PS4TLA: Privacy Support for the Total Learning Architecture

Specification Document, version 0.1:

Operational Characteristics

Author information

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

Project team: David Cherry, Daricia Wilkinson, Saadhika Sivakumar, Reza Ghaiumy Anaraky, Moses Namara, Dr. Erin Ash, Rosey Davis

Technical point of contact: Dr. Elaine Raybourn, Sandia National Laboratories

Recommended citation: Knijnenburg, B.P., Cherry, D., Wilkinson, D., Sivakumar, S., Ghaiumy Anaraky, R., Namara, M., Ash, E., Davis, R.M., and Raybourn, E.M. (2017) "Privacy Support for the Total Learning Architecture: Operational Characteristics". PS4TLA Specification Document, v0.1, Clemson University, Clemson, SC. Available at: http://www.usabart.nl/PS4TLA/spec0.1.pdf.

Executive summary

The purpose of this document is to inform ADL and other TLA performers about the Operational Characteristics that impact users' privacy concerns and to make recommendations for implementing a Privacy by Design (PbD) model where privacy decisions of the system are made in its initial developmental stages. The set of recommendations put forth in this document will allow ADL and other TLA performers to select the characteristics that best alleviate users' concerns.

In the first section, **User Characteristics**, the variance in individuals' privacy concerns are considered. Personal characteristics of users, including decision-making mechanisms, cognitive-processing practices, and communication styles, have significant implications for their privacy management behaviors. To address these characteristics, it is recommended that TLA tailor to different privacy management strategies and communication styles in its design. For example, users should be given options for selective sharing of their information and outcomes, and both social-network style and direct, chat-style interaction features should be available.

A section on **Input Data Characteristics** addresses issues surrounding the collection of users' personal information. Data collection is fundamental to providing personalized learning experiences that adapt to users. However, it is important to consider the privacy implications of collecting the potentially sensitive data about users that is necessary for adaptive systems to function effectively. In some cases, the TLA may make incorrect predictions or predictions that users may be uncomfortable with. To solve this issue, users should be allowed to scrutinize and correct potential mistakes in system predictions, as well as to venture beyond the personalized recommendations.

The **Output Characteristics** section offers recommendations for presenting adaptations to users. Personalized notifications are an important feature of the TLA architecture, which specifies three types of adaptations, specifically providing individualized recommendations to switch from one Learner Activity to another, determining the next learning activity within a single activity provider, and adapting learning content within a single learning activity. For these mechanisms to be effective, they must be accurate without being intrusive or inconvenient. Notifications should be carefully planned to prevent interrupting a user's current task. Systems should also be designed to prevent leaking personal information in social settings by providing only generic notifications or, where possible, tailoring notifications to specific social settings.

A section titled **Data Location and Ownership** addresses questions of data location and ownership from a user-privacy perspective. Because the TLA is inherently decentralized by design, it is important decisions about where collected data will be stored and what entities have access to it must be made. Such decisions should reflect the spirit of "open" learning models by giving users ownership over their data. Of course, employers and apps will necessarily have access to some data, but access should be limited to narrowly specified purposes and take steps

to maintain user privacy, such as de-identifying data when possible. Another strategy may be to allow users to designate a "data steward" to manage their data in accordance with their privacy preferences. The TLA should also make user models portable so users can take their data with them as they move between employers.

The **Data Sharing** section outlines how recipients of user data can preserve user privacy when using the data for various purposes. It is important for managers to communicate secondary data usage practices to users; users should be aware of what information collected about them is used and how. Managers should also act responsibly regarding placement and promotion decisions by being transparent about the guidelines they use to assess potential conflicts between competencies and preferences and to prevent discriminatory practices. Finally, Institutional Research Board (IRB) guidelines for research, which require data to be anonymized to the degree possible, should be followed.

The last section, **Privacy Support Mechanisms**, discusses techniques for user-tailored privacy (UTP). UTP moves beyond a "one-size-fits-all" approach to privacy design by accounting for the high variability and context-dependency of people's privacy decisions. UTP aims to strike a balance between giving users no control over their privacy, which may elicit privacy concerns, and giving users full control over their privacy, which is often unmanageably complex for the typical user. Successful implementation of UTP requires taking such steps as using accessible privacy controls, and using users' behavioral patterns to make privacy-related adaptations.

Given the complexity of privacy in advanced distributed learning systems, upcoming versions of this document will delve deeper into the idea of user-tailored privacy as a decision-support mechanism for TLA. For the final document, we will seek consensus among TLA performers regarding the operational characteristics 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. As the TLA specifications are being developed, there exists an opportunity to implement Privacy by Design (PbD), where privacy is treated as a fundamental part of the system, and taken into account throughout the entire development lifecycle of the system, starting at the early stages of design and development [51, 207, 308, 314, 335]. This document therefore describes the potential impact of the Operational Characteristics (OCs) of TLA-based systems on users' privacy concerns. The OCs are aspects of TLA-based systems that can be implemented in various ways. The purpose of this document is to allow ADL and other TLA performers to select the operational variants that best alleviate users' privacy concerns.

Scope

This document describes the operational variations of privacy-relevant OCs of distributed learning systems in general—and specifically of TLA—as well as their impact on users' privacy concerns. Where possible, an attempt is made to juxtapose the privacy concerns with the benefit of the described operational variant.

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 OCs and their variants may therefore go beyond any currently envisioned specification and implementation of TLA.

OCs that are not privacy-relevant are not discussed in this document. Most notably, this document does not concern the security of the TLA. Security-related OCs are only discussed where they intersect with user privacy concerns.

Several types of actors are involved in the development of the TLA specifications, and the implementation and operation of distributed learning systems based on these specifications. To aid different actors in navigating this document, we highlight parts of it that are particularly relevant for specific audiences. These audiences are described in Table 1.

Where possible, the document contains concrete recommendations. Further recommendations will be added after intensive discussion with ADL and other TLA performers during the development of version 1.0 of this document.

Table 1: Actors that form the potential audiences for this document.

Icon	Description
	**Training Manager**—Responsible for evaluation, promotion, mission planning, user data management and research
	**Activity Developer**—TLA End User Application Developers who develop and implement the TLA user facing apps
	**TLA Backend Developer**—Develops the Data Core and TLA Processors
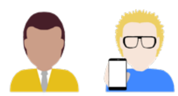	**Applicable to Training Manager + Activity Developer**
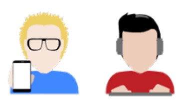	**Applicable to Activity Developer + TLA Backend Developer**
	**Applicable to all**

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 [310]. 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 [98]. 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) [370]. 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.


Privacy by Design (PbD) is a design philosophy in which privacy aspects are addressed early in the system design and development process, rather than after the system has been developed ("post hoc privacy") [51, 207, 308, 314, 335]. While post hoc privacy solutions typically try to mitigate privacy problems that exist within a system, PbD tries to avoid privacy problems from occurring at all. This document addresses PbD by analyzing the proposed operational characteristics of the TLA from a privacy perspective.

An **Operational Characteristic (OC)** is an aspect of TLA that influences the users' experience. OCs are compositional, in that each OC (e.g. "input data") consists of underlying sub-OCs (e.g. "learner runtime activity", "smartphone tracking data"). An OC can be implemented in multiple ways across different operational dimensions (e.g. learner runtime activity may be tracked in a "granular" or "aggregated" manner, either in "real time" or "asynchronous")—these are called operational *variants*. A *privacy-relevant* OC is defined as an OC whose variants have an impact on users' privacy concerns.

Personally Identifiable Information (PII) is information that reveals a person's real-life identity, e.g. their name, social security number, or (in most cases) their primary email address.

De-identification is the practice of removing Personally Identifiable Information (PII) from a set of data. Given that data that are in themselves not personally identifiable may be designated as PII when used in combination, the term *k*-anonymity is used to characterize a dataset as containing no less than *k* exemplars of a certain combination of values. Generally speaking, a dataset is considered de-identified when all PII is removed, and *k*-anonymity is guaranteed for all combinations of non-PII data.

Pseudonymity is a means to identify a person in a system without revealing any links to their true identity outside the system. Pseudonymity is usually implemented by allowing users to choose a username that deviates from their real name.

The **Health Insurance Portability and Accountability Act (HIPAA)** establishes data privacy and security provisions for safeguarding medical information.

Overview

Privacy threats have shown to be an important barrier to the adoption of personalized systems [21, 54, 105, 193, 289, 317, 342, 357, 366], 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 [48, 59, 68, 79, 209]. 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 [7, 155, 157, 185, 205, 225, 231], leaving them vulnerable to perceived and real privacy threats.

Fortunately, the TLA specifications and reference implementation are still in the early stage of development, which presents an opportunity implement Privacy by Design (PbD). This document supports a comprehensive implementation of PbD by systematically investigating the impact of the Operational Characteristics (OCs) of TLA-based distributed learning systems on users' privacy. This allows ADL and other TLA performers to make informed decisions about which operational variations present the optimal tradeoff between privacy and other considerations. In cases where less-than-ideal privacy solution may be preferred for other reasons, the specification suggests mitigating (post hoc) solutions to limit the impact on users' privacy.

This document considers the following OCs and sub-OCs:

1. **User characteristics**[1] (learners' privacy decision-making practices)
   1.1. Decision-making
   1.2. Elaboration likelihood
   1.3. Communication style
2. **Input data characteristics** (data collection by the TLA Data Core)
   2.1. Levels of identifiability
   2.2. Collection of various data types
   2.3. Inferences made based on collected data
3. **Output characteristics** (mechanisms for conveying learning adaptations to the users)
   3.1. Recommendation presentation methods and mechanisms
   3.2. Output modalities and devices
   3.3. Feedback and conversation about recommendations
4. **Data location and ownership** (learner data management within TLA-based architectures)
   4.1. Managing meta-, macro-, and micro- adaptations
   4.2. Data ownership and stewardship
5. **Data sharing** (social and organizational aspects of distributed learning systems)
   5.1. Scrutability and the quantified self
   5.2. Social learning experiences
   5.3. Assessment, promotion, and mission planning
6. **Privacy support mechanisms** (supporting learners' privacy decision-making)
   6.1. Privacy notices

---

[1] User characteristics are of course outside the control of the system developers, but provide important parameters that need to be considered in the design of the TLA's privacy features.

6.2. Control mechanisms
6.3. Privacy nudging
6.4. User-tailored privacy

Each OC and sub-OC is further unpacked, and the tradeoff between privacy and other considerations are described for all operational variations. Where possible, concrete recommendations are made. Further recommendations will be added after intensive discussion with ADL and other TLA performers.

1 User characteristics

Problem: What is the user's perspective? The main purpose of this document is to define operational parameters for the TLA that are acceptable for its users from a privacy perspective. How do users react to privacy-related decisions, and how does their perspective come about?

Current state of the art: Little focus on user characteristics. A lot of the existing privacy and security literature focuses on *technical* solutions to privacy, and often disregards the complexities of the behavior of users who operate within this technical landscape.

Solution: Study user characteristics. In this section, we acknowledge that users vary extensively in their information disclosure behavior, as evidenced by the following research:

- Recurring privacy surveys by Westin and Harris Interactive that started in the early '80s consistently find a substantial diversity in users' extent of privacy concerns. They identify three types of users: fundamentalists, the unconcerned, and a pragmatic majority [130, 131, 379, 380].
- Recent research shows that users' disclosure behavior is multi-dimensional [184], i.e., users differ not just in the *amount* of information that they disclose, but also in the *kind* of information that they are most and least likely to disclose.
- Research shows that even for the same person, the disclosure decision depends on the context in which it is made [28, 32, 64, 143, 158, 160, 214, 220, 260, 262, 266, 276, 356, 377].
- Indeed, the variability and context-dependency of privacy preferences is at the core of many privacy theories such as Altman's *privacy regulation theory* [12], Nissenbaum's *contextual integrity* [259, 260], and Petronio's *communication privacy management* [282, 283].

This section analyzes how these differences in privacy concerns and behaviors come about, which results in important PbD recommendations for TLA-based systems. In this section, you will learn about the user characteristics that affect privacy concerns and behaviors, with specific consideration of individuals':

- Decision-making mechanisms
- Cognitive processing practices
- Communication styles

Key findings and recommendations are presented in Table 2.

Table 2: Key findings regarding the user characteristics

	Key Findings	Recommendations
Decision-Making (1.1)	– Balance privacy risks and relevance – Users are not always rational – Trust increases acceptance of data collection and tracking	– Survey users – Build trust – Highlight relevance
Cognitive Processing (1.2)	– Decision practices range from heuristic to effortful (rational) – Motivation and ability influence processing style – Motivation and ability can be instilled	– Cater to both heuristic and effortful processing styles – Encourage users to make rational decisions through effortful processing
Communication Style (1.3)	– Social network users selectively apply privacy management strategies – Social network users use different communication styles	– Tailor to different privacy management styles – Tailor to different communication styles

1.1 Decision-making

It is important to first acknowledge the mechanisms by which users make privacy-related decisions. While the benefits of adopting TLA specifications may be abundant, there are also various challenges that exist with the collection of large amounts of data. Advances in storage capabilities and data mining abilities enables the TLA Data Core to have a deeper analysis of the preferences and behavior of its users by collecting a vast amount of data [237]. This includes very detailed real-time information from users' cellphone and other type of devices that could reveal information about the decision-making process or personal stances on sensitive topics that they normally would not share with other people or other systems.

Research shows that users acknowledge the benefit of data collection for personalization [366] but when taken too far, the same data collection can deter users from using the system extensively, or even dissuade them from using the system at all. This subsection discusses the research that quantifies this phenomenon, which has been labeled the *personalization-privacy paradox* [21, 54, 104]. We highlight the value of doing user research, building trust, and highlighting the relevance of the information that is being collected (see Table 3).

Table 3: Recommendations regarding decision-making

Survey Users
– Perform scenario-based experiments to quantify the effects of various data collection practices
– Conduct in-depth interviews to uncover users' privacy-related attitudes
– Perform controlled user experiments to detect potentially deleterious effects of heuristic decision practices
Build Trust
– Ensure that the learning applications originate from trustworthy sources
– Employ sensible data collection practices and a privacy by design philosphy from the outset
Highlight Relevance
– Highlight the potential improvements in content relevance, time saving, enjoyment and novelty
– Refrain from asking for information in situations in which the relevance is not readily apparent

Existing work recommends balancing privacy risks and relevance

Starting with the basic research on information disclosure, one of the most-used (cf. [223, 280, 325]) conceptualizations of users' conscious process behind their information disclosure decisions is the "privacy calculus" [211, 212]. This conceptualization has been used by many researchers to investigate the antecedents of information disclosure [69, 70, 80, 128, 170, 220, 246, 281, 386, 394, 396].

The privacy calculus is a privacy-specific instance general human decision-making theories [21, 223, 304, 338], which argue that people gather information about various aspects of each choice option, assign a value to each of these aspects, trade off the different aspects, and then choose the option that maximizes their utility [34, 97, 320]. What are the aspects that people trade off in privacy decisions? Two aspects are mentioned repeatedly in existing work: perceived risk and perceived relevance.

Perceived risk—Privacy risk is the "potential loss of control over personal information, such as when information about you is used without your knowledge or permission" [95]. This loss of control can lead to unintended uses and distribution of the information [265, 315, 369]. The *perception* of risk is the fear that these unintended consequences will happen [148, 223]. The following research quantifies the effects of perceived risk:

- Several studies found a direct effect of perceived risk on disclosure intentions [219, 220, 262], indicating that risk perceptions may lead us to restrict access to our personal information [221, 281].
- Consumer surveys have found that between 58.2% [242] and 72% [137] of all respondents cite risk as a reason not to disclose their personal information.
- Comparing effect sizes between studies, Dinev & Hart [80] note that privacy risk may even be more likely to dissuade people from making an e-commerce transaction than the economic risk of the transaction (see also [35]).
- Extending this argument, research shows that risk may indeed influence users' intention to transact in a web shop [172, 279], or their intention to adopt an online service [95].
- White argues that "Marketers' efforts may be wisely directed at attempts to mitigate any perceived "downside risks" associated with disclosure." [383].

Moving specifically towards research on privacy in personalized systems, several researchers show that privacy risks may inhibit the use of such services:

- In a study on ubiquitous commerce (u-commerce), Sheng et al. [317] showed that personalization induced privacy concerns, and that users consequently would feel less inclined to use personalized (rather than non-personalized) u-commerce services, unless the benefits were overwhelming (i.e., providing help in an emergency).
- Awad and Krishnan [21] showed that users' privacy concerns inhibited their use of personalized services and advertising.

- Sutanto et al. [342] demonstrated that privacy concerns can prevent people from using a potentially beneficial personalized application.

