Communication Privacy Management in Electronic Commerce. *Journal of Computer-Mediated Communication* 12, 2: 335–361. <https://doi.org/10.1111/j.1083-6101.2007.00328.x>

233. Miriam J Metzger. 2004. Privacy, Trust, and Disclosure: Exploring Barriers to Electronic Commerce. *Journal of Computer-Mediated Communication* 9, 4. <https://doi.org/10.1111/j.1083-6101.2004.tb00292.x>
234. Miriam J Metzger. 2006. Effects of Site, Vendor, and Consumer Characteristics on Web Site Trust and Disclosure. *Communication Research* 33, 3: 155–179.
235. Sandra J Milberg, Sandra J Burke, H. Jeff Smith, and Ernest A Kallman. 1995. Values, Personal Information, Privacy and Regulatory Approaches. *Communications of the ACM* 38, 12: 65–74. <https://doi.org/10.1145/219663.219683>
236. George R Milne and Mary J Culnan. 2004. Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don't Read) Online Privacy Notices. *Journal of Interactive Marketing* 18, 3: 15–29.
237. Caroline Lancelot Miltgen and Dominique Peyrat-Guillard. 2014. Cultural and generational influences on privacy concerns: a qualitative study in seven European countries. *European Journal of Information Systems* 23, 2: 103–125. <https://doi.org/10.1057/ejis.2013.17>
238. Deirdre Mulligan and Jennifer King. 2012. Bridging the Gap Between Privacy and Design. *University of Pennsylvania Journal of Constitutional Law* 14, 4: 989. Retrieved from <http://scholarship.law.upenn.edu/jcl/vol14/iss4/4>
239. Deirdre Mulligan and Ari Schwartz. 2000. Your Place or Mine?: Privacy Concerns and Solutions for Server and Client-Side Storage of Personal Information. In *Tenth conference on Computers, Freedom and Privacy*, 81–84.
240. Moses Namara, Henry Sloan, Priyanka Jaiswal, and Bart P. Knijnenburg. 2018. The Potential for User-Tailored Privacy on Facebook. In *IEEE Symposium on Privacy-Aware Computing*.
241. Aravind Narayanan and Vitali Shmatikov. 2008. Robust De-anonymization of Large Sparse Datasets. In *2008 IEEE Symposium on Security and Privacy*, 111–125. <https://doi.org/10.1109/SP.2008.33>
242. Thuy Ngoc Nguyen and Francesco Ricci. 2017. Combining Long-term and Discussion-generated Preferences in Group Recommendations. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*, 377–378. <https://doi.org/10.1145/3079628.3079645>
243. Valeria Nikolaenko, Stratis Ioannidis, Udi Weinsberg, Marc Joye, Nina Taft, and Dan Boneh. 2013. Privacy-preserving Matrix Factorization. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (CCS '13)*, 801–812. <https://doi.org/10.1145/2508859.2516751>
244. Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Washington Law Review* 79: 119–157.
245. Helen Nissenbaum. 2011. A Contextual Approach to Privacy Online. *Daedalus* 140, 4: 32–48. https://doi.org/10.1162/DAED_a_00113
246. Helen Fay. Nissenbaum. 2009. *Privacy in context : technology, policy, and the integrity of social life*. Stanford Law Books, Stanford, CA.
247. Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs* 41, 1: 100–126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>
248. Nadia Olivero and Peter Lunt. 2004. Privacy versus willingness to disclose in e-commerce exchanges: The effect of risk awareness on the relative role of trust and control. *Journal of Economic Psychology* 25, 2: 243–262. [https://doi.org/10.1016/S0167-4870\(02\)00172-1](https://doi.org/10.1016/S0167-4870(02)00172-1)
249. Judith S. Olson, Jonathan Grudin, and Eric Horvitz. 2004. *Toward Understanding Preferences for Sharing and Privacy*. Microsoft Research, Redmond, WA.