Perceived relevance—Whereas perceived risk describes the negative side of the privacy calculus, the positive side appears to be governed by the *perceived relevance* of disclosure. The following research quantifies the effects of perceived relevance:

- Stone was the first to consider the effect of the perceived relevance of information requests on privacy-related behaviors [337], and this effect has since been demonstrated empirically [219].
- Phelps et al. note that people's purchase intentions go down when a service requests information that does not serve the purpose of the request. They therefore argue that "marketers need to resist asking for such information in situations in which the relevance is not readily apparent" [289].

In the realm of personalized services, research shows that concerns mainly exist when these services fail to provide useful benefits for which the disclosed information is relevant:

- Scientific research into consumer perceptions shows that people are willing to give up privacy for personalization [128, 265], as long as this gives them benefits [289].
- Deeper investigations into this phenomenon show that users particularly value the benefits of content relevance, time savings, enjoyment and novelty to an extent that may have them ignore their initial privacy concerns [135, 144].
- Consequently, certain researchers claim that "privacy isn't the issue" [123] as long as the benefits are clear [174].

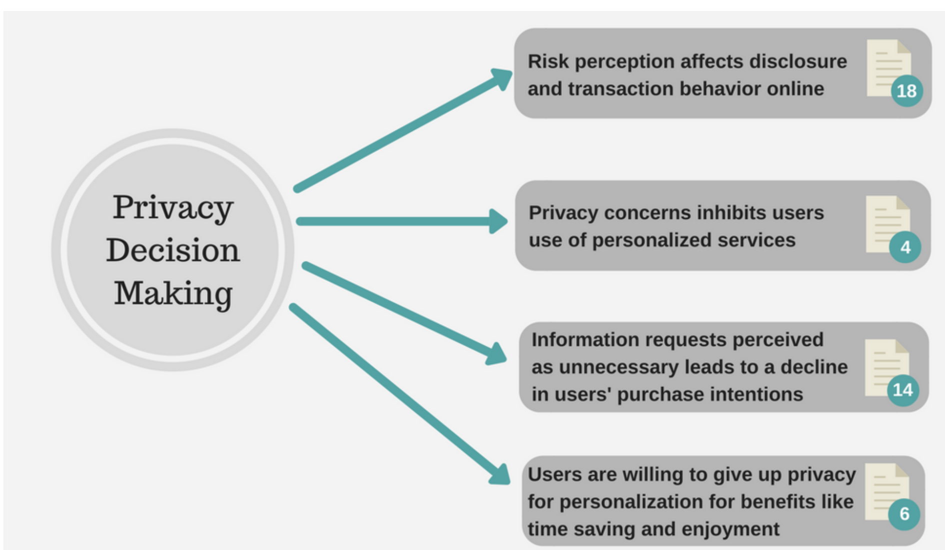

Figure 1: Summary of findings from literature review of privacy decision-making

Integrating these streams of research, existing work has predominantly shown that risk and relevance are both important in determining users' willingness to adopt and provide personal information to personalized services, and researchers therefore claim that they should both meet a certain threshold [357], or that they at least should be in balance [54, 394, 396] (see Figure 1). This finding has been depicted in Figure 2.

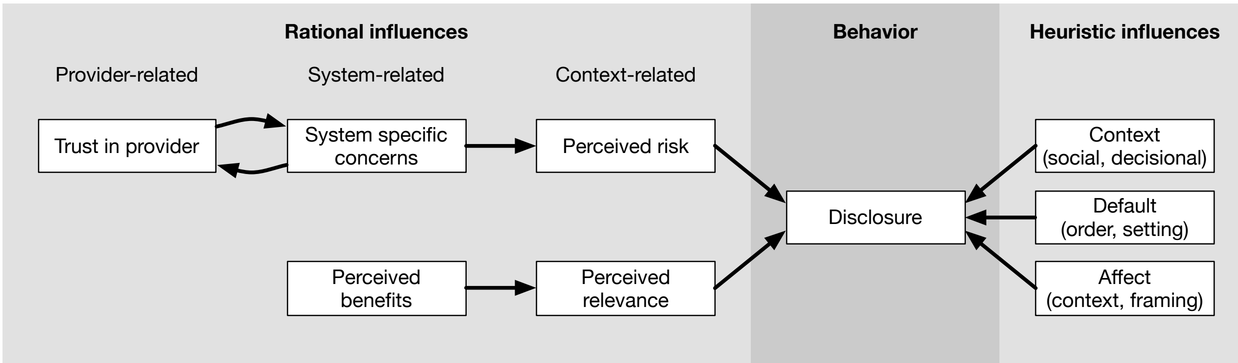

Figure 2: A diagrammatic representation of users' privacy decision process.

Users are not always rational; in-depth investigations are required

One critique of this existing work on the privacy-personalization paradox is that it often fails to truly investigate the tradeoff between risk and relevance as a concrete behavioral decision, because their outcome measure is a more generic form of an intention (i.e., it is measured with generic questionnaire items such as “How likely would you provide your personal information (including your location) to use the M-Coupon service?”). Such stated intentions arguably do not directly relate to observable privacy behaviors (cf. Spiekermann et al. [336] and Norberg et al. [262], who show that privacy preferences and actual behavior tend to be weakly related at best).

Indeed, the privacy calculus itself has been criticized for making unrealistic assumptions about the rationality of decision-makers [168, 169]. Rather than being rational, people’s privacy decisions are influenced by various heuristics, such as:

- Information on others’ privacy decisions (i.e. “social proof” [7])
- The order of sensitivity in which decisions are being made (“foot in the door” and “door in the face” [7])
- The overall professionalism of the privacy-setting user interface (“affect heuristic” [155])
- The available options to choose from (“context non-invariance” [185])
- The default setting and phrasing of privacy-related requests (“default” and “framing” effects [181, 204]).

Future work on disclosure behavior—including investigations of TLA users’ privacy behaviors—should conceptualize perceived risk as contextualized privacy concerns (i.e., concerns about the possible consequences of disclosing a specific piece of information to a specific recipient [70, 232, 289, 326]) and perceived relevance as contextualized benefit: the perceived benefit of disclosing a specific piece of information to a specific recipient [183, 219]. Initial work at the level of individual privacy decisions (a yes/no decision for multiple disclosures) has been successful in separating the rational tradeoff from irrational influences, quantifying their relative contribution [7, 182, 183, 194]. The distinction between general (system-related) concerns/benefits and contextualized (information-related) risk/relevance is also depicted in Figure .

Trust increases users' acceptance of data collection and tracking

Aside from highlighting the relevance of the data collection/tracking practices, there are several ways to convince users to disclose more information. Some of these methods (and their shortcomings) will be described in Section 6. Here we address the topic of trust. The following research quantifies the effect of trust on information disclosure:

- Several researchers suggest that concern/risk is a mediator between trust and disclosure intentions [232, 369, 395, 406]. This suggests that trust may reduce perceived risk, which in turn increases disclosure.
- Dinev et al. argues the opposite effect, i.e. that the effect of concern/risk is (partially) mediated by trust [79, 80]. Similarly, Knijnenburg and Kobsa showed that disclosure behavior in a demographics- and context-based recommender system was determined by trust in the company and concern/risk, with trust (partially) mediating the effect of concern/risk [182]. This suggests that trust itself can be built by reducing the perceived risk of information disclosure.
- Kobsa et al. show that trust can be a rational influence (rooted in risk and system-specific concerns) as well as a heuristic influence (rooted in the affect heuristic) [194].

Figure 2 shows the interplay between trust and concern/risk.

***Recommendations:** survey users, build trust, highlight relevance*

In sum, while privacy concerns are cause for hesitation in the unfettered collection of personal information, the TLA processors rely on the collection of such information to provide accurate personalization. This leads to a *privacy-personalization paradox*, i.e., a conflict between the user's perceived benefit of using TLA-based learning systems and their perceived concern regarding the disclosure of requisite information. The Federal Trade Committee suggests that addressing this paradox is essential for the success of personalized services [104]. Based on the analysis in this sub-section, we therefore make the following recommendations to ADL and other TLA performers:

- **Survey users**—As users' privacy behaviors are rooted system- and context-dependent perceptions of risk and relevance, it is important to continuously measure these perceptions as TLA-based systems and their data collection practices evolve. At design-time, TLA implementers should perform scenario-based multi-factorial experiments to quantify the effects of various data collection practices on perceived risk, perceived relevance, and disclosure behaviors. At deployment-time, they should conduct in-depth interview studies to uncover users' privacy-related attitudes and their potentially unanticipated antecedents and consequences. Moreover, implementers should perform controlled user experiments to detect potentially deleterious effects of heuristic decision practices on users' overall privacy concerns.

- **Build trust**—Our analysis shows that trust in the provider of a TLA-based system is an important factor in determining users' system-specific privacy concerns and perceived disclosure risk. Trust can be built heuristically through favorable name-brand associations, and it is thus important that all providers within the interconnected network of learning applications that constitute a TLA implementation are highly trusted by its users. So, while the TLA specifications may suggest an "open" learning platform, each implementation should ensure that the learning applications originate from trustworthy sources. Trust can also be built rationally by making sure that users have minimal privacy concerns at any point of time while using the system. TLA-based systems should therefore employ sensible data collection practices and a privacy by design philosophy from the onset. The suggestions in this document are instrumental in this endeavor.
- **Highlight relevance**—The privacy-personalization paradox and the privacy calculus suggest that far-reaching data collection practices are admissible, so long as the user understands the relevance of the data collection. At every step of the way, a TLA implementation should therefore highlight the potential improvements in content relevance, time savings, enjoyment and novelty that the collection of data can provide. Section 6 discusses different means of communicating relevance to the user. Moreover, the TLA activity providers should refrain from asking for information in situations in which the relevance is not readily apparent.

1.2 Elaboration likelihood

The previous subsection demonstrated that privacy decisions are influenced by both rational and heuristic antecedents, suggesting that users sometimes elaborate on their privacy-related behaviors, while at other times take decisional shortcuts. This subsection analyzes the factors that determine the relative importance of these two types of influences. We conclude that TLA-based systems should cater to both rational and heuristic decision-making practices, and that they can try to empower users to take more active control over their privacy (see Table 4).

Table 4: Recommendations regarding elaboration likelihood

Cater to Both Routes
– Provide detailed privacy control mechanisms for central route decision-making
– Provide sensible default settings to aid peripheral route decision-making
– Provide both heuristic and rational sources of trust
Empower Users
– Provide contextualized controls and comic-based information

Users’ decision practices range from heuristic to effortful

Outside the context of privacy, the decision-making literature has long realized that users’ decision practices range from heuristic to effortful, and have attempted to create “dual process theory” models that reconcile these different types of decision processes. One such model is the Elaboration Likelihood Model (ELM) [285, 287], which argues that people—to a varying extent—use two routes of processing: a central route (high elaboration) and a peripheral route (low elaboration):

- Central route processing is most in line with the privacy calculus, as it involves a more effortful elaboration process [403], in which users form their attitudes about a product based on a careful assessment of the most relevant available information [44, 226, 286], such as objective information about risks and benefits of disclosure [15, 227, 398, 407] and the availability of advanced privacy protection mechanisms [194].
- Peripheral route processing involves a more heuristic evaluation, which relies on superficial but easily accessible cues [285, 286, 324], such as website reputation [194, 313], ostensible privacy guarantees [398, 407], and design quality [24], which is in line with the heuristic accounts of privacy decision making discussed earlier [14, 155, 222].

Users’ motivation and ability influence their elaboration likelihood

ELM specifies two variables that determine the extent to which someone uses the central or peripheral route: motivation and ability. Motivation can be a personal characteristic (i.e., certain people are just generally more motivated to make privacy decisions), or it can depend on the situation. (i.e., certain people are more motivated to make privacy decisions when the dealing with a particular type of application or a particular type of data). Similarly, the ability to process presented information can be a personal trait (i.e., certain people have more privacy knowledge) or depend on situational factors (i.e. people are likely to make more elaborate privacy decisions when they have sufficient time and no distractions) [44, 286]. The following research provides evidence for the effect of motivation and ability on elaboration likelihood in privacy decision-making:

- Privacy researchers have used *general privacy concerns* as a measure of one’s motivation to engage in cognitive elaboration when making privacy-related decisions [24, 194, 407]. Privacy issues are of central importance to people with high levels of concern, and those individuals will thus be more motivated to make systematic use of issue-relevant cues and information.
- People with low levels of concern, on the other hand, will be more likely to use ostensive yet superficial cues in their evaluation process [24, 194, 407]. Indeed, privacy scholars have argued that some people use shortcuts and heuristics because they are not motivated to spend the effort needed to make an elaborate decision [61].

- In a similar vein, privacy researchers have used *privacy self-efficacy* (a person's belief in her cognitive resources required to cope with privacy-related problems [210]) as a measure of one's ability to engage in cognitive elaboration of privacy-related cues and information [24, 194, 407]. People who are equipped with more knowledge and resources are more able to engage in extensive elaboration.
- In contrast, People with low levels of self-efficacy will elaborate less and are more likely to rely on decisional shortcuts [194, 324]—cues that help them decide without needing to engage in cognitively elaborate processes. Indeed, privacy scholars have argued that some people use shortcuts and heuristics because they are incapable of making an elaborate decision [225, 231].

The effect of motivation and ability on elaboration likelihood is displayed in Figure 3.

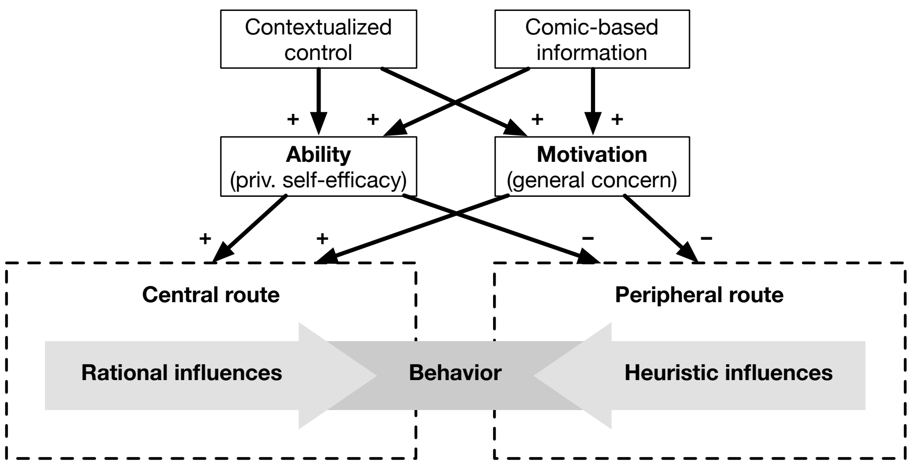

Figure 3: The effects of ability and motivation on central and peripheral route processing in privacy decision-making.

Motivation and ability can be instilled

We must acknowledge the harsh reality that users typically lack the motivation [61] and ability [225, 231] to make elaborate privacy decisions. So, while we should be careful not to overburden users with privacy control, it often serves to try and motivate and/or enable users to take a more central processing route in their privacy decision process.

One way to do this is to provide highly contextualized privacy controls, which may increase users' self-efficacy [183] (see Section 6.2). Another way to encourage central route processing of privacy-related information is the use of privacy comics [178] (see the Section 6.1). Figure 3 further displays the (potential) effects of contextualized control and comic-based information.

Recommendations: *cater to both routes, empower users*

In sum, TLA users may not only differ in the extent of information disclosure (as mentioned in the introduction of this section), but also in the way in which they make privacy decisions. The TLA will have to deal with users who are highly capable and motivated to make privacy-related decisions, and users for whom this is decidedly not the case. Based on the analysis in this sub-section, we therefore make the following recommendations to ADL and other TLA performers:


- **Cater to both routes**—It is important to realize that not all users will be able to make elaborate privacy decisions at all times. TLA-based systems should therefore cater to both central and peripheral route privacy decision-making. For example, a TLA providers can provide detailed privacy control mechanisms for central route decision-making, but also provide sensible default settings to aid peripheral route decision-making. Similarly, TLA providers should provide both heuristic and rational sources of trust.
- **Empower users**—While this is not always possible, it is better if users make rational rather than heuristic decisions. Hence, TLA-based systems should provide contextualized controls and comic-based information in an effort to increase users' motivation and ability to make more rational privacy-related decisions.

1.3 Communication style

Previous subsections covered users' privacy behaviors towards personalized systems. However, privacy in TLA-based systems extends beyond personalization; it is also relevant to the interpersonal ("social networking") aspects of these systems. Social networks typically provide a plethora of mechanisms to manage one's privacy beyond disclosure [89, 390], and research finds that users tend to employ a wide variety of strategies to limit their disclosure [272, 390]. This subsection describes these strategies, and how they can be supported in a TLA-based system (see Table 5).

Table 5: Recommendations regarding communication style

Tailor to Different Privacy Management Strategies
– Give Selective sharers the ability to selectively share information with specific apps and groups of people
– Give Self-Censors non-personalized mechanisms for selecting material, and restricted forms of sharing
– Allow Time Savers to opt out of active notifications and social features
– Give Privacy Maximizers all of the functionality described above
– Give privacy balancers mechanisms for curation, blocking, and avoiding direct interaction
– Make sure that Privacy Minimalists can maximally benefit from the adaptive and social functionalities of TLA
Tailor to Different Communication Styles
– Employ automatic social-network style sharing for FYI communicators
– Employ direct, chat-style interaction for non-FYI communicators
– Pay special attention to effects of integrating different communication styles within a single application

Social network users selectively apply privacy management strategies

Wisniewski et al. [389, 391] identified ten distinct privacy behaviors on Facebook: withholding basic or contact information, selective sharing through customized friend lists, blocking people, blocking apps or event invitations, restricting chat availability, limiting access to or visibility of one's Timeline/Wall, untagging or asking a friend to take down an unwanted photo or post, and altering one's News Feed. Moreover, they demonstrated that users use these strategies selectively. Specifically, they classified participants into six categories (see Figure 4) with distinct privacy management strategies:

- **Privacy Maximizers** use almost all of the available privacy features on the social network.
- **Self-Censors** use very few of the available privacy features, but primarily protect their privacy via the traditional method of withholding information.
- **Selective Sharers** share much more information, but they protect their privacy by sharing this content selectively, using custom friend lists.
- **Privacy Balancers** exhibit moderate levels of privacy management behaviors. Follow-up work shows that this class of SNS users contains both "informed balancers" (who carefully select the privacy mechanisms that suit their personal preferences) and "uninformed balancers" (who simply make do with the few mechanisms they are aware of).
- **Time Savers/Consumers** use Facebook primarily for passively consuming other people's posts, and take precautions to limit or avoid direct interaction with other users (e.g. through chat).
- **Privacy Minimalists** use only a few common privacy features, but are generally very open in their disclosure.

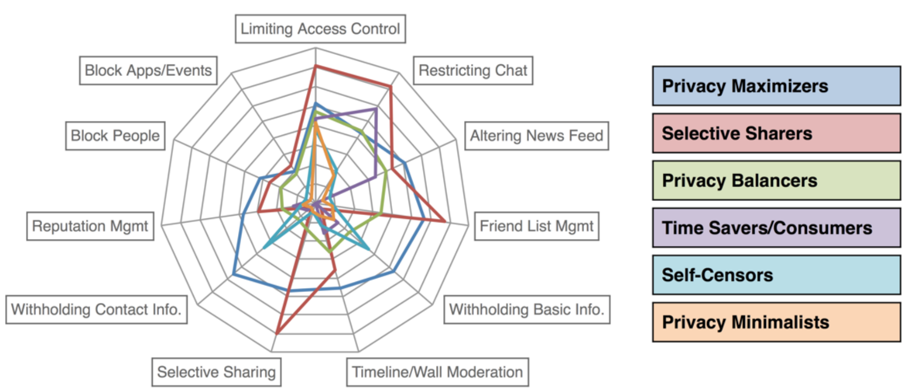

Figure 4: The six privacy management strategies uncovered by Wisniewski et al. [389, 391]
See http://www.usabart.nl/chart for an interactive version

Social network users use different communication styles

Page et al. [270] suggests that users choose their social network based on their preferred communication style. They argue that services that broadcast implicit social signals (e.g. location-sharing social networks) are predominantly used by users who are predisposed to "FYI (For Your Information) communication". FYI communicators prefer to keep in touch with others through posting and reading status updates, i.e., without actually having to interact with them. They tend to benefit from the implicit social interaction mechanisms provided by broadcast-based social network systems. People who are not FYI communicators, on the other hand, would rather call others, or otherwise interact with them in a more direct manner, rather than passively reading about them on social media. They thus tend to benefit more from systems that promote more direct interaction.

Recommendation: *tailor to different privacy management strategies and communication styles*

TLA users are likely to expect the system to have a wide variety of ways to communicate with other users and manage their social privacy, and that different users will use these mechanisms in different ways. Based on the analysis in this subsection, we therefore make the following recommendations to ADL and other TLA performers:

- **Tailor to different privacy management strategies**—In a recent paper [385], we explored how the different privacy management profiles uncovered by Wisniewski et al. can be applied to TLA-based systems. We refer the interested reader to the paper, and provide a summary of our analysis here:
    - Give Selective sharers the ability to curate and selectively share their personal information and training outcomes with specific applications and groups of people.
    - Give Self-Censors non-personalized mechanisms for the selection of learning material, and highly restricted forms of sharing learning outcomes.
    - Allow Time Savers to opt out of active notifications and social features.
    - Give Privacy Maximizers all of the functionality described above.
    - Give privacy balancers mechanisms for curation, blocking, and avoiding direct interaction.
    - Make sure that despite these mechanisms, Privacy Minimalists can maximally benefit from the adaptive and social functionalities of TLA.
- **Tailor to different communication styles**—As users with different communication styles prefer different mechanisms for interacting with each other, TLA-based systems should support these different mechanisms. Specifically, the TLA should employ automatic social-network style sharing for FYI communicators. These users will maximally benefit from the "social awareness" that results from seeing the implicit activity of other TLA

users. At the same time, TLA-based systems should employ direct, chat-style interaction for non-FYI communicators. This is in line with research that shows that learners are interested in seeing who is online and messaging them when they want to [399]. Since the communication styles of FYI and non-FYI communicators is at odds, user research should pay special attention to effects of integrating different communication styles within a single application.

2 Input data characteristics

Problem: What data should TLA collect? The TLA specifications envision a highly adaptive learner model that proactively mines and tracks a variety of information sources to provide personalized learning experiences. The goal of this learner model is to train employees on the job, adapting presented training modules to personal capabilities, mission requirements, and available time and other resources [98, 296]. Like many other adaptive systems, TLA-based systems thus rely on the collection of potentially privacy-sensitive information to provide its personalized learning services [105, 193, 289, 317, 342, 357, 366]. What kind of data should TLA-based systems collect, and what should they refrain from collecting?

Current state of the art: Widespread data collection envisioned. Currently, the TLA specification supplies APIs for the following kinds of user data [310, 329]:

- Learner runtime activity (Learner Experience Facts; xAPI): detailed tracking of specific learning activities
- Competency/mastery/evidence (Learner Profile; pAPI): data users' competencies, completed objectives, evidence, expiration dates, etc.
- Learning activity descriptions (Activity Index; iAPI): most importantly, "paradata" that includes user ratings, comments, and usage statistics about the learning activity)
- The learner's context (Context; cAPI): aspects of the learner's physical situation, computation equipment, schedule, etc.

Solution: Study the privacy implications of collecting various types of data. This section considers the privacy implications of the collection of these and other types of data by TLA-based systems. Notably, we add a discussion of social connections, physiological / psychological / medical data, skills and competences acquired outside the system, and users' learning ambitions. It also covers the privacy implications of the potential inferences that the TLA and/or its underlying learning systems can make based on this data. Formal questions about data ownership and the transmission of data to other parties are covered Sections 4 and 5, but are referred to in this section whenever relevant.

This section describes the privacy implications of collecting data about learner activity within and outside of the learning system, with specific attention to:

- Levels of identifiability
- Collection of various data types
- Inferences made based on data collection

Key findings and recommendations are presented in Table 6.

Table 6: Key findings regarding the input data characteristics

	Key Findings	Recommendations
Levels of identifiability (2.1)	– Full anonymity is impossible; pseudonymous users can be re-identified – De-identifying server data is still a good security practice – Pseudonymity has consequences for social interaction	– Use de-identification but don't rely on it for privacy purposes – Tailor users' identifiability based on the formality of the environment
Collection of various data types (2.2)	– Learner runtime activity is essential for operation, but can be sensitive – Continuously tracking the learner's context can create a digital panopticon – Social connection data can be used re-identify users – Detailed physiological data is sensitive, and tracking it may create an unwanted power dynamic – HIPAA prohibits the collection and sharing of medical data	– Allow users to correct/appeal competency data – Allow users to add outside skills – Allow users to submit their learning ambitions
Inferences made based on collected data (2.3)	– Users don't like incorrect predictions – Even correct predictions may at times be unwanted – Users are more than the sum of their data	– Allow users to correct and move beyond the personalized recommendations

2.1 Levels of identifiability

Of all the types of data that can be collected by TLA-based systems, Personally Identifiable Information (PII) deserves special attention, because the use and sharing of PII presents the risk of revealing the identity of users to other parties. PII can be defined as any information that could be used on its own or with a combination of other details to identify, contact or locate a person or to identify a person in context. The potentially classified nature of military identities makes identifiability a particularly important problem in military applications [312].

This subsection explores the limits of de-identification, and discusses the situations in which pseudonymity should be used or avoided (see Table 7).

Table 7: Recommendations regarding levels of identifiability

Use De-Identification
– Use—but do not rely on—de-identification for privacy purposes
Tailor Users' Identifiability
– Creative and (self-)evaluative environments should use pseudonymity – Formal and diplomatic settings should enforce a real name policy

Full anonymity is impossible; pseudonymous users can be re-identified

Beyond security concerns, requesting PII also induces privacy concerns. In a seminal paper, Ackerman et al. demonstrated that most users are uncomfortable disclosing PII, such as their social security number (99%), credit card number (97%), phone number (89%), address (56%) and full name (46%)—in contrast to information that does not personally identify them, such as their favorite snack (20%) or favorite TV show (18%) [3].

A possible mitigation of these privacy concerns is to allow users to remain fully anonymous. Anonymous interaction means that there are no persistent identifiers associated with the user. Fully anonymous interaction with TLA-based systems is difficult though, since the personalization functionality inherent in the TLA specifications crucially depends on the systems' ability to recognize the user across interactions [309]. More realistically, users can be allowed to interact with the system under a pseudonym [19, 196]. However, scholars have debated the value of de-identifying personal data stating that anonymized data may still be at risk of being re-identified [264], due to the high dimensionality and sparsity of the data typically collected by personalization learning systems [254]. In this sense, the combination of various data that are not directly personally identifiable (e.g. a combination of the user's favorite snacks, TV shows, and other preferences) can effectively be used to identify them in a dataset. An overview of these mitigation techniques is presented in Figure 5.

This re-identification threat can be reduced by not giving others access to any of the user data. Note, though, that even without such access, it may be possible for a third party to make inferences based on the output of the system. Calandrino et al. [45] employ such a "reverse re-identification scheme" using a fake user accounts that are similar to the account of a target user. An adversary using this technique can use the recommendations provided to the fake accounts to isolate the target user's data. A means to overcome this problem is differential privacy, a privacy model that inserts carefully calibrated noise into the user profile computation. The noise masks the influence that any difference in a particular record could have on the outcome of the computation [230, 241, 300, 408].

Interestingly, while pseudonyms and anonymity may reduce privacy concerns, many systems increasingly require users to use their real name [409] (presumably to combat the increasing number of fake accounts), and even some governments require their citizens to verify their real name before signing up on certain popular websites (presumably to counter rumors and defamation of politicians during the election cycle) [58]. A learning system that may be used to make deployment and promotion decisions may similarly require users to authenticate with their PII to prevent and combat fraudulent use (e.g. cheating). In this case, it is good practice to hash the requisite PII.

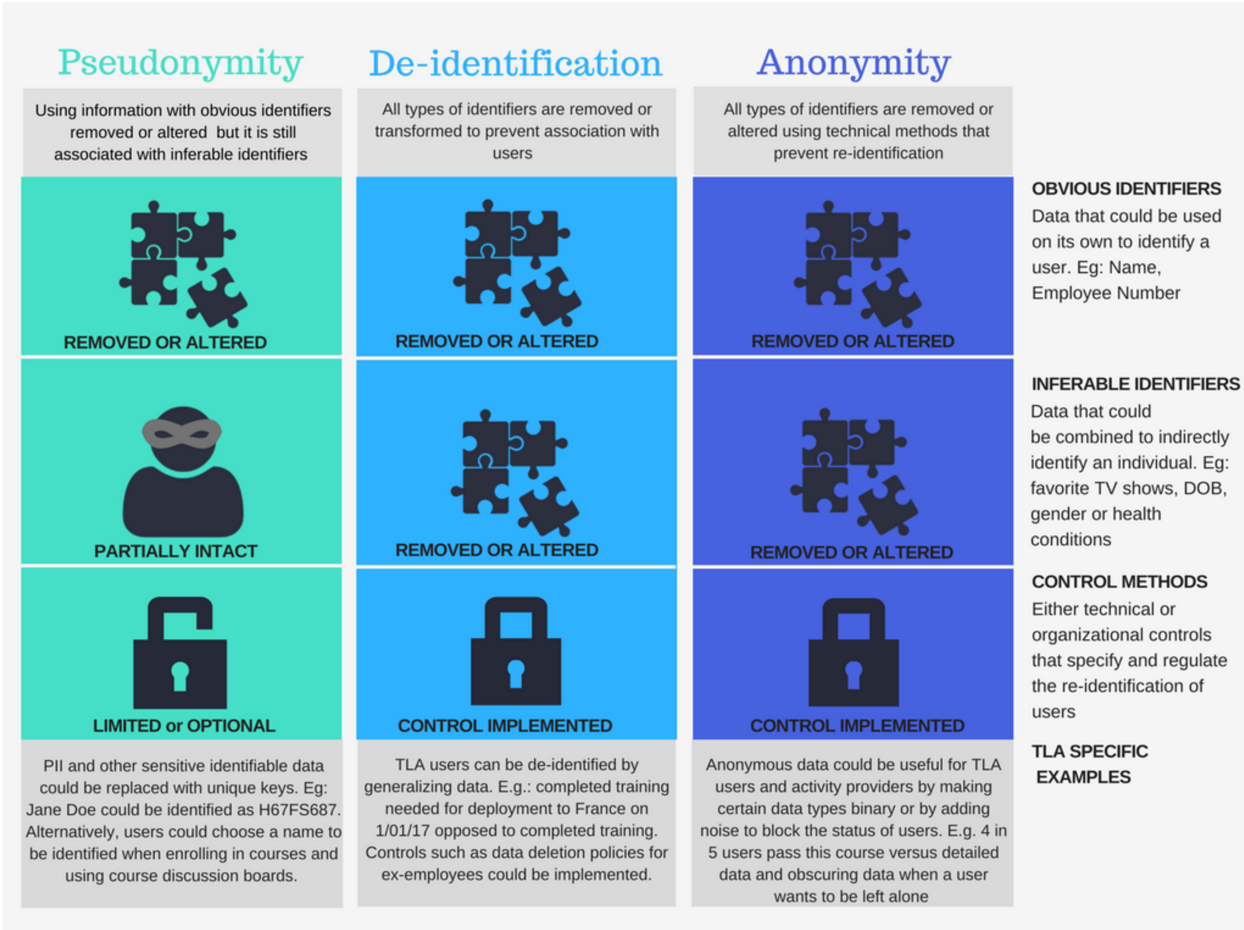

Figure 5: Identification techniques and their application to TLA

De-identifying server data is still a good security practice

Although pseudonyms do not guarantee that users can never be re-identified, researchers have argued that de-identification of server data can be a valuable but not foolproof method for minimizing privacy and security risks. For instance, Masiello and Whitten argued that while anonymized information will always carry some risk of re-identification, most of the privacy risks occur only when there is *certainty* in re-identification [234]. Removing or hashing PII and introducing differential privacy introduces uncertainty in re-identification, thereby removing the most prominent privacy risks.

These mechanisms could for instance be helpful to in the instance of a data breach: Personal information such as upcoming lessons could possibly reveal where a soldier would be deployed next, which would be a privacy risk. But if the soldier's PII is removed or hashed, and if the stored lesson data contains a certain amount of noise, then the identity of the soldier and his/her upcoming deployments cannot be determined with certainty. Targeted attacks to uncover the soldier's PII may still succeed, but the de-identification practice prevents the data breach from automatically compromising the PII of all the data subjects involved in the breach. Routinely de-identifying data could also prevent rouge employees from having direct and effortless access to data that exposes the identity of users.

Pseudonymity has consequences for social interaction

Anonymity and pseudonymity not only influence users' privacy; research also suggests that runtime identifiability (i.e., whether users get to know the identity of other users they interact with while using the "social" features of a system) has an influence on user behavior [316]. Specifically, in creative environments (e.g. creative thinking exercises, brainstorming activities), the absence of a name allows users to produce content more freely, which increases creativity [57, 341], reduces conformity and inhibition [60], and increases the opportunity for intimacy and the sharing of secrets [371]. The latter makes anonymity and pseudonymity useful for self- and peer evaluation exercises.

Unfortunately, pseudonymity also induces a certain dissociation between the members of an online community [341], which is obviously bad for team building exercises and other team activities. Real name requirements can avoid such problems, and have also been shown to reduce profanity and anti-normative expressions in online social networks, especially among more-frequently participating users [58]. Formal and diplomatic settings may thus benefit from a real name policy.