250. Judith S. Olson, Jonathan Grudin, and Eric Horvitz. 2005. A study of preferences for sharing and privacy. In *CHI '05 Extended Abstracts*, 1985–1988. <https://doi.org/10.1145/1056808.1057073>
251. Will Oremus. 2014. Facebook's Privacy Dinosaur Wants to Make Sure You're Not Oversharing. *Slate*. Retrieved May 6, 2014 from http://www.slate.com/blogs/future_tense/2014/03/25/facebook_privacy_dinosaur_privacy_checkups_take_aim_at_oversharing.html
252. Rita Orji, Regan L. Mandryk, Julita Vassileva, and Kathrin M. Gerling. 2013. Tailoring Persuasive Health Games to Gamer Type. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 2467–2476. <https://doi.org/10.1145/2470654.2481341>
253. Babajide Osatuyi. 2015. Personality Traits and Information Privacy Concern on Social Media Platforms. *Journal of Computer Information Systems* 55, 4: 11–19. <https://doi.org/10.1080/08874417.2015.11645782>
254. Xinru Page, Bart P. Knijnenburg, and Alfred Kobsa. 2013. FYI: communication style preferences underlie differences in location-sharing adoption and usage. In *Proceedings of the 2013 ACM international joint*

- conference on Pervasive and ubiquitous computing*, 153–162. Retrieved March 25, 2017 from <http://dl.acm.org/citation.cfm?id=2493487>
255. Xinru Page, Alfred Kobsa, and Bart P. Knijnenburg. 2012. Don't Disturb My Circles! Boundary Preservation Is at the Center of Location-Sharing Concerns. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, 266–273. Retrieved June 24, 2012 from <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM12/paper/view/4679>
 256. Peiyu Pai and Hsien-Tung Tsai. 2016. Reciprocity norms and information-sharing behavior in online consumption communities: An empirical investigation of antecedents and moderators. *Information & Management* 53, 1: 38–52. <https://doi.org/10.1016/j.im.2015.08.002>
 257. 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>
 258. Yue Pan and George M. Zinkhan. 2006. Exploring the impact of online privacy disclosures on consumer trust. *Journal of Retailing* 82, 4: 331–338. <https://doi.org/10.1016/j.jretai.2006.08.006>
 259. Alexandros Paramythi, Stephan Weibelzahl, and Judith Masthoff. 2010. Layered evaluation of interactive adaptive systems: framework and formative methods. *User Modeling and User-Adapted Interaction* 20, 5: 383–453. <https://doi.org/10.1007/s11257-010-9082-4>
 260. Sameer Patil and Alfred Kobsa. 2005. Uncovering Privacy Attitudes in Instant Messaging. In *Proceedings of the 5th ACM Conference on Supporting Group Work*, 101–104.
 261. Sameer Patil and Jennifer Lai. 2005. Who Gets to Know What when: Configuring Privacy Permissions in an Awareness Application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 101–110. <https://doi.org/10.1145/1054972.1054987>
 262. Sameer Patil, Greg Norcie, Apu Kapadia, and Adam J. Lee. 2012. Reasons, rewards, regrets: privacy considerations in location sharing as an interactive practice. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS '12)*, 5:1–5:15. <https://doi.org/10.1145/2335356.2335363>
 263. 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>
 264. Sameer Patil, Roman Schlegel, Apu Kapadia, and Adam J. Lee. 2014. Reflection or Action?: How Feedback and Control Affect Location Sharing Decisions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, 101–110. <https://doi.org/10.1145/2556288.2557121>
 265. Nathaniel Persily. 2017. The 2016 U.S. Election: Can Democracy Survive the Internet? *Journal of Democracy* 28, 2: 63–76. <https://doi.org/10.1353/jod.2017.0025>
 266. Joseph Phelps, Glen Nowak, and Elizabeth Ferrell. 2000. Privacy Concerns and Consumer Willingness to Provide Personal Information. *Journal of Public Policy & Marketing* 19, 1: 27–41. <https://doi.org/10.1509/jppm.19.1.27.16941>
 267. Huseyin Polat and Wenliang Du. 2005. Privacy-Preserving Collaborative Filtering. *International Journal of Electronic Commerce* 9, 4: 9–35. <https://doi.org/10.1080/10864415.2003.11044341>
 268. Irene Pollach. 2007. What's Wrong with Online Privacy Policies? *Communications of the ACM* 50, 9: 103–108. <https://doi.org/10.1145/1284621.1284627>
 269. Clay Posey, Paul Benjamin Lowry, Tom L. Roberts, and T. Selwyn Ellis. 2010. Proposing the online community self-disclosure model: the case of working professionals in France and the U.K. who use online communities. *European Journal of Information Systems* 19, 2: 181–195. <https://doi.org/10.1057/ejis.2010.15>
 270. 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 '16)*.
 271. 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
 272. 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>

273. 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. Sottolare, Arthur Graesser, Xiangen Hu and Heather Holden (eds.). U.S. Army Research Laboratory, Orlando, FL, 217.
274. Steffen Rendle and Lars Schmidt-Thieme. 2008. Online-updating Regularized Kernel Matrix Factorization Models for Large-scale Recommender Systems. In *Proceedings of the 2008 ACM Conference on Recommender Systems (RecSys '08)*, 251–258. <https://doi.org/10.1145/1454008.1454047>
275. Daniele Riboni and Claudio Bettini. 2012. Private context-aware recommendation of points of interest: An initial investigation. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops*, 584–589. <https://doi.org/10.1109/PerComW.2012.6197582>
276. Francesco Ricci, Lior Rokach, and Bracha Shapira (eds.). 2015. *Recommender Systems Handbook*. Springer US.
277. Nora J Rifon, Robert LaRose, and Sejung Marina Choi. 2005. Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures. *Journal of Consumer Affairs* 39, 2: 339–360. <https://doi.org/10.1111/j.1745-6606.2005.00018.x>
278. Peter Rosen and Donald Kluemper. 2008. The Impact of the Big Five Personality Traits on the Acceptance of Social Networking Website. *AMCIS 2008 Proceedings*. Retrieved from <http://aisel.aisnet.org/amcis2008/274>
279. Bernard L. Rosenbaum. 1973. Attitude toward invasion of privacy in the personnel selection process and job applicant demographic and personality correlates. *Journal of Applied Psychology* 58, 3: 333–338. <https://doi.org/10.1037/h0036294>
280. Paul R. Rosenbaum and Donald B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1: 41–55. <https://doi.org/10.1093/biomet/70.1.41>
281. Craig Ross, Emily S. Orr, Mia Sisic, Jaime M. Arseneault, Mary G. Simmering, and R. Robert Orr. 2009. Personality and motivations associated with Facebook use. *Computers in Human Behavior* 25, 2: 578–586. <https://doi.org/10.1016/j.chb.2008.12.024>
282. Tracii Ryan and Sophia Xenos. 2011. Who uses Facebook? An investigation into the relationship between the Big Five, shyness, narcissism, loneliness, and Facebook usage. *Computers in Human Behavior* 27, 5: 1658–1664. <https://doi.org/10.1016/j.chb.2011.02.004>
283. 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>
284. Yukiko Sawaya, Mahmood Sharif, Nicolas Christin, Ayumu Kubota, Akihiro Nakarai, and Akira Yamada. 2017. Self-Confidence Trumps Knowledge: A Cross-Cultural Study of Security Behavior. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*, 2202–2214. <https://doi.org/10.1145/3025453.3025926>
285. Peter Schaar. 2010. Privacy by Design. *Identity in the Information Society* 3, 2: 267–274. <https://doi.org/10.1007/s12394-010-0055-x>
286. J. Ben Schafer, Joseph A Konstan, and John Riedl. 2001. E-Commerce Recommendation Applications. *Data Mining and Knowledge Discovery* 5: 115–153.
287. Sae Schatz. 2016. *The Total Learning Architecture: Rationale and Design*. Advanced Distributed Learning Initiative, Alexandria, VA.
288. Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. 2002. Methods and Metrics for Cold-start Recommendations. In *Proceedings of the 25th International Conference on Research and Development in Information Retrieval (SIGIR '02)*, 253–260. <https://doi.org/10.1145/564376.564421>
289. Shalom H. Schwartz. 1994. Beyond individualism/collectivism: New cultural dimensions of values. In *Individualism and collectivism: Theory, method, and applications*, U. Kim, H. C. Triandis, Ç. Kâğıtçıbaşı, S. -C and G. Yoon (eds.). Sage Publications, Inc, Thousand Oaks, CA, US, 85–119.
290. Stuart S Shapiro. 2009. Privacy by design: moving from art to practice. *Commun. ACM* 53: 27–29. <https://doi.org/10.1145/1743546.1743559>
291. Hong Sheng, Fiona Fui-Hoon Nah, and Keng Siau. 2008. An Experimental Study on Ubiquitous commerce Adoption: Impact of Personalization and Privacy Concerns. *Journal of the Association for Information Systems* 9, 6: 344–376. Retrieved from <http://aisel.aisnet.org/jais/vol9/iss6/15>
292. Shlomi Sher and Craig R. M. McKenzie. 2006. Information leakage from logically equivalent frames. *Cognition* 101, 3: 467–494. <https://doi.org/10.1016/j.cognition.2005.11.001>

293. Thomas B. Sheridan and William L. Verplank. 1978. *Human and Computer Control of Undersea Teleoperators*. MIT Cambridge Man-Machine Systems Lab. Retrieved from <http://www.dtic.mil/docs/citations/ADA057655>
294. Kent Shi and Kamal Ali. 2012. GetJar Mobile Application Recommendations with Very Sparse Datasets. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12)*, 204–212. <https://doi.org/10.1145/2339530.2339563>
295. Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and Creepiness in App Space: Perceptions of Privacy and Mobile App Use. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*, 2347–2356. <https://doi.org/10.1145/2556288.2557421>
296. Reza Shokri and Vitaly Shmatikov. 2015. Privacy-Preserving Deep Learning. 1310–1321. <https://doi.org/10.1145/2810103.2813687>
297. Solveig M. Singleton and Jim Harper. 2002. *With A Grain of Salt: What Consumer Privacy Surveys Don't Tell Us*. Social Science Research Network, Rochester, NY. Retrieved May 20, 2014 from <http://papers.ssrn.com/abstract=299930>
298. Manya Sleeper, Lorrie Faith Cranor, and Sarah K. Pearman. 2017. Exploring Topic-Based Sharing Mechanisms. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*, 6973–6985. <https://doi.org/10.1145/3025453.3025840>
299. Aaron Smith. 2014. 6 new facts about Facebook. *Pew Research Center*. Retrieved April 25, 2014 from <http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/>
300. H. J Smith, Sandra J Milberg, and Sandra J Burke. 1992. Concern for Privacy Instrument. *Working document*: School of Business, Georgetown University, Washington, DC.
301. H. Jeff Smith, Sandra J Milberg, and Sandra J Burke. 1996. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. *MIS Quarterly* 20, 2: 167–196. <https://doi.org/10.2307/249477>
302. 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>
303. Barry Smyth. 2007. Case-Based Recommendation. In *The Adaptive Web: Methods and Strategies of Web Personalization*, Peter Brusilovsky, Alfred Kobsa and Wolfgang Nejdl (eds.). Springer Verlag, Berlin, 342–376.
304. Daniel J. Solove. 2007. "I've Got Nothing to Hide" and Other Misunderstandings of Privacy. *San Diego Law Review* 44: 745. Retrieved September 1, 2013 from <http://papers.ssrn.com/abstract=998565>