Recommendation: *use de-identification; tailor users' identifiability*

Due to the intricate relationship between privacy, security, and the social consequences of pseudonymity and anonymity, it is difficult to make a simple recommendation regarding the identifiability of TLA users. Based on the analysis in this subsection, we can make the following recommendations to ADL and other TLA performers:

- **Use de-identification**—Recent re-identification threats show that simply removing or encrypting all PII in a system does not guarantee that users cannot be identified. However, de-identifying server data makes identification uncertain, which removes the most prominent privacy risks. The data storage protocols of the TLA should thus use—but do not rely on—de-identification for privacy purposes, either by not collecting any PII at all, or by removing, encrypting or securely key-coding the PII that the system is required to collected (e.g. for authentication purposes).
- **Tailor users' identifiability**—Runtime identifiability has an influence on user behavior, with anonymity, pseudonymity and real name policies each having both desirable and undesirable social consequences. The social components of the TLA should thus use anonymity, pseudonymity, and real name policies selectively, where socially useful and appropriate. Specifically, creative and (self-)evaluative environments should use

pseudonymity, because it increases creativity and earnest. Conversely, formal and diplomatic settings should enforce a real name policy, because it reduces profanity and anti-normative expressions.

2.2 Collection of various data types

TLA-based systems collect a wide array of data that is used to offer learners adaptive guidance and to help teachers and institutions manage their learning ecosystem. Long-term persistent data tracking allows TLA-based systems personalize learning, build competency models, coach the learner, and discover helpful insights about its user base. Detailed in-situ tracking enables the development of increasingly smart learning activities and personal assistants that guide, coach, assess, and give feedback to learners. The collected data might even help to identify cognitive states and traits that contribute to a large number of competencies and thereby offer new generalizations of existing methods to teach and assess [310].

This subsection analyzes the privacy implications of the various data types that TLA-based systems are currently envisioned to collect and/or may collect in future iterations of the specifications (see Table 8).

Table 8: Recommendations regarding the collection of various data types

Carefully Protect Learner Runtime Activity
– Protect learner runtime activity using a combination of strict access control, encryption, de-identification and obfuscation – Provide easy-to-use notice and control mechanisms for users to control the boundary between leisure and learning – Test the mechanisms presented in Figure 4
Treat Social Connection Data as PII
– Protect social connection data as if it were personally identifiable information
Be Careful Not to Create a Panopticon
– Reduce unfettered context tracking to prevent the creation of a digital panopticon
Keep Some Data Local
– Process it and use it locally
Allow Users to Add Outside Skills
– Allow users to selectively add skills and competences acquired outside the system
Allow Users to Submit Their Learning Ambitions
– Provide a comprehensive manual self-reporting system – Provide a way to test or otherwise provide evidence for skills and competences

Learner runtime activity is essential for operation, but can be sensitive

Tracking learner runtime activity is essential for TLA-based systems to enable personalized learning with smart learning activities. Adaptive learning modules can use runtime learning activity to track users' abilities as they learn, and adapt the topic and difficulty level of the training to the user's current knowledge level and pace of learning. Moreover, the analysis of highly granular learning behavior arguably allows training department managers to glean superior detailed insights about users' overall learning progress, the effectiveness of specific training modules, and the capabilities available in their division [310]. For these reasons, users should arguably not be allowed to opt out of tracking their runtime activity, as doing so would undermine the very purpose of the TLA specifications.

Runtime activity may include very sensitive data, though. For example, detailed data from cyber range practices may reveal battlefield tactics, and training data from top diplomats may reveal weaknesses that can be exploited in negotiations. To alleviate users' (and supervisors') privacy concerns, these data thus need to be protected by a combination of strict access control, encryption, de-identification and obfuscation.

A complication in the tracking of runtime learner activity is the fact that tracking may occur outside the traditional learning channels; TLA-based systems can give users credit for learning other activities, such as playing a tactical game, or reading a relevant blog post. Such non-traditional learning activities blur the boundaries between leisure and learning, and require runtime learner activity tracking to be in an "always on" mode. Consequently, users may find their real-world activities tracked, which arguably more privacy-invasive (and difficult to control) than the tracking of in-system behaviors [261], especially when it happens in a pervasive and unobtrusive manner. The mining and tracking activities may also be regulated by government privacy regulations [83].

Given this fluid boundary between learning activities and real-world activities, the TLA should give the user easy-to-use notice and control mechanisms to control this boundary. Figure 6 displays five of such mechanisms, at different levels of automation:

- **Control center widget**—This mechanism uses an opt-in paradigm, as it requires users to activate the learning activity before starting it. The widget could potentially provide a time-out functionality that automatically disables the tracking after a certain time window has passed, and the widget could also automatically turn on during business hours, and off outside business hours (these options are similar to Apple's Night Shift functionality). As an opt-in mechanism, the widget is the most private mechanism, but also the most error prone: users may forget to turn on the learning activity tracking, or may not be aware that something counts as "learning". Compared to other mechanisms, this is the only one that does not require any background tracking.
- **Opt-in toast**—This mechanism also uses an opt-in mechanism. But it monitors users' activity in the background to detect potential learning activities. When a learning activity

is detected, the notification appears, allowing the user to start the learner activity tracking with a single tap (alternatively, the user can ignore the pop-down toast to avoid tracking). The observed background activity used by this method (as well as the three remaining methods) does not have to be permanently stored, and it should be made clear to the user that this data is used for observation only. If client-side methods (see Section 4) are used to monitor background activity, then this data does not need to be transmitted to a server at all.

- **Opt-out toast**—This mechanism is similar to the opt-in toast, but uses an opt-out mechanism: learner activity tracking is automatically started unless the user cancels the tracking with a single tap. This mechanism is slightly more privacy sensitive, as users may overlook the notification.
- **Pop-up message**—This mechanism uses a forced-choice paradigm, as it forces users to confirm whether an activity is considered a learning activity. The automated pop-up overcomes the problem of forgetting to turn on the tracking, but the pop-up itself can be perceived as intrusive.
- **Recording banner**—This mechanism shows the user a pervasive banner to indicate that the system is recording the user's learning activity. Tapping the banner brings the user to a screen where the recording can be ended and/or adjusted. This mechanism can be combined with any of the other mechanisms.

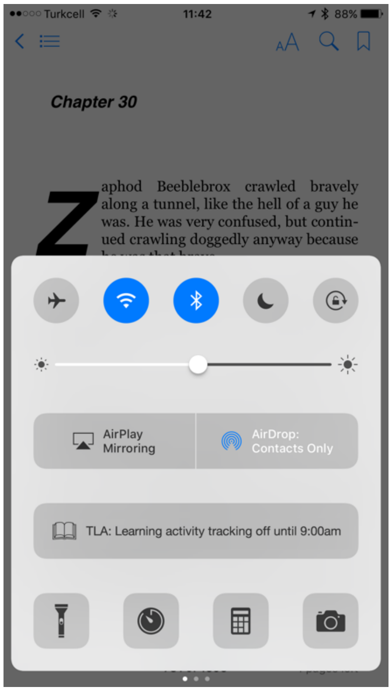

Control center widget

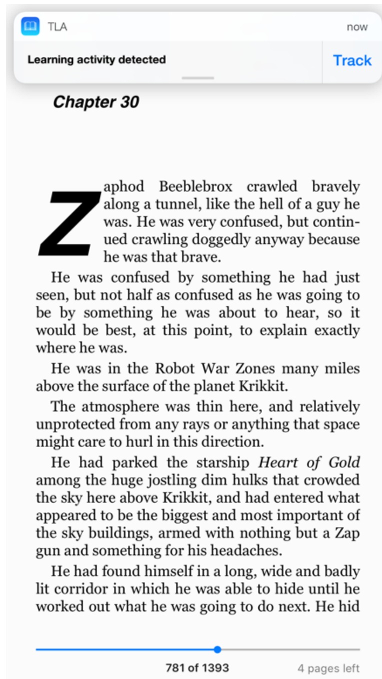

Opt-in toast

Opt-out toast

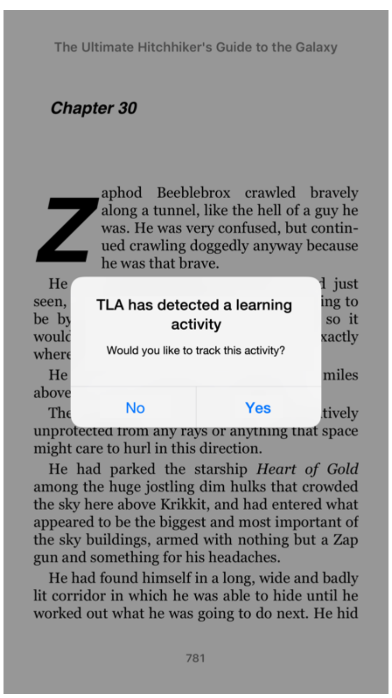

Pop-up message

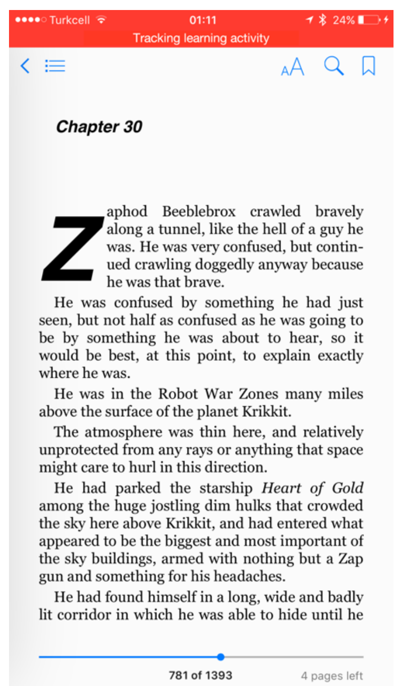

Recording banner

Figure 6: Five candidate mechanisms to control the tracking of runtime learner activity

Users should be allowed to correct/appeal competency data

Competency data is essential for research, deployment, promotion decisions. The ethical and privacy implications of using data for this purpose is discussed in Section 4. Competency data can be regarded as an aggregate form of learner runtime activity data, i.e., it only captures the outcome of a learning activity (and in some cases the evidence for this outcome). Most people feel comfortable to share such information, especially in situations where it is directly relevant [183].

Beyond that, it may also contain subjective evaluations of the user’s competence, e.g. assessments by peers, superiors, or training officers. This includes the identification of personality traits and other psychological factors [310]. The inclusion of such a subjective component requires a mechanism for users to appeal decisions or evaluations that they deem unfair or incorrect. In fact, even for objective competency data the correctness cannot always be guaranteed, so an appeals process may also be instrumental in correcting glitches in the recording of objective competencies.

Learning activity descriptions are not very sensitive

Learning activity descriptions are arguably the least sensitive type of data that TLA-based systems can collect. The main user-generated part of this data—the “paradata” that contains user feedback regarding the Learning Activities—is a type of “preference data”. Preference data can be collected in various different ways, including question-answering [257], attribute weighing [187, 189, 190], and item-based feedback—the latter can be subdivided into implicit feedback [187, 191, 299], rating [108, 173, 333], and example critiquing [56, 239, 294]. Users usually do not mind providing preference-related feedback to a system [366], but direct preference measurement is not always ideal. For example, users may not always be motivated to give explicit preference feedback [192], and their feedback may not always accurately reflect their preferences [13], especially when they are novices in the recommendation domain [187, 189]. Implicit feedback, on the other hand is easier to gather, but can result in overspecialization [191].

Research [192] shows that privacy concerns can reduce users’ intention to give explicit preference feedback, but that this intention will increase with choice satisfaction and system effectiveness. In other words, a responsive adaptive system can overcome privacy concerns and encourage users to contribute preference information. Privacy concerns regarding implicit tracking of preferences (e.g. by monitoring recommendation browsing behavior) are also low: recent research found that between 80% [194] and 87% [182] of users allow this type of tracking.

Continuously tracking the learner's context can create a digital panopticon

The field of context-aware recommender systems has shown that context information (such as a user's interaction with other users, location, calendar events, etc.) can be used to improve the accuracy of predictions about users' tastes and preferences, which could improve the personalized presentation of learning activities [10]. Context can also be used to adapt the presentation of the current learning activity to the user's situation [100, 310, 399].

One benefit of context data is that it can be collected continuously and unobtrusively (compared to e.g. users' demographics or preferences, which may have to be explicitly elicited) [215]. This is at the same time also the biggest problem of context data, because it has been shown that users are more concerned about personal information that is collected automatically compared to manually provided information [182, 184, 194]. Particularly, users fear that the system could make incorrect inferences about their situation, or use the collected data for unintended purposes. While users are relatively okay with an adaptive system tracking their location (85%), phone model (85%), and general app usage (82%), they are much less willing to have an app track their Web browsing (48%), email messages (37%) and mobile credit card purchases (20%) [182].

An important reason for users' worries is that the system may make incorrect inferences based on the data [182], or that the system may reveal embarrassing contextual information to other users of the system [272]. Another problem with context data is that it typically concerns behavior that is not directly representative of users' tastes and preferences, which makes it difficult for the system to highlight its relevance. Since context tracking is not essential for the correct operation of the core TLA-based personalization services, users should be allowed to opt out of continuous context tracking (or, alternatively, context tracking should be disabled by default, allowing users to opt in for a better personalized experience; see Section 6.3).

The continuous tracking of the learner's context makes it more difficult for users to lie about their activities and whereabouts [126]. While increasing honesty may seem like a desirable goal, studies in computer-mediated interactions show that users sometimes lie as a privacy preservation tactic [125]. A user may for example tell a friend that she has fallen ill, rather than telling the friend that she does not want to go out with her that evening [271]. Researchers recommend that social information systems allow users to make white lies; a functionality that has been dubbed "plausible deniability" [17, 36, 213]. Page et al. [271] demonstrate that in systems that create a "panopticon" (cf. [22, 298]) by pervasively tracking, the practice of lying indeed *increases* the privacy concerns of the liar. The problem of lying in information systems is a complex issue that involves balancing the opportunity for users to lie with the moral responsibility of creating honest digital experiences. TLA developers need to be acutely aware of this issue, since TLA-based systems may expose—or further exacerbate—users' lies.

Social connection data can be used re-identify users

Access to social connection data allows TLA learning applications to create powerful collaborative or competitive social learning experiences. The privacy implications of such experiences are discussed in Section 5.2. Here we consider the privacy of the social connection data itself.

In an increasingly networked world, it is important to realize that social connection data can reveal a lot of information about a user. Indeed, social connection data can be used to re-identify anonymous users [253], and "neighborhood attacks" can be used to infer unrevealed traits about a user from friends' traits [404, 405].

Detailed physiological data is sensitive, and tracking it may create an unwanted power dynamic

Leveraging novel sensing technologies that are increasingly incorporated into consumer devices, TLA-based systems could potentially have continuous access to a wide range of physiological metrics, such as sleep patterns, weight, physical exertion, and heart rate. Detailed runtime physiological tracking can be used to make real-time adjustments to combat and fitness training routines, pushing users to—and beyond—their personal limits. Moreover, in aggregate, such data allow TLA-based systems to track the health of its users, and recommend physical training programs that match their current physical condition. Knowing users' overall health and fitness also supports supervisors' deployment decisions.

An increasingly interesting use of physiological sensing technology is *biometric authentication* [76]. TLA based applications could use this method to provide access to authorized personnel without the need for passwords. Note that some forms of biometric authorization, such as face scans [75], fingerprints [149], and iris scans [247], can be compromised with the right tools.

It remains a question whether TLA users will feel comfortable with having the system tracking their biometrics and physiological activity. Studies show that relationships between different types of physiological data can give very detailed insights into the user's life [122], so these tracking applications and wearable devices are rapidly becoming an important source of privacy and security leaks [26, 121]. Insights into the user's sexual activity and bodily functions can for example be gleaned from this type of data [26]; users may not want such information to be available to their employers. Aggregation reduces these problems. We therefore suggest a hybrid solution where detailed runtime physiological data is used for adaptations at the client side only, and aggregated before transferring it to the server, where it is used for tracking users' general health and physical condition (see also Section 4).

Surveying research on pervasive tracking in elderly care, Alemdar and Ersoy find that the use of sensors creates an interesting power dynamic: while the monitoring benefits the elderly as well as the people taking care of them, only the elderly themselves suffer the downsides of constantly being monitored [11]. Analogously, the pervasive tracking of physiological activity may create an unwanted power dynamic between users and their supervisors. To wit, the military already places important restrictions on soldier's activities, and strong demands their physique, for the purpose of combat preparedness [372]. While continuous tracking may support a soldier in pushing their boundaries and striving for perfection, it also creates an implicit expectation of 24/7 commitment to such goals, which can be a source of unwanted pressure which can negatively impact the employee and their family [382]. This can be countered by increasing employee choice and flexibility over work demands.