305. Daniel J. Solove. 2013. Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126: 1880–1903.
306. Luís Fernandes Sousa. 2009. Privacy Policy Dynamics in Location Sharing Applications. University of Lisbon, Department of Computer Science, Lisbon, Portugal. Retrieved December 13, 2010 from <http://docs.di.fc.ul.pt/handle/10455/3286>
307. E. Isaac Sparling and Shilad Sen. 2011. Rating: How Difficult is It? In *Proceedings of the Fifth ACM Conference on Recommender Systems*, 149–156. <https://doi.org/10.1145/2043932.2043961>
308. Sarah Spiekermann. 2012. The Challenges of Privacy by Design. *Commun. ACM* 55, 7: 38–40. <https://doi.org/10.1145/2209249.2209263>
309. 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. <https://doi.org/10.1145/501158.501163>
310. Jan-Benedict E.M Steenkamp and Inge Geyskens. 2006. How Country Characteristics Affect the Perceived Value of Web Sites. *Journal of Marketing* 70, 3: 136–150. <https://doi.org/10.1509/jmkg.70.3.136>
311. Stilgherrian. 2014. Why big data evangelists need to be reprogrammed. *ZDNet*. Retrieved October 28, 2015 from <http://www.zdnet.com/article/why-big-data-evangelists-need-to-be-reprogrammed/>
312. 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.
313. Fred Stutzman, Ralph Gross, and Alessandro Acquisti. 2013. Silent Listeners: The Evolution of Privacy and Disclosure on Facebook. *Journal of Privacy and Confidentiality* 4, 2. <https://doi.org/10.29012/jpc.v4i2.620>

314. S. Shyam Sundar and Sampada S. Marathe. 2010. Personalization versus Customization: The Importance of Agency, Privacy, and Power Usage. *Human Communication Research* 36, 3: 298–322. <https://doi.org/10.1111/j.1468-2958.2010.01377.x>
315. Juliana Sutanto, Elia Palme, Chuan-Hoo Tan, and Chee Wei Phang. 2013. Addressing the Personalization-Privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users. *MIS Quarterly* 37, 4: 1141–1164. Retrieved from <http://aisel.aisnet.org/misq/vol37/iss4/9/>
316. 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>
317. Karen Tang, Jialiu Lin, Jason Hong, Daniel Siewiorek, and Norman Sadeh. 2010. Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing. In *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing*, 85–94. <https://doi.org/10.1145/1864349.1864363>
318. Martin Tanis and Tom Postmes. 2005. A Social Identity Approach to Trust: Interpersonal Perception, Group Membership, and Trusting Behavior. *European Journal of Social Psychology* 35: 413–424. <https://doi.org/10.1002/ejsp.256>
319. David Taylor, Donna Davis, and Ravi Jillapalli. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electronic Commerce Research* 9, 3: 203–223. <https://doi.org/10.1007/s10660-009-9036-2>
320. Shelley E. Taylor and Marci Lobel. 1989. Social comparison activity under threat: Downward evaluation and upward contacts. *Psychological Review* 96, 4: 569–575. <https://doi.org/10.1037/0033-295X.96.4.569>
321. Max Teltzrow and Alfred Kobsa. 2004. Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study. In *Designing Personalized User Experiences for eCommerce*, Clare-Marie Karat, Jan Blom and John Karat (eds.). Kluwer Academic Publishers, Dordrecht, Netherlands, 315–332.
322. Omer Tene and Jules Polonetsky. 2013. A Theory of Creepy: Technology, Privacy and Shifting Social Norms. *Yale Journal of Law and Technology* 16: 59–102. Retrieved from <http://digitalcommons.law.yale.edu/yjolt/vol16/iss1/2/>
323. Luis Terán and Aigul Kaskina. 2016. Enhancing Voting Advice Applications with Dynamic Profiles. In *Proceedings of the 9th International Conference on Theory and Practice of Electronic Governance (ICEGOV '15-16)*, 254–257. <https://doi.org/10.1145/2910019.2910043>
324. Richard H Thaler and Cass Sunstein. 2008. *Nudge : improving decisions about health, wealth, and happiness*. Yale University Press, New Haven, NJ & London, U.K.
325. Nava Tintarev and Judith Masthoff. 2011. Designing and Evaluating Explanations for Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira and Paul B. Kantor (eds.). Springer US, Boston, MA, 479–510. Retrieved October 1, 2011 from <http://www.springerlink.com/content/x61ll765673660g7/>
326. Eran Toch, Justin Cranshaw, Paul Hanks Drielsma, Janice Y. Tsai, Patrick Gage Kelley, James Springfield, Lorrie Cranor, Jason Hong, and Norman Sadeh. 2010. Empirical models of privacy in location sharing. In *Proceedings of the 12th ACM international conference on Ubiquitous computing (UbiComp '10)*, 129–138. <https://doi.org/10.1145/1864349.1864364>
327. Eran Toch, Norman M. Sadeh, and Jason Hong. 2010. Generating default privacy policies for online social networks. In *Ext. Abstracts CHI 2010*, 4243–4248. <https://doi.org/10.1145/1753846.1754133>
328. Horst Treiblmaier and Irene Pollach. 2007. Users' Perceptions of Benefits and Costs of Personalization. In *ICIS 2007 Proceedings*. Retrieved from <http://aisel.aisnet.org/icis2007/141>
329. Janice Y. Tsai, Patrick Kelley, Paul Drielsma, Lorrie Faith Cranor, Jason Hong, and Norman Sadeh. 2009. Who's viewed you?: the impact of feedback in a mobile location-sharing application. In *Proceedings of the 27th international conference on Human factors in computing systems*, 2003–2012. <https://doi.org/10.1145/1518701.1519005>
330. Jannice Y. Tsai, Sergei Egelman, Lorrie F. Cranor, and Alessandro Acquisti. 2010. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 21, 1: 1–18. <https://doi.org/10.1287/isre.1090.0260>
331. Zeynep Tufekci. 2008. Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites. *Bulletin of Science, Technology & Society* 28, 1: 20–36. <https://doi.org/10.1177/0270467607311484>

332. Michael A Turner and Robin Varghese. 2002. *Making sense of the privacy debate: a comparative analysis of leading consumer privacy surveys*. Privacy & American Business.
333. S. Tyagi and K. K. Bharadwaj. 2012. Trust-enhanced recommender system based on case-based reasoning and collaborative filtering. In *2012 2nd International Conference on Power, Control and Embedded Systems*, 1–4. <https://doi.org/10.1109/ICPCES.2012.6508112>
334. David Vallet, Arik Friedman, and Shlomo Berkovsky. 2014. Matrix Factorization without User Data Retention. In *Advances in Knowledge Discovery and Data Mining*, Vincent S. Tseng, Tu Bao Ho, Zhi-Hua Zhou, Arbee L. P. Chen and Hung-Yu Kao (eds.). Springer International Publishing, 569–580.
335. Craig Van Slyke, J. T. Shim, Richard Johnson, and James J. Jiang. 2006. Concern for Information Privacy and Online Consumer Purchasing. *Journal of the Association for Information Systems* 7, 1. Retrieved March 15, 2013 from <http://aisel.aisnet.org/jais/vol7/iss1/16>
336. Rich Viano. 2015. OSLM. *ADL Net*. Retrieved January 25, 2016 from <http://adlnet.gov/oslm/>
337. Jesse Vig, Shilad Sen, and John Riedl. 2009. Tagsplanations: Explaining Recommendations Using Tags. In *Proceedings of the 14th International Conference on Intelligent User Interfaces (IUI '09)*, 47–56. <https://doi.org/10.1145/1502650.1502661>
338. Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative Deep Learning for Recommender Systems. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*, 1235–1244. <https://doi.org/10.1145/2783258.2783273>
339. Na Wang, Jens Grossklags, and Heng Xu. 2013. An Online Experiment of Privacy Authorization Dialogues for Social Applications. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*, 261–272. <https://doi.org/10.1145/2441776.2441807>
340. Na Wang, Heng Xu, and Jens Grossklags. 2011. Third-party apps on Facebook: privacy and the illusion of control. In *Proceedings of the 5th ACM Symposium on Computer Human Interaction for Management of Information Technology (CHIMIT '11)*, 4:1–4:10. <https://doi.org/10.1145/2076444.2076448>
341. W. Wang and I. Benbasat. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems* 23, 4: 217–246. <https://doi.org/10.2753/MIS0742-1222230410>