HIPAA prohibits the collection and sharing of medical data

As an extension beyond physiological data, TLA researchers could help identify, prevent, and or mitigate health risks. Platforms such as Apple's HealthKit and WebMD demonstrate that basic online diagnosis of health issues is becoming more prevalent today [99, 109]. This could be a useful feature for users working in areas where contagious diseases are common, or in remote locations without access to medical care. TLA-based learning systems would also benefit from having applications that educate users about personalized preventative health practices. Such applications could also obtain input from the user to help them identify any ailments they may be suffering from [82, 249]. Finally, such functionality could reduce some of the immense pressure on the Veterans Affairs to care for veterans' medical health.

As the HIPAA privacy law [15] prohibits the sharing of medical information with employers and other third parties, such applications should be accessed only by the user and should under no circumstances be shared with their employer. Client-side methods could be used to implement this requirement (see also Section 4).

Allow users to add skills and competences acquired outside the system

TLA users may have skills and competences that were not acquired under the auspices of TLA, either because these skills were acquired before they became TLA users, or because they were acquired through learning applications that are not part of TLA. Users may want to "import" these skills and competences into their TLA-based system to demonstrate their diverse skillset to their employer (e.g. for promotion purposes). Likewise, for TLA-based systems it is also useful to be aware of these skills and competences. Indeed, the design rationale for TLA envisions it to be fully backwards compatible with traditional LMS-based learning environments [310].

Note that users might not want to input *all* the skills they have acquired outside TLA. For example, users may fear that a skill they acquired at a previous job but that is not aligned with their current interests may inadvertently cause their supervisor to change their current job to make use of these other skills. One way to resolve this problem is to allow users to *selectively* add previously-acquired skills. Another resolution lies in the ethical dilemma of the tradeoff between organizational needs and users' personal interests. This dilemma is further discussed in Section 5.3.

Allow users to submit their learning ambitions

As adaptive systems become increasingly more common, there exists a fear of "over-automation" and loss of control among their users [208, 275]. In TLA-based adaptive learning systems, this could cause a loss of perceived ownership over the user's learning process—a situation that may reduce users' motivation and learning effectiveness. In response to this problem, recent research has suggested to give users of adaptive decision support systems a more meaningful role in the decision-making process [188]. TLA-based systems could promote a similar philosophy, by allowing users to submit their learning ambitions (e.g. whether the user would like to build a specific specialization, or transition into a management position), and take these into account in providing learning recommendations. Arguably, a system that acknowledges these ambitions can leverage users' intrinsic motivation to learn, which is much more powerful mechanism than extrinsic motivation (e.g. recognition, promotion) [77]. This idea requires an ethical discussion of the tradeoff between these ambitions and organizational needs, which is provided in Section 5.3).

***Recommendation:** treat each data type in an appropriate manner*

TLA-based systems may collect a wide variety of input data to provide personalized learning experiences. Each type of data should carefully be considered in an appropriate manner. Based on the analysis in this subsection, we can make the following specific recommendations to ADL and other TLA performers:


- **Carefully protect learner runtime activity**—Learner runtime activity can reveal a lot of sensitive details about users, and compromise security. It is therefore important to protect learner runtime activity using a combination of strict access control, encryption, de-identification and obfuscation. This learner activity may overlap with leisure activity, so the TLA should provide easy-to-use notice and control mechanisms for users to control the boundary between leisure and learning. To this goal, we recommend conducting a user experiment to test the mechanisms presented in Figure 6.
- **Treat social connection data as PII**—Social connection data can be used to re-identify users. The TLA should thus protect social connection data as if it were personally identifiable information.
- **Be careful not to create a panopticon**—Context tracking and pervasive monitoring of physiological data may improve personalization, but it also restricts user freedom. The TLA should reduce unfettered context tracking to prevent the creation of a digital panopticon.

- **Keep some data local**—Fine-grained physiological data can be very revealing about users' most personal activities, so users may not want to share it. HIPAA prevents medical data from being shared. If learning applications want to take advantage of this data, they should process and use it locally (i.e., on the user's device).
- **Allow users to add outside skills**—Not all skills and competences are acquired within the context of the TLA. The system may still benefit from knowing about these skills and competences. To the extent that users want to share this data, the TLA should allow users to selectively add skills and competences acquired outside the system.
- **Allow users to submit their learning ambitions**— A system that acknowledges these users' ambitions can leverage their intrinsic motivation. The TLA should thus provide a comprehensive manual self-reporting system. It should possibly also provide a way to test or otherwise provide evidence for skills and competences the user claims to have acquired outside the TLA.

2.3 Inferences made based on collected data

As the previous subsection has already alluded to, the privacy implications of data collection in personalized systems extend beyond the collected data itself, to the potential (and actual) inferences the TLA Processors are able to make based on the combination of different data sources. Users are intuitively aware of this threat of aggregation, and indeed seem to get increasingly wary as disclosures accumulate within a given system [27, 182, 202].

This subsection analyzes the impact of automatic inferences on users' privacy. Our main recommendation is to allow users to scrutinize and correct inferences made by the personalized learning system, and to give them a more active role in the process of curating learning activities (see Table 9).

Table 9: Recommendations regarding inferences made based on collected data

Allow Scrutiny and Corrections
– Give users the opportunity to scrutinize and correct potential mistakes
Build Trust
– Allow users to venture beyond the personalized recommendations
– Give TLA users a more meaningful role in the decision-making process

Users don't like incorrect predictions

First and foremost, users get annoyed when personalized systems make an incorrect prediction or inference about them [326]. One famous example of this involved a man whose TiVo (a digital video recorder with a built-in recommender system) started exclusively recording TV shows with Gay themes [402]—arguably after "overfitting" a previously encountered information pattern [299]. In the TLA incorrect predictions can lead to the system recommending a training at the wrong difficulty level (which in some cases may lead to physical injuries), presenting a training in a modality that does not match the user's context (e.g. presenting video-based learning material while the user is driving, or an audiobook in a noisy environment), or presenting training material that does not match the user's preferences or learning goals (leading to boredom and wasted time). Moreover, incorrectly stereotyped recommendations can lead to embarrassing situations when other people (e.g. team members or classmates) get to observe these recommendations. Users typically have an urge to correct and/or compensate for mistaken predictions [65]. In effect, researchers suggest that users should have the opportunity to scrutinize [163] and correct [103] potential mistakes.

Even correct predictions may at times be unwanted

Even when inferences are correct, they may not always be in the user's best interest. Some of the recommendations that the TLA Processors will be able to make may simply be perceived as "creepy" [318, 348]. For example, Phelan et al. find that Facebook users intuitively dislike the fact that their data is being tracked, even if they have no rational objections against it [288]. This intuitive dislike may reduce users' trust in the system.

Moreover, the use of data that to most users deem innocuous in isolation (e.g. preferences [3]) may in aggregate result in inferences about personality or lifestyle that the user is uncomfortable disclosing. For example, it has been shown to be possible to predict someone's sexual orientation based on 5-10 Facebook likes [200]. A related fear is that such inferences might transpire in the user's recommendations, which, if consumed in the presence of others, may "out" the user. Examples of this are the secretly pregnant teenager who received personalized Target baby advertisement brochures at her parents' address [84], and the closeted lesbian mom whose private Netflix viewing history was re-identified using her public IMDB profile [256].

Another problem is that stereotypical inferences could result in discriminatory practices. For example, the TLA specifications envision data mining capabilities that might be able to identify cognitive states and traits that contribute to a large number of competencies, thereby offering new generalizations of existing methods to teach and assess [310]. The fear is that such data could be used in a negative sense as well. For example, Schneider et al. ask "What if that Soldier misses out on a promotion, key assignment, award, or superior evaluation because the algorithm has determined that he is at risk for suicide-related behavior? Is this outcome fair? Does it violate the Soldier's right to privacy? Will uninformed use of this data actually increase the Soldier's risk of self- harm?" [312]. Section 5.3 provides an ethical discussion of the use of such data, which is a first step in answering these questions.

Users are more than the sum of their data

It is important for realize that predictions made by TLA-based systems may never be perfect, because users are more than the sum of their data. This means that the personalization aspects of the TLA specifications should not be taken too far, and that users should always be able to venture beyond the personally recommended content. Indeed, researchers have argued that heavily filtered content may isolate us from a diversity of viewpoints, content, and experiences, and thus make us less likely to discover and learn new things (a phenomenon known as the "Filter Bubble" [275]). The Filter Bubble can be thought of as a privacy threat because it intrudes upon our ability to experience the world from an unbiased perspective [330]. When users are encouraged to follow the recommendations of their adaptive learning system, stereotyping can even create a "positive feedback loop" [208], where users increasingly try to fit the stereotype. This leads to a very worrying concern that recommender algorithms may gradually replace human creativity and understanding; a scenario reminiscent of the seminal privacy novel *1984* [267]. As mentioned earlier, a good remedy against this concern is to give users of adaptive decision support systems a more meaningful role in the decision-making process [188].

> ***Recommendation:*** *Allow users to correct and move beyond the personalized recommendations*

In sum, users may not always welcome the inferences their adaptive learning system may make based on the available input data, regardless of whether these inferences are correct or not. TLA-based personalized system should thus recognize the limitations of personalization, and allow users to more actively engage with the system and its content. Based on the analysis in this subsection, we can make the following recommendations to ADL and other TLA performers:


- **Allow scrutiny and corrections**—Providing users a personalized learning experiences will not be without problems. The TLA Processors may make incorrect predictions, or predictions that users may be uncomfortable with. TLA-based systems should therefore give users the opportunity to scrutinize and correct potential mistakes in their predictions.
- **Support self-actualization**—As people tend benefit from exploring their interests beyond the beaten path, TLA-based systems should allow users to venture beyond the personalized recommendations. One way to support this is to give TLA users a more meaningful role in the decision-making process.

3 Output characteristics

Problem: How should TLA present adaptations? As TLA envisions both adaptations about apps (meta-adaptations) as well as within apps (macro- and micro-adaptations), the TLA Providers should coordinate the presentation of adaptations to create a consistent experience throughout the learning architecture. What should these adaptations look like, and how should they be timed, so that users experience minimal intrusion from these adaptations?

Current state of the art: Very little work on adaptation presentation. The TLA architecture specifies three types of adaptations [329]:

- Meta-adaptations are individualized recommendations to switch from one Learner Activity to another that are based on the learner's specific needs and progress.
- Macro-adaptations determine the next learning activity inside a single activity provider.
- Micro-adaptations adapt learning content within a single learning activity.

For users, such adaptations make it easier to find Learning Activities and activity providers that fit their current needs, and help them explore and expand their learning interests. For User Facing Application developers, good adaptations result in trust and continued use. To be efficient and useful for both parties, the adaptation mechanism needs to be accurate without being intrusive or inconvenient. However, very little work to date has considered the presentation of adaptations as a main focus of user-centric research [176, 179].

Solution: Study the timing and presentation of adaptations. The previous section discussed the intrusiveness of various types of input data; this section describes factors that impact the effectiveness of recommendations and adaptations within and across learning activities, including:

- Adaptation and presentation mechanisms
- Output modalities and devices

We discuss the convenience (or potential inconvenience) caused by the recommendations themselves. Specifically, we argue that adaptations should be:

- Carefully timed, potentially based on contextual input regarding the interruptability of the user.
- Carefully explained, without being overly persuasive.
- Conservative in how much information they provide, limiting the potential for leaking classified information.

Key findings and recommendations are presented in Table 10.

Table 10: Key findings regarding the output characteristics

	Key Findings	Recommendations
Adaptation and presentation mechanisms (3.1)	– Adaptations can serve multiple purposes – Recommendations can be "pulled" by the user or "pushed" to the user – Explanations can persuade users to follow recommendations	– Provide multi-purpose adaptations – Carefully time pushed recommendations – Explain recommendations without being overly persuasive
Output modalities and devices (3.2)	– Smartphones are ideal for Just-In-Time learning, but can be distracting – Wearables are less disruptive, but may feel more intrusive – Notifications can leak personal information	– Do not disturb the user – Prevent leaking information in social settings

3.1 Adaptation and presentation mechanisms

Making correct inferences about the user is only half the job of an adaptive system: those inferences need to be turned into actionable adaptations, and presented to the user in a way that is impactful but not intrusive or overly persuasive.

This subsection deals with the best presentation mechanisms for multi-purpose adaptations. We argue that adaptations should be presented timely, explained carefully, and that apps should avoid pressuring users into engaging into learning activities they do not want to engage in (see Table 11).

Table 11: Recommendations regarding adaptation and presentation mechanisms

Provide Multi-Purpose Adaptations
– Carefully balance different adaptation purposes
– Allow users to weight adaptation purposes
Carefully Time Pushed Recommendations
– Use (client-side) context-awareness to detect the optimal time to make a recommendation
– Provide users with timely feedback about their learning performance
Explain Recommendations Without Being Overly Persuasive
– Explain the implemented adaptations to the users
– Avoid pressuring the users into accepting adaptations that they do not want to accept
– Give users various options to choose from and help them understand the value of each option

Adaptations can serve multiple purposes

Recommender systems typically model their recommendations after users' predicted behaviors or preferences. In a learning environment, adaptations can be based on suggested lesson sequences, user goals, team needs, or mission objectives. In other words, the adaptations serve multiple purposes, both from the users' viewpoint as well as the providers' viewpoint [152]. In line with Jannach and Adomavicius' "Purposeful Evaluation Framework" [152], Table 12 illustrates the learning tasks that TLA-based adaptations can support.

The different learning purposes are not always aligned. Promoting behavioral change means breaking with users' current preferences, and creating group consensus may mean individual users have to compromise.

Even focusing on the users' preferences, it is possible that the users' current behaviors and future aspirations are not aligned [88], or that their current preferences are uninformed due to the limited viewpoint that their "filter bubble" provides [208, 275]. As such, adaptations could focus on allowing users to explore and understand their own learning preferences, rather than replacing this process algorithmically [188]. Giving users an active role in deciding what to learn reduces their dependency on the TLA. Moreover, it will likely result in a more thorough understanding of a user's learning preferences, something that is very useful since users' learning preferences are typically not singular, but rather multi-faceted and only loosely connected [151].

Table 12: User tasks TLA-based adaptations can support

Item/User Task	Description	Generic Recommender	TLA Recommender
Exploration	Proposing things that vary from current preferences	Proposing a new song from a genre the user usually does not listen to	Proposing a course on a topic the user is not yet familiar with
Recommended Sequence	Recommending the best sequence of items	Recommending a sequence of books	Recommending a daily "couch to 5K" training sequence
Finding a better fit (Goal Oriented)	Suggesting things that better aligns with users' goals	Suggesting a movie based on the plot keywords of previous choices	Suggesting a more detailed security course for a security expert
Promoting behavioral change	Making a suggestion with the purpose of changing users' behavior	Suggesting the user to increase their workout goals	Suggesting the user to take a more challenging course
Task Specific	Supporting users while they complete other tasks	Suggesting alternative shirts to the one being displayed	Recording an ad-hoc Learning Activity
Novelty	Recommending novel items	Recommending breaking news articles	Recommending a new training program
Context Specific	Recommending different learning styles or methods	Recommending a restaurant nearby	Recommending an audio-based training when the user is driving

Recommendations can be "pulled" by the user or "pushed" to the user

Traditionally, recommendations are requested by the user, e.g. via search or navigation [41]. User-requested recommendations are usually shown on a page, e.g. a "Top-N" [78] or as "related items" [224]. Importantly, such recommendations do not get pushed to the users; rather, users "pull" these learning recommendations by visiting a portal (e.g. meta-adaptations can be made to Learning Activities that are presented a learning "App Store").

Alternatively, a recommender system can "push" adaptations to the user using mobile technologies such as text messages or smartphone notifications[2]. Pushed adaptations have the advantage of being timely (i.e. they support Just-In-Time learning [310]). Moreover, they are more suitable for recommending sequences, and can more easily adapt themselves to the task and context (see Table 12). That said, it is harder to give users a choice when using push-based recommendations, and research shows that users' privacy concerns for push-based adaptations are significantly higher than for pull-based adaptations [396].

Timing is an important aspect of push-based recommendations [203, 290, 306, 361]. Specifically, recommendations should be made only when users are available, e.g. when they are transitioning from one task to the next [71, 72]. Complex adaptive methods exist for determining when users are most interruptable [141].

Timing is also important for giving users feedback about their performance on a certain Learning Activity. Giving users clear and timely feedback about their performance maximizes their potential to learn from their mistakes, reduces evaluation anxiety, and increases users' trust in the system's subsequent recommendations.

Explanations can persuade users to follow recommendations

Adaptations are only useful if the user cares to listen to them, and in many cases this means that they need to be carefully explained. The ability to effectively explain results or reasoning could be incredibly important for TLA users when they are faced with difficult choices. Explanations contribute to increased levels of perceived system competency by making interfaces easier to use, understand, and trust [293, 354, 373].

An example of how explanations could be used within TLA is illustrated using PERLS. Figure 7 shows an 'action card' that persuades learners to set goals. Using the explanatory criteria mentioned above, the interface allows *transparency* in explaining how the recommendation was chosen using a conversational tone that users can connect with. The interface also has a very prominent design. The 'action card' is salient as it is the only item being shown on the screen and which increases the effectiveness of the recommendation. The framing of the recommendation is also very persuasive, which may be helpful for some learners but annoying for others.

[2] Note that email is an intermediate format between "push" and "pull"-based recommendations.

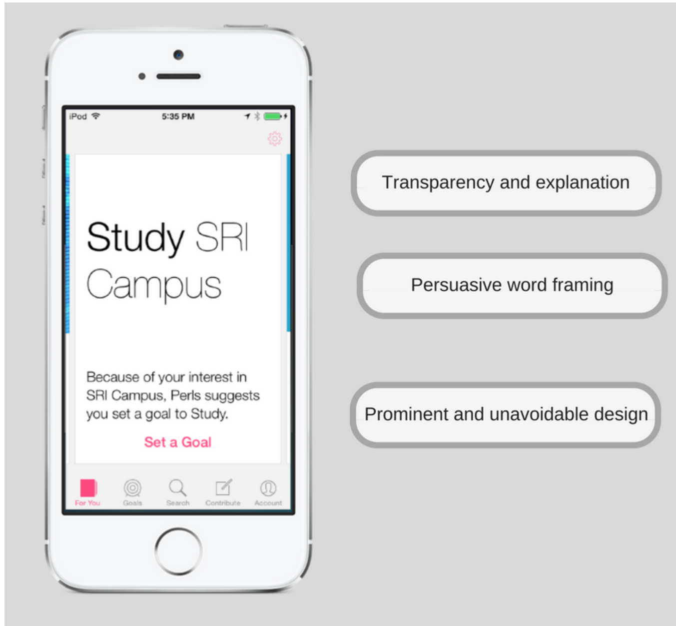

Figure 7: Action card used in PERLS. Image Source: http://perls.sri.com/

Knijnenburg [176] argues that explanations of adaptive functionalities can be provided by a human-like agent. They argue that the benefits of a human-like character can be twofold: Firstly, an agent can explicitly explain the occurrence of an adaptation, by stating what has changed and why it changed. But more importantly, an agent implicitly explains the adaptive behavior by representing the autonomous behavior of the system. When an adaptation is made, the agent can explain that *it*, instead of the system, performed the change. The agent then appears to be an autonomous body that monitors the users' interaction, reasons about their domain knowledge and choice goals, and adjusts the system accordingly. In other words, its human-like appearance can be used as an instant metaphor for autonomy and intelligent adaptiveness. The initial results of Knijnenburg's study show, however, that virtual agents had a *negative* effect on the acceptability of adaptation explanations.

An important caveat to explanations is that system developers should be careful not to "nudge" users into the direction of the recommendation too forcefully, especially when the recommendation serves purposes other than the user's own benefit. Adaptations have much in common with nudges (see Section 6.3), in that they provide a subtle yet persuasive cue or suggestion regarding the optimal user behavior [349]. Like nudges, recommendations have been shown to have *persuasive* qualities [67, 179]: users are prone to agree with a recommender's predicted ratings [65] and to follow a its advice [115]. This creates a "positive feedback loop" [275]: rather than going through the trouble of developing our own unique taste, we take the default setting and simply consume whatever the system serves us [188]. This consequence of "soft paternalism" has been criticized by both decision theorists and privacy scholars for violating the user's right to make their own decision [327, 332]. To avoid such criticism, recommendations should be framed in a way that expands rather than restricts the user's choice options. Using this philosophy, careful explanations may actually empower users to make better decisions on their own.

***Recommendation:** Provide carefully timed, well-explained, multi-purpose adaptations*

Adaptations can serve multiple purposes, not all of which are in service of the user themselves. It is therefore important not to "shove these adaptations down the user's throat", but instead respectfully involve the user in the recommendation process. Based on the analysis in this subsection, we can make the following recommendations to ADL and other TLA performers:

- **Provide multi-purpose adaptations**—The TLA Providers, which calculate and distribute adaptations, should carefully balance different adaptation purposes. Different adaptation purposes may conflict with each other. Therefore, if possible, the TLA providers should allow users to weigh adaptation purposes relative to one another.
- **Carefully time pushed recommendations**—Pushing recommendations may be more appropriate for Just-In-Time learning than requiring users to pull them. User Facing Apps should use (client-side) context-awareness to detect the optimal time to make a recommendation, to not bother users when they are busy. Likewise, apps should avoid building up evaluation anxiety, and provide users with timely feedback about their learning performance.
- **Explain recommendations without being overly persuasive**—To increase trust and confidence, apps should explain the implemented adaptations to the users. However, while apps may make recommendations with the purpose of promoting behavioral change, their explanations should avoid pressuring users into accepting adaptations that they do not want to accept, and rather give users various options to choose from and help them understand the value of each option.

3.2 Output modalities and devices

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. The meta- and macro-adaptations provided by TLA-based applications can also be pushed to (or accessed by) the user through these various devices.

This subsection discusses the pros and cons of using different devices for the display of Learning Activity recommendations (meta- and macro-adaptations). We suggest that notifications of such recommendations should be planned carefully, so as to not intrude upon users' privacy and/or leak potentially classified learning activities (see Table 13).

Table 13: Recommendations regarding output modalities and devices

Do Not Disturb the User
– Plan notifications carefully
– Do not interrupt a user's current task
– Provide easy controls for notification urgency
– Adapt notification timing to the user's context
Prevent Leaking Personal Information in Social Settings
– Provide generic notifications that do not reveal (potentially classified) details
– Change the amount of information provided in each notification depending on the number of people that are near the user

Smartphones are an ideal device for Just-In-Time learning, but can be distracting

Seventy-two percent of U.S. adults own a smartphone (slightly above the average of 68% in advanced economies) [292]. Smartphones are ideal for providing notifications, updates, schedule changes, and other information to users almost instantly, since most people carry their smartphones with them. This makes smartphones an ideal device for learning, and researchers have indeed made a push for using smartphones in what has been called "m-learning" [139].

A TLA-specification based learning system can provide users with trainings, information and personalized Learning Activity recommendations (i.e. meta- and macro-adaptations) anywhere and anytime. This ubiquitous availability allows for Just-In-Time learning, and is ideal for users who are not bound to a specific location to do their work [310]. If any new Learning Activity recommendations or scheduled trainings become available, they can appear instantly on all users' phones.

Another benefit of having a learning system on users' smartphones, is that the smartphone sensors can be used to contextualize individual Learning Activities (i.e. micro-adaptations). For example, the smartphone can find out whether the user is on the move (GPS), walking or driving (accelerometer), or in a crowded environment (microphone), and adjust the training recommendations accordingly. If users maintain other information on their smartphone as well (e.g. their calendar, social networks, and email), adaptations can use this information to inform adaptations as well. Such context-aware recommendations [10] can be very powerful, but they can also result in privacy issues [182] (see Section 2.2).

Another problem is that notifications can cause unwanted interruptions of existing tasks [203, 290, 306, 361]. In the previous subsection we therefore suggested to time notifications carefully, based on contextual cues [71, 72].

Wearables are less disruptive, but may feel more intrusive

Wearables are a more recent advancement in mobile technology that can be used for learning applications [101]. We discuss the privacy implications of tracking health data through wearable technology in Section 2.2. Here we focus on the ability of wearables to notify users of available Learning Activities. Smart watches are ideally suitable for this purpose. Glancing at a notification on a smart watch is less disruptive than having to look at one's phone, but such notifications are also harder to ignore. Moreover, a smart watch screen is small, so conveying detailed information and making privacy and interruptibility settings is challenging [250].

Apple has created several settings for its watch to reduce intrusiveness, including a Silent Mode, Do Not Disturb, and most recently Theatre Mode [29, 263] (see Figure 8). These settings allow users to receive notifications with a level of urgency that matches their current activity. Like with smartphones, adaptive methods to detect interruptibility would shift some of the burden of managing notification intrusiveness from the user to the device itself [71, 72].

Figure 8: Apple's Silent Mode and Do Not Disturb features

Notifications can leak personal information

TLA-based applications should take care not to "leak" personal information through their notifications. Overly public notifications can lead to security threats and embarrassing situations when other people (e.g. family members or visitors) get to observe these notifications [40]. Having the option to control how particular devices notify users of recommendations could avoid potential leakage of personal information.

For example, notifications could be muted when displayed on a communal display device (e.g. a smart TV) or played over a set of speakers (e.g. in the car). Rather than announcing "Would you like to start reading a newly available e-book on combat in the middle east?", which reveals a (potentially classified) learning goal of the user, the system could prompt the user by first announcing that there is an update, allowing the user to respond to get more details or alternatively to dismiss the notification.

Similar mechanisms can be used to address privacy for shared devices [127], and to address privacy for mobile devices in the case of "shoulder surfing" [132].

***Recommendation:** adapt notifications to the user's context*

The TLA is envisioned to support a multitude of devices for the consumption of Learning Activities. These devices can also be used to notify users for activity recommendations (meta- and macro-adaptations). In this subsection, we argued that these notifications should be "smart" or adaptive themselves as well. Specifically, based on the analysis in this subsection, we can make the following recommendations to ADL and other TLA performers:

- **Do not disturb the user**—Smartphones and smartwatches provide a means to push adaptations to users instantly. TLA End User Application developers should plan notifications carefully, so that they do not interrupt users' current task. Systems can either provide easy controls for notification urgency, or adapt notification timing to the user's context.
- **Prevent leaking personal information in social settings**—TLA-based apps may be used on devices that are visible to, or shared by, multiple people. In such situations, systems should provide generic notifications that do not reveal (potentially classified) details unless the user asks for them. Again, systems can use contextual cues to measure the social setting, and change the amount of information provided in each notification depending on the number of people that are near the user.

4 Data location and ownership

Problem: Who owns the data, and where does it reside? The TLA specifications enable the creation of distributed learning systems that are inherently decentralized in nature. This raises questions about where exactly the collected data resides, and which components can access and process this data. Moreover, it raises questions about who owns the training data and user models that are collected and constructed by the ecosystem of connected learning applications.

Current state of the art: No clear specification of data location and ownership. The standardized web service specifications that comprise TLA are explicitly developed for assembling component products into enterprise e-learning solutions. They provide a decentralized means to connect external learning applications, augmented with a layer of data collection and adaptation. This allows developers in creating ecosystems for self-directed life-long learners who expand their competences as they progress through their career [310]. Within these ecosystems, data is stored in the "TLA Data Core", which accumulates the data collected by various applications, and subsequently allows these applications (and the TLA Processors) to use this data for adaptation purposes (Figure 9). Questions of ownership, usage rights, and storage (beyond the central TLA Data Core) remain unanswered in the existing specifications [310].

Solution: Specifically address questions of location and ownership in the TLA architectural specification. This section addresses these questions of data location and ownership from a user-privacy perspective—in addition to discussing the effects of privacy-preserving solutions on the system's security and adaptation capabilities—through:

- Managing meta-, macro-, and micro-adaptations
- Data ownership and stewardship

We conclude by building access control for macro-adaptations and client-side methods for micro-adaptations directly into the TLA architectural specifications. Key findings and recommendations are presented in Table 14.

Table 14: Key findings regarding data location and ownership

	Key Findings	Recommendations
Managing meta-, macro-, and micro-adaptations (4.1)	– TLA Processors and Data Core should operate at the appropriate level – User Facing Apps may want to do their own macro- and micro-adaptation – Giving apps access to TLA Data Store impacts privacy – Access models for TLA may be difficult to understand – Client-side methods are ideal for micro-adaptations – Users are worried about loss of client-side data	– Implement the TLA Processors and Data Core at the appropriate level – Regulate access of individual apps to the TLA Data Core – Use client-side methods for micro-adaptations
Data ownership and stewardship (4.2)	– User data can be treated like a 401(k) – User data can be owned by multiple entities at once – A designated "data steward" can make decisions regarding user data – Using the Two-Person Concept can prevent leaks and attacks – Portable models are essential for life-long learning	– Give users ownership over their data – Give employers and apps limited co-ownership of data – Allow users to designate a "data steward" – Make user models portable

4.1 Managing meta-, macro-, and micro-adaptations

This subsection discusses the various ways in which TLA can implement its data collection and storage facilities, addresses the adaptation capabilities they rule out or enable, and analyzes their impact on users' privacy perceptions. Essentially, the TLA has three components that are relevant to the collection, storage, and processing of personal information [329] (see Figure 9):

- The **TLA Data Core** stores all the data collected from the user, including Learner Experience Facts, the Learner Profile, and Context data.
- The **TLA Processors** provide centralized meta-adaptation capabilities on top of the data core. In this setup, individual learning applications use personalization as a service [119, 374] through the aAPI.
- Individual learning applications, known as **User Facing Apps** generate data about the user that feeds into the data core through the xAPI. These apps may also provide their own macro- and microadaptations, thereby creating a distributed personalization architecture. This architecture either requires each application to access the TLA Data Core through the xAPI and the cAPI.

This subsection outlines the benefits and drawbacks of this adaptation approach, and then proposes a hybrid architecture that uses the current mechanism of TLA processors for meta-adaptations (using the aAPI), specifies a privacy-controlled distributed personalization architecture for macro-adaptations (clamping down on the xAPI), and requires context-based micro-adaptations to be implemented on the client-side, thereby avoiding the need to collect granular runtime learner activity data and context data (see Table 15).

Table 15: Recommendations regarding managing meta-, macro-, and micro-adaptations

Implement the TLA Processors and Data Core at the Appropriate Level
– Should thus operate under the auspices of a trusted entity
– Support the portability of learning models
– Allow for interoperability of TLA processors through the aAPI
Regulate Access of Individual Apps to the TLA Data Core
– Allow user facing apps to do their own macro- and micro-adaptation
– Put user access control mechanisms in place to regulate the use of the xAPI
Use Client-Side Micro-Adaptation
– Use client-side mechanisms for micro-adaptations and adaptive recommendation presentations to prevent the storage of this highly sensitive information
– Use client-side data in an ephemeral manner to prevent data loss or theft

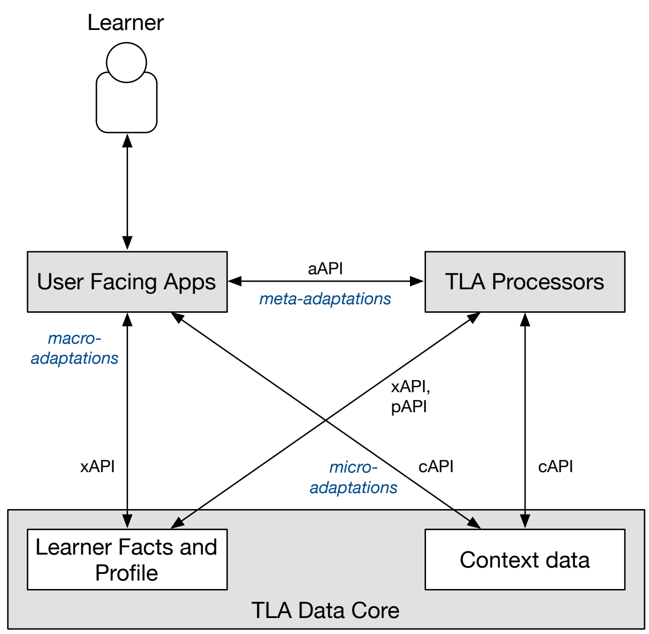

Figure 9: TLA Spring 2017 Objective Architecture (simplified for our current purposes).

> ***The TLA Processors and Data Core should operate at the appropriate level***

The TLA Processors and Data Core process and store the users' learner data and personal information, and expose the outcomes of this process (i.e. adaptations) to User Facing Apps through the aAPI [329]. This mechanism enables so-called "meta-adaptations", which are essentially recommended learning activities that span the entire spectrum of user facing apps.

Since the TLA Processor and Data Core components deal with the lion share of data collection, distribution, and use, building trust between the user and these components is extremely important. Therefore, it would be advisable to put these components under strict control of a trusted entity, such as the user's training department. However, implementing these components at the departmental level may shield them from important insights that can be gained from data collected in other departments or divisions. It also makes TLA users' mobility within the organization more cumbersome.

At the other end of the spectrum, one can imagine a situation where all TLA implementations share the same TLA Processors and Data Core. However, this would lead to performance issues [374], and it conflicts with the idea of the TLA as a specification rather than an actual infrastructure. Users may also have issues with the idea of a single entity that collects the data of all TLA users, especially after the recent hack of the Office of Personnel Management (OPM), the agency that recruits and vets prospective federal employees. In this hack, the personal information of approximately five million current and former federal employees and government contractors was stolen [112].