342. Yang Wang and Alfred Kobsa. 2007. Respecting Users' Individual Privacy Constraints in Web Personalization. In *User Modeling 2007 (Lecture Notes in Computer Science)*, 157–166. https://doi.org/10.1007/978-3-540-73078-1_19
343. Yang Wang and Alfred Kobsa. 2013. A PLA-based privacy-enhancing user modeling framework and its evaluation. *User Modeling and User-Adapted Interaction* 23, 1: 41–82. <https://doi.org/10.1007/s11257-011-9114-8>
344. 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>
345. 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. <https://doi.org/10.1145/2487788.2488038>
346. Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. 2011. "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*, 10:1–10:16. <https://doi.org/10.1145/2078827.2078841>
347. Yang Wang, Gregory Norice, and Lorrie Faith Cranor. 2011. Who is concerned about what? A study of American, Chinese and Indian users' privacy concerns on social network sites. In *International Conference on Trust and Trustworthy Computing*, 146–153. Retrieved August 3, 2016 from http://link.springer.com/10.1007%2F978-3-642-21599-5_11
348. J. Watson. 2015. Predicting privacy settings with a user-centered approach. In *2015 International Conference on Collaboration Technologies and Systems (CTS)*, 499–500. <https://doi.org/10.1109/CTS.2015.7210443>
349. Jason Watson, Andrew Besmer, and Heather Richter Lipford. 2012. +Your circles: sharing behavior on Google+. In *Proceedings of the 8th Symposium on Usable Privacy and Security*, 12:1–12:10. <https://doi.org/10.1145/2335356.2335373>

350. 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>
351. 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.
352. Alan F. Westin and Danielle Maurici. 1998. *E-Commerce & Privacy: What the Net Users Want*. Privacy & American Business, and PricewaterhouseCoopers LLP.
353. White House. 2012. *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Economy*. White House, Washington, D.C.
354. Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android Permissions Remystified: A Field Study on Contextual Integrity. In *Proceedings of the 24th USENIX Conference on Security Symposium (SEC'15)*, 499–514. Retrieved July 30, 2018 from <http://dl.acm.org/citation.cfm?id=2831143.2831175>
355. Primal Wijesekera, Arjun Baokar, Lynn Tsai, Joel Reardon, Serge Egelman, David Wagner, and Konstantin Beznosov. 2017. The Feasibility of Dynamically Granted Permissions: Aligning Mobile Privacy with User Preferences. Retrieved from https://www.ftc.gov/system/files/documents/public_comments/2016/09/00018-129026.pdf
356. Daricia Wilkinson, Saadhika Sivakumar, David Cherry, Bart P. Knijnenburg, Elaine M. Raybourn, Pamela Wisniewski, and Henry Sloan. 2017. User-Tailored Privacy by Design. In *Proceedings of the Usable Security Mini Conference (USEC '17)*. <http://dx.doi.org/10.14722/usec.2017.23007>
357. Martijn C. Willemsen, Mark P. Graus, and Bart P. Knijnenburg. 2016. Understanding the role of latent feature diversification on choice difficulty and satisfaction. *User Modeling and User-Adapted Interaction* 26, 4: 347–389. <https://doi.org/10.1007/s11257-016-9178-6>
358. Dave Wilson, Jeffrey Proudfoot, and Joseph Valacich. 2014. Saving Face on Facebook: Privacy Concerns, Social Benefits, and Impression Management. *ICIS 2014 Proceedings*. Retrieved from <http://aisel.aisnet.org/icis2014/proceedings/SocialMedia/3>
359. Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor, Jay Springfield, Sae Young Jeong, and Arun Balasubramanian. 2013. Privacy manipulation and acclimation in a location sharing application. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing (UbiComp '13)*, 549–558. <https://doi.org/10.1145/2493432.2493436>
360. 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. <https://doi.org/10.17705/1CAIS.03810>
361. Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies* 98: 95–108. <https://doi.org/10.1016/j.ijhcs.2016.09.006>
362. Pamela Wisniewski, Bart P. Knijnenburg, and H. Richter Lipford. 2014. Profiling Facebook Users' Privacy Behaviors. In *SOUPS2014 Workshop on Privacy Personas and Segmentation*.