A good tradeoff is therefore to put these components at a level that is "low" enough for users to trust, but high enough to allow efficient mobility and user modeling synergies. Mobility problems can be further reduced through portability requirements for the Learner Profile. Moreover, to share useful learning insights, the TLA processors of different departments/divisions can be interconnected through their aAPIs.

User Facing Apps may want to do their own macro- and micro-adaptation

A centralized adaptation architecture may not be the best solution for macro- and micro-adaptations. Macro-adaptations are recommendations regarding learning activities *within* User Facing Apps. While these are similar in format to meta-adaptations, they may rely on logic that the provider of the app deems proprietary business information. For example, companies like Netflix consider their recommendation algorithms to be one of their most valuable business assets [110]. So, while it is entirely possible to let the TLA Processors handle this type of adaptation, it is politically inadvisable to require this structure.

Similarly, micro-adaptations depend on intimate knowledge of the learning activities, and so it is not only politically inadvisable, but also technically cumbersome to put the logic behind the context-based adaptation of every single learning activity into the TLA Processors and the TLA Data Core.

Giving apps access to the TLA Data Store has an impact on users' privacy

User Facing Apps that want to do their own macro- and micro-adaptation will need more direct access to the users' data than through the aAPI alone. Specifically, they may need access to the xAPI (for macro-adaptations) and the cAPI (for micro-adaptations). Allowing such access has an enormous impact on users' privacy, because it moves the TLA from a situation where all user data is stored and processed by a single entity (i.e. the entity that controls the TLA processors and Data Core) to a situation where individual apps have access to the user's data.

The TLA can deal with this situation in two different ways: On the one hand, they can provide the User Facing Apps unfettered access to the users' data. Research has shown, though, that users are likely to trust different personalization providers to a much different extent, and that this trust can be a very personal decision (e.g. while user X may trust app A more than app B, the opposite might be true for user Y) [194]. Therefore, some sort of access control mechanism is needed to allow applications to optimally utilize the users' data while at the same time respecting each user's privacy preferences [18, 218].

Access models for TLA may be difficult for users to understand

The access control model for the TLA is potentially very difficult for users to understand. The reason for this is the complex ways in which data can be collected, stored and used by different parts of and implemented TLA architecture. Notably:

- An app that *collects* data does not *store* this data; that part is done centrally, in the TLA Data Core. This means that a (potentially less-trusted) User Facing App *mediates* the data collection practices of the (hopefully well-trusted) TLA Data Core. This may cause confusion on the users' side, causing them to disclose less information.
- Any of the centrally stored data may be used by the TLA Processors to produce meta-adaptations. Even if the TLA Processors are well-trusted, this recombination of data may result in extremely accurate recommendations [288, 348] that may at times be perceived as "creepy". It would be difficult for users to anticipate such potential usage of the collected data [55].
- Moreover, any app could potentially request access to any data collected by any other app, for the purpose of macro- and micro-adaptations. This kind of cross-domain adaptation is difficult to understand, and hard to regulate [47].

Given the potential abundance of data types, User Facing Apps, and connections among and between them, a typical "who gets to see what and when" access control mechanism would arguably be too complicated for most users [214, 277]. This argument, as well as potential alternatives, will be explored in Section 6. We believe that the sheer complexity of this situation may make User-Tailored Privacy the only viable solution.

Client-side mechanisms are ideal for micro-adaptations

In Section 2, we noted how fine-grained learner runtime data and context data can contain extremely sensitive information about the user. This type of data is typically used for micro-adaptations that occur within a single learning activity. Can we support such micro-adaptations without collecting a vast amount of sensitive data?

A technical solution that has recently become popular abandons the assumption that personal data must be sent to a remote server for adaptation to take place. Rather, this "client-side" solution enables all necessary calculations take place on the user's own device [49, 159, 252]. Research in recommender systems shows that users prefer client-side methods as a means to alleviate privacy concerns [194, 342].

Although client-side adaptation mechanisms are typically limited in their ability to leverage data from other users, distributed and hybrid versions of collaborative filtering algorithms do exist [46, 319, 367]. Preventing anyone from accessing personal data enhances user privacy [331]. However, client-side adaptation methods can only use limited inference methods (e.g. if-then rules, simple classification) that can be executed directly on the user's device. They also do not contribute to the TLA Data Core, although hybrid methods exist [25].

The lack of generalized learning is less problematic for micro-adaptations, because they are typically app-specific anyway. That said, transferrable insights can still be shared with other applications without having to share the data itself. An additional benefit of client-side micro-adaptations, is that they can operate even when the user is offline, such as on a plane or in remote regions with limited cellular coverage.

Users are worried about the potential loss of client-side data

Research has shown that client-side personalization is not without problems. Specifically, users are concerned that their data can be hacked if their device is stolen, and that their user model is lost forever in case they lose or break their device [195]. Micro-adaptations may however not suffer these consequences, as they are usually ephemeral: they rely only on the *current* learner runtime and/or context data. It is thus best to implement this mechanism without storing any of such data on the user's device.

Recommendation: *Use the TLA processors for meta-adaptations, individual apps for macro-adaptations, and client-side methods for micro-adaptations*

It is impossible to provide high-quality personalized training recommendations without collecting, storing, and processing some data server-side, especially when centralized goals are expected to be taken into account in the adaptation process. A good design compromise would be a three-tier adaptation approach: On the first tier, resources, mission goals, and users' previous learning outcomes are used by the TLA processors to decide what training applications to recommend to the user (meta-adaptation). On the second tier, individual training applications can use similar data—albeit with strictly regulated access control—to make app-level adaptations (macro-adaptation). Finally, on the third tier, client-side mechanisms can use fine-grained learner runtime data and behavioral tracking to make subtle adjustments to the learning experience (micro-adaptation). Such client-side mechanisms can also be used to decide upon the ideal presentation and timing of the learning recommendations themselves (part of the Recommendation UI). These recommendations to ADL and other TLA performers are depicted in Figure 10, and further specified below:

- **Implement the TLA Processors and Data Core at the appropriate level**—These components deal with a large amount of potentially sensitive user data, and **should thus operate under the auspices of a trusted entity**. If this means that separate processors and data cores are needed for each department/division, then the TLA specification should **support the portability of learner models** and **allow for interoperability of TLA processors through the aAPI**.
- **Regulate access of individual apps to the TLA Data Core**—Since apps may consider their internal adaptation strategy a business asset, the TLA specification should **allow user facing apps to do their own macro- and micro-adaptation**. This requires access to the TLA Data Core, and the TLA specification should **put access control mechanisms in place to regulate the use of the xAPI**.
- **Use client-side micro-adaptation**—Micro-adaptations and presentation choices for the learning recommendations themselves are usually based on fine-grained learner runtime data and contextual information. The TLA can **use client-side mechanisms for micro-adaptations and adaptive recommendation presentations to prevent the storage of this highly sensitive information**. This mechanism should **use client-side data in an ephemeral manner to prevent data loss or theft**.

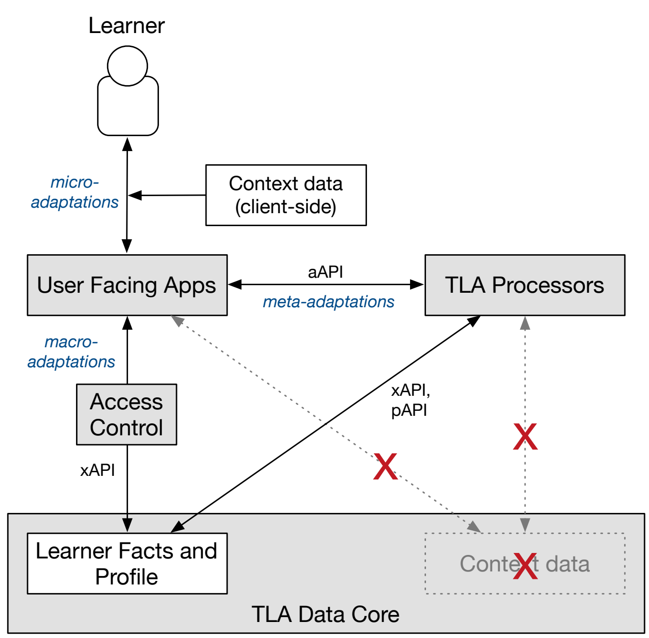

Figure 10: Proposed Architecture with Access Control for macro-adaptations and client-side micro-adaptations.

4.2 Data ownership and stewardship

As different entities contribute to the TLA Data Core, the collected data can be owned by multiple entities at once. To improve openness and mobility, we propose to treat users' data like a 401(k):

- This allows users to take active ownership over their data and decisions involving their data.
- It allows this ownership to be partially shared with other contributors.
- It allows users to delegate control to a fiduciary, or "data steward".
- It enables users to move their data from one organization (one TLA instance) to the next.

This subsection discusses the mechanisms of shared ownership, stewardship, and data mobility. Moreover, we discuss how sharing and processing decisions that involve multiple stakeholders can be implemented using the "Two-Person Concept" [365] as a means to prevent data leaks and extortion/social engineering attacks (see Table 16).

Table 16: Recommendations regarding data ownership and stewardship

Give Users Ownership Over Their Data
– Give users the right to peruse their raw data and user models
– Structure data ownership like a 401(k)
Give Employers and Apps Limited Co-Ownership
– Allow employers and apps to co-own the data
– Request minimal amounts of data, avoid duplicate storage, and de-identify data
Allow Users to Designate a "Data Steward"
– Allow users to delegate responsibilities to a "data steward" to manage the user's data under a strict fiduciary policy
– Implement the Two-Person Concept
Make User Models Portable
– Enable users to take their data with them to their new job
– Retain limited access to ex-employees' data
– Implement Private Equality Testing

A TLA user's data can be treated like a 401(k)

The end-user license agreement (EULA) of most modern online services claim full ownership over the personal information they collect about their users. The legality of this claim is questionable, though: the legal concept of "owning information" is still new, and laws are still being written about this topic [236, 401].

Preliminary debates and investigations among users show that there are merits in granting end users ownership over the personal information that is collected about them [302, 351]. Indeed, granting the user the right to peruse their raw data and user models is in line with TLA's "open" philosophy. Giving users ownership over their data also expedites the movement of data among different TLA instances—something that is very desirable given the decentralized nature of TLA and its focus on quantified self and lifelong learning [310].

Conceptually, data ownership can be structured like a 401(k): users formally own the data, but allow their employer to manage and contribute to the data. If a user moves, the data can move with them. Over time, data from different sources culminate into a well-rounded profile of the user's certifications and other capabilities.

A TLA user's data can be owned by multiple entities at once

Data ownership is not exclusive, and it may be desirable to give other entities partial co-ownership over the user's data. For example, the user's employer—who provides the user access to its TLA and the connected training applications—should also have a right over some of the data that is collected about its employee. This particularly holds true for training data itself, since it enables the company to do learning analytics, and to utilize the data in making promotion decisions. As discussed in the previous subsection, such usage may occur at a higher level in the organization, and so the user should be aware of the possibility that their data may be shared laterally within the organization for analytics and promotion purposes.

Similarly, individual training applications may use internally generated data—as well as data requested from the TLA Data Core—for macro-adaptations and internal analytics. For internally generated data, this practice mimics typical industry practices [236, 401]; for data requested from the TLA Data Core, users should be asked for permission first.

In any case, co-owners should treat user data with care. In contrast to the Big Data "collect everything mentality" [384] which permeates the current online landscape, they should request minimal amounts of data, avoid duplicate storage, and de-identify data where feasible.

A designated "data steward" can make decisions regarding TLA users' data

Data ownership puts an important responsibility on the shoulders of users. Users can decide to play an active role in making sharing decisions about their data (e.g. "who gets to see what and when" [214, 277]), but not all users may be motivated and capable of taking on this responsibility (see Section 1.2). Expanding upon the analogy of a 401(k), the TLA user should be allowed to partially delegate the responsibility of making decisions regarding their data to a fiduciary, such as their training department manager. As a "data steward" this fiduciary is allowed to make decisions about the data on the user's behalf.

Like with a 401(k), data stewards should adhere to a strict policy that outlines the intent behind their decisions and the limits of their powers. Such a policy may be a generic organizational policy, but it could also be created in a way that keeps each individual user's control preferences in mind. As such a policy can become rather complex, steps need to be taken to improve the transparency of the policy (see Section 6.1).

The fiduciary policy can outline several practices (e.g. sharing rules, processing rules) that are always allowed, never allowed, or require the explicit consent of the user. In the latter case, such consent should not just be a notice with an option to "opt out" [157, 204, 205] (Section 6.3 explains why this practice does not meet the standards of informed consent). Rather, it should ask the user to formally opt-in to the proposed practice. Another option is to combine opt-in and opt-out consent practices by algorithmically anticipating the individual user's likely response (this is in line with the idea of User-Tailored Privacy, as discussed in Section 6.4).

Using the Two-Person Concept can prevent leaks and attacks

The consent procedure of the data steward's fiduciary policy implements a Two-Person Concept solution (a concept proposed by US Air Force Instruction 91-104 [365]) that prevents any single person from intentionally or unintentionally leaking data or becoming victimized by extortion or social engineering attacks [400].

This principle also works the other way around: TLA can be implemented in such a way that the employer must give their formal approval when an employee wants to share their data with other entities pertaining to the training they did while working for that employer.

Portable user models are essential for life-long learning

Throughout their career, users may move between different employers, gaining experience, certifications and capabilities along the way. Figure 11 addresses the privacy of user models that are *portable*, i.e., that can move with the user from one employer to the next.

Different employers may use different instances of TLA. User data should therefore be specified in a standard format that allows it to be portable between TLA instances. On top of this, clear policies must be in place for when a user transfers out of their current unit [312], both in terms of what data can still be used by the former employer, as well as what data can transfer to the new employer.

In terms of the former employer, the data collected during the user's employment should still be accessible for analytics purposes even after the user leaves. At this point, though, new updates to the user's data should no longer be propagated to the former employer. Moreover, insofar as the former employee's identity is not needed for analytics purposes, the data of this employee can be de-identified, and any data that does not contribute to the analytics practices could be removed.

From a practical perspective, it may be useful to “purge” the data of ex-employees with a bit of a delay, because even if they own a portable copy of their user data, it may not always be correct, and users may initially have to come back for clarifications and corrections. Moreover, users may wish to ask their employers for a letter of recommendation, which would likely be based on their training data. For requests for recommendations that happen after the data has been de-identified, users could temporarily re-grant their former employer access to their data.

Not all data may be transferred to the new employer—the user may have certain “classified capabilities” that cannot be transferred if their new employer does not have clearance to know about these capabilities. The Two-Person Concept prevents the user from accidentally disclosing classified training data to entities without clearance. Alternatively, the concept of Private Equality Testing (PET) can be used to disclose classified capabilities without leaking them [20, 92, 150]. The user themselves may also decide to redact certain information.

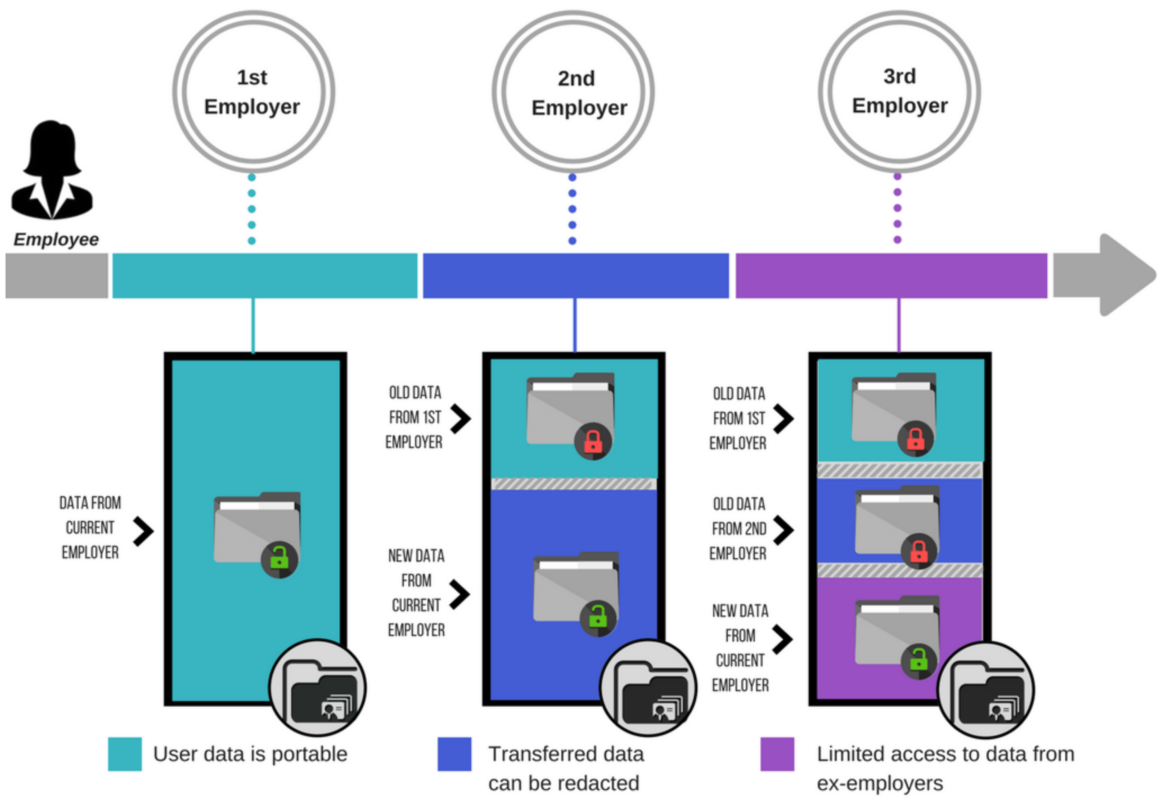

Figure 11: Managing privacy in portable user models

Recommendations: implement portability, co-ownership, and stewardship

As legal aspects of data ownership are still being debated, we recommend that TLA developers proactively decide on this question in the development of the TLA protocols and service specifications. This subsection recommends a user-centric ownership model that has facilities for portability, delegation, and shared ownership. Specifically, based on the presented analysis, we can make the following recommendations:

- **Give users ownership over their data**—In the spirit of "open" learning models that support mobility and lifelong learning, TLA should give users the right to peruse their raw data and user models, e.g. by structuring data ownership like a 401(k).
- **Give employers and apps limited co-ownership**—TLA should allow employers and apps to co-own the data for narrowly specified purposes, provided that they request minimal amounts of data, avoid duplicate storage, and de-identify data where feasible.
- **Allow users to designate a "data steward"**—TLA managers should allow users to delegate responsibilities to a "data steward", such as their training department manager. This data steward should manage the user's data under a strict fiduciary policy, that might be tailored to the user's privacy preferences. To keep users in the loop, and to prevent leaks and attacks, important decisions should implement the Two-Person Concept, where both involved parties have to authorize new data practices.
- **Make user models portable**—As users move between employers, TLA should enable users to take their data with them to their new job. The former employer may retain limited access to ex-employees' data for analytical purposes. The TLA can implement Private Equality Testing to disclose classified capabilities to authorized parties without leaking them.

5 Data sharing

Problem: How can data be shared with people and organizations? The previous section covered the exchange of user data between TLA-based applications. Data collected in the TLA Data Core can however also be used by people and organizations. What are the consequences of such social data sharing?

Current state of the art: Privacy implications of social and organizational data use are unknown. The TLA specifications call for an Open Social Learner Model (OSLM) that allows learning materials, activities, and outcomes to be shared across learners (enabling peer interactions) [370]. Moreover, the TLA design rationale puts a lot of emphasis on learning research [310, 329]. Finally, the collected TLA data can be used to make mission planning and promotion decisions. The consequences of these applications of data collected by TLA-based systems are however currently unknown.

Solution: Study the consequences of the social and organizational use of data in TLA-based systems. This section covers how recipients may use user data for various purposes while keeping the user's privacy in mind. We make the following observations:

- Sharing data with the users themselves increases trust, and can enable powerful "quantified self" experiences.
- Some user data may be shared with other users to create social learning experiences. The FERPA laws must be respected here, and care must be taken to create social learning experiences that are meaningful and not overwhelming or discouraging.
- Employers can use data about their employees to do research, and make mission planning and promotion decisions. Employers should carefully adhere to laws and regulations that protect users from unethical treatment.

Key findings and recommendations are presented in Table 17.

Table 17: Key findings regarding data sharing

	Key Findings	Recommendations
Scrutability and the quantified self (5.1)	– Scrutable Profiles increase trust and quality – Promoting the quantified self turns users into active learners	– Implement scrutability – Leverage TLA for the quantified self to motivate users
Social learning experiences (5.2)	– FERPA prevents disclosure of educational records – Users may fear "social overload" – Communication styles and social comparison styles influence social dynamics – Peer assessment depends on social dynamics	– Give users control over what to share – Allow users to limit their connections – Implement a "learning buddy" recommender – Involve users in the determination of the peer evaluation procedures
Research, promotion, and mission planning (5.3)	– IRBs require anonymized data – For placement, competency and preferences may be at odds – Algorithmic promotion decisions could obscure unwanted biases	– Let users know about secondary data use – Follow IRB guidelines for research – Act responsibly regarding placement and promotion decisions

5.1 Scrutability and the quantified self

The first and foremost entity with whom TLA can share its data is the user themselves. In this subsection, we argue that giving users insight into their user model can increase their trust, and even empower them to become more active learners (see Table 18).

Table 18: Recommendations regarding scrutability and the quantified self

Implement Scrutability
– Allow users to inspect and correct their profiles
Leverage TLA for the Quantified Self
– Use the quantified self as a motivator for data collection – Turn users into active learners using compelling infographics that establish unique behavioral connections – Use gaming elements to enable users to push themselves further – Avoid turning the quantified self features into a source of unwanted pressure

Scrutable Learner Profiles increase trust and data quality

TLA's Learner Profiles are based on advanced analysis of user behavior by the TLA Processors. The results of this analysis may not always be intuitively understandable to the user. Moreover, in some cases the insights or the facts on which they are based are incorrect, and it may be difficult for users to correct such mistakes. Building the Learner Profiles in line with the principles of "scrutability" makes them easier for users to understand and correct.

Several researchers have argued that explanations and control are important qualities of an intelligent system:

- Höök et al. were among the first to suggest a “glass box” model for adaptive hypermedia systems [140].
- Tintarev and Masthoff suggest to “explain how the system works” and to “allow users to tell the system it is wrong” [353, 354]
- Kay and Lum further unpack the idea of providing explanations, suggesting to explain why individual elements and relations in the underlying model have particular values [164].

Some researchers have shown that providing explanations and control indeed improves users’ understanding:

- Herlocker argues that “exposing the reasoning behind a recommendation” provides transparency [133] (see also [73, 74]).
- Tintarev and Masthoff show that explanations make it easier to judge the quality of recommendations [355].
- Sinha and Swearingen demonstrate that users rate systems that provide detailed information about items as more useful and easier to use [323] (see also [114]).
- Knijnenburg et al. show that mechanisms that increase transparency and control both contribute to the perceived recommendation quality and users’ satisfaction with the system [177].

Finally, research shows that explanations and control increase trust:

- Cramer et al. and Felfernig argue that explanations increase users’ trust in the recommendations [66, 96].
- Guy et al. and Wang and Benbasat show that explanations increase the perceived competence of a system [120, 373].
- Finally, Knijnenburg demonstrates that users’ understandability of and control over the personalization process influence their perceived trust and privacy threat [175].

Implementations of “scrutability” in learner models can take several levels of complexity, but it is best for the user to keep things simple. An example of a very simple scrutable user model is Google’s Ad Personalization page (Figure 12). This page shows the topics that Google has derived the user is interested in. It allows users to see how these insights were generated (“Where did these come from?”), and gives users the option to add or remove individual topics.

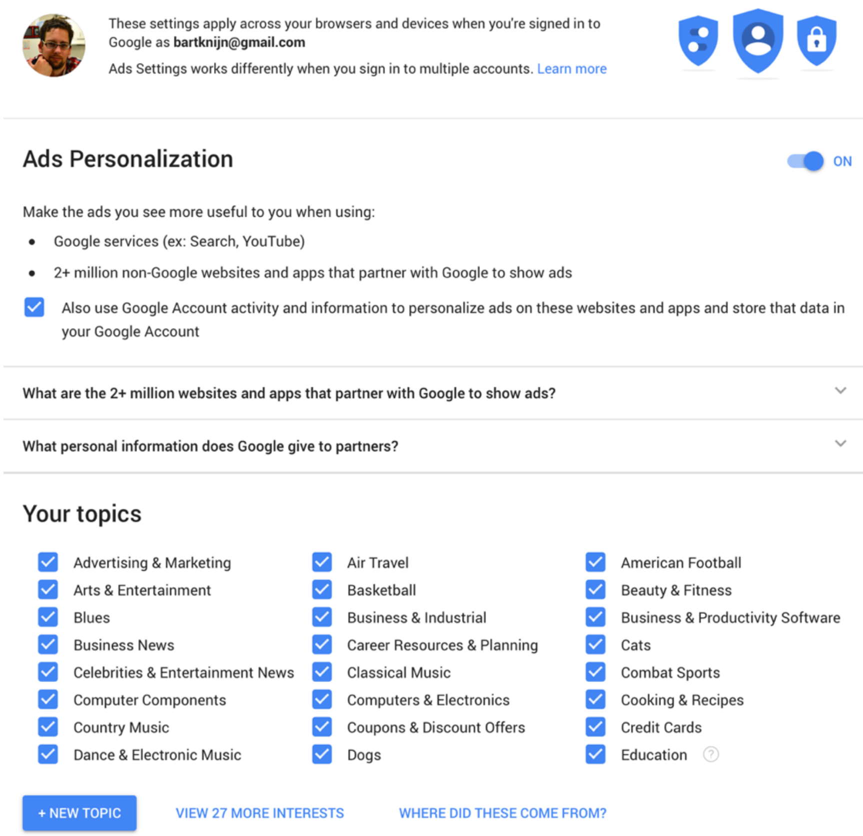

Figure 12: Google's Ad Personalization page implements a simple type of scrutability.

Promoting the quantified self turns users into active learners

Taking scrutability to a higher level, the learner modeling insights can be used to create a "quantified self" experience. The quantified self is a movement of users tracking information about themselves and using it to form insights for self-improvement. A good quantified self experience makes it easy to get data about everyday activities without having to consciously think about the process of data acquisition [215]. As mentioned in the TLA design rationale, the TLA could help individuals learn about themselves by facilitating the empirical measurement and manipulation of individual experience [310]. This way, TLA-based systems can help the user to improve their lifestyle.

Since the quantified self experience helps users to improve themselves, it is a reason for many people to accept the potential privacy intrusion that comes with wearable technology and constant tracking [26, 121]. As such, the quantified self can be a motivating factor behind TLA's data collection efforts.

Similarly, the quantified self can turn users into "active learners" through a process of self-actualization [188]. At a general level, the developers of some commercial recommender systems (e.g. OkCupid, The EchoNest[3]) have recently started to share fascinating insights into consumer tastes on their company blogs. These analyses often use compelling infographics to highlight surprising preference dynamics, sometimes broken down by state, gender, age or other demographic dimensions. Could such analyses be personalized? This would allow users to gain insights from patterns in their behavior that show a previously unknown connection (e.g. “I seem to get tired when my carbohydrate consumption is high... maybe my eating behavior causes my sleepiness”). Carefully constructed personalized infographics can allow users to explore the common and unique sides of their identity, and—if comparable across users—provide a starting point for establishing sub-cultures with similar abilities and limitations that they can explore or exploit together (see Section 5.2).

Finally, the quantified self can be a catalyst for learning. Translating self-tracked parameters into a game-like structure can create new motivational and pedagogical support structures that encourage and enable users to push themselves further [63, 107]. Games can be addictive, though [284, 334], and a system that urges the user for perfection and constantly pushes their boundaries could become a source of unwanted pressure on the individual to perform [372].

Recommendations: *Implement scrutability, and leverage TLA for the quantified self*

Sharing TLA Learner Profile info with the users themselves can keep them in the loop, help them understand the system, and increase trust. By presenting connections between data dimensions, TLA-based systems can turn users into active learners. Based on the presented analysis, we can make the following recommendations:


- **Implement scrutability**—TLA’s Learner Profiles are based on complex inferences. The TLA Data Core should therefore have facilities to allow users to inspect and correct their profiles as needed. Scrutability increases users’ understanding of the recommendation process, and is instrumental in building trust.
- **Leverage TLA for the quantified self**—Learner Profiles are an excellent source of information for the user to gain insights about themselves. As such, TLA can use the quantified self as a motivator for data collection. This paradigm can turn users into active learners using compelling infographics that establish unique behavioral connections. Moreover, TLA could use gaming elements to enable users to push themselves further. However, TLA should avoid turning the quantified self features into a source of unwanted pressure.

[3] http://blog.okcupid.com, http://blog.echonest.com

5.2 Social learning experiences

Research suggests that cooperative learning improves learning performance [156], even when it occurs in a computer-mediated fashion [166]. Communication, coordination, and collaboration can help learners support each other and improve themselves. Even competition [156] may result in cooperative benefits when implemented in a careful manner. TLA specification-enabled learning applications can provide a digital environment to support these benefits of cooperative learning.

That said, cooperative learning can also result in privacy problems: the social learning environment may disproportionately promote certain personalities and communication styles, may overload users, and may result in unfriendly interactions. This subsection addresses these privacy concerns as well as potential mitigations. We make suggestions for networking facilities that promote inclusion, foster healthy social dynamics, and limit social overload (see Table 19).

Table 19: Recommendations regarding social learning experiences

Give Users Control Over What to Share
– Refrain from sharing any learning outcomes with others by default
– Require an explicit decision from users before sharing learning outcomes with others
Allow Users to Limit Their Social Connections
– Allow users to limit their connections to those they deem relevant for each application
Implement a “Learning Buddy” Recommender
– Pair learners with similar communication styles
– Pair learners with compatible social comparison styles
Involve Users in the Determination of the Peer Evaluation Procedures
– Peer evaluations that are simple and allow for feedback
– Create a cooperative culture around peer evaluation
– Determining the peer evaluation procedures by consensus

FERPA prevents TLA from disclosing educational records

Enterprise social networking research shows that social networking helps to establish social norms [368], foster connections [268], and catalyze innovation [216, 217] among employees. These benefits seem to extend to learning as well: A study conducted at National Central University shows that learners are interested in seeing who is online and messaging them when they want to [399]. Among other things, status awareness can help with participation [93] and social navigation [94].

Note that sharing the user’s current learning status can in some cases be considered a violation of Family Educational Rights and Privacy Act (FERPA) and state laws, which prevent educational institutions from disclosing educational records to the public. Care should thus be taken that the user (not the system) makes the decision to disclose such information.

Users may fear “social overload”

How should social networks within a TLA-based learning environment be established? One possibility is to leverage users’ existing social networks. A problem with this implementation is that users may not consider all their existing social connections to be “close friends”—users of e.g. Facebook have a median of 200 contacts [1] and average seven new contacts a month [124]. If all these connections are shared with learning applications to create social learning experiences, users may rightfully fear that they could become bothered by an overload of social activity [270]. As a potential protection mechanism, we suggest that the TLA allows users to limit the sharing of social network connections to only those connections that the user deems relevant for each specific learning application.

Communication styles and social comparison styles influence social learning dynamics

Beyond allowing users to restrict their interactions to a limited number of connections, the TLA processors can also play an active role in helping users to select an appropriate learning community. This model can prevent the community from becoming unbalanced in a way that can lead to a skewed contribution model. The discussion boards of MOOCs, for example, seem to have problems with a large number of “lurkers” that may post questions but never answer them [38, 134, 201], while enterprise social networks occasionally show the opposite problem of excessive contributors that do not consume the produced content of others [269]. A careful mix of consumers and contributors prevents a social learning network from becoming ineffective.

Communication style can be another criterion for learning community selection. Moreover, the communication mechanisms provided to the network can be tailored to the predominant communication style. Referring to Section 1.3, it is important to note that messaging facilities are a typical non-FYI communication solution, while status awareness fits with FYI communicators. So conversely, the status awareness functionality may not be suitable for non-FYI communicators, and the direct messaging functionality may eventually irritate FYI communicators [270]. A possible solution, then, would be to build social learning system in a way that not only supports different learning styles [117], but also different communication styles [273].


Another thing to consider is the social dynamic involved in the creation of pairs or groups of learners. From a privacy perspective, it is better not to let everyone compare themselves against everyone else: this is overwhelming and likely ineffective. Rather, social psychology tells us that people engage in social comparison processes when they feel uncertain about their performance [138], and that some prefer to engage in upward social comparison (comparing themselves against aspirational peers), while others prefer to engage in downward social comparison (comparing themselves against trailing peers) [346]. Therefore, pairs of one upwards comparator and one downwards comparator may result in the most ideal social dynamic.

Peer assessment and social dynamics

Peer assessment is formative feedback that will help and motivate users to perform better. A socially-capable TLA specification-based implementation can provide ample opportunity for peer evaluation. To limit intrusion into the time of the evaluator, it is necessary that peer assessment methods are easy to understand and do not take too much time [340]. To be fair to the person being evaluated, it is important to allow them to reply to the evaluation [118] (see Section 2.2).

Another question is whether peer evaluation should be anonymous