363. 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>
364. 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>
365. Allison Woodruff, Vasyl Pihur, Sunny Consolvo, Lauren Schmidt, Laura Brandimarte, and Alessandro Acquisti. 2014. Would a Privacy Fundamentalist Sell Their DNA for \$1000...If Nothing Bad Happened as a Result? The Westin Categories, Behavioral Intentions, and Consequences. In *Proceedings of the Symposium On Usable Privacy and Security: SOUPS '14*.
366. Luke Wroblewski. 2008. *Web Form Design: Filling in the Blanks*. Rosenfeld Media, Brooklyn, NY.
367. Hongchen Wu, Bart P. Knijnenburg, and Alfred Kobsa. 2014. Improving the prediction of users' disclosure behavior... by making them disclose more predictably? In *Symposium on Usable Privacy and Security (SOUPS)*. Retrieved January 22, 2015 from <http://www.ics.uci.edu/~kobsa/papers/2014-SOUPS-PPS-Kobsa.pdf>

368. Jierui Xie, Bart Piet Knijnenburg, and Hongxia Jin. 2014. Location Sharing Privacy Preference: Analysis and Personalized Recommendation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces (IUI '14)*, 189–198. <https://doi.org/10.1145/2557500.2557504>
369. Heng Xu, Hock-Hai Teo, and Bernard C Y Tan. 2005. Predicting the Adoption of Location-Based Services: The Role of Trust and Perceived Privacy Risk. In *Proceedings of the International Conference on Information Systems*, 861–874. Retrieved from <http://aisel.aisnet.org/icis2005/71>
370. Heng Xu, Hock-Hai Teo, Bernard C. Y. Tan, and Ritu Agarwal. 2009. The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services. *Journal of management information systems* 26, 3: 135–174. <https://doi.org/10.2753/MIS0742-1222260305>
371. Heng Xu, Na Wang, and Jens Grossklags. 2012. Privacy-by-ReDesign: Alleviating Privacy Concerns for Third-Party Applications. In *ICIS 2012 Proceedings*.
372. Bo Yan and Guanling Chen. 2011. AppJoy: Personalized Mobile Application Discovery. In *Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services (MobiSys '11)*, 113–126. <https://doi.org/10.1145/1999995.2000007>
373. Alyson L. Young and Anabel Quan-Haase. 2009. Information Revelation and Internet Privacy Concerns on Social Network Sites: A Case Study of Facebook. In *Proceedings of the Fourth International Conference on Communities and Technologies*, 265–274. <https://doi.org/10.1145/1556460.1556499>
374. Yuchen Zhao, Juan Ye, and Tristan Henderson. 2014. Privacy-aware Location Privacy Preference Recommendations. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MOBIQUITOUS '14)*, 120–129. <https://doi.org/10.4108/icst.mobiquitous.2014.258017>
375. Yuchen Zhao, Juan Ye, and Tristan Henderson. 2016. The Effect of Privacy Concerns on Privacy Recommenders. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16)*, 218–227. <https://doi.org/10.1145/2856767.2856771>
376. Tao Zhou. 2012. Understanding users' initial trust in mobile banking: An elaboration likelihood perspective. *Computers in Human Behavior* 28, 4: 1518–1525. <https://doi.org/10.1016/j.chb.2012.03.021>
377. Tianqing Zhu, Yongli Ren, Wanlei Zhou, Jia Rong, and Ping Xiong. 2014. An effective privacy preserving algorithm for neighborhood-based collaborative filtering. *Future Generation Computer Systems* 36: 142–155. <https://doi.org/10.1016/j.future.2013.07.019>
378. J. Christopher Zimmer, Riza Ergun Arsal, Mohammad Al-Marzouq, and Varun Grover. 2010. Investigating online information disclosure: Effects of information relevance, trust and risk. *Information & Management* 47, 2: 115–123. <https://doi.org/10.1016/j.im.2009.12.003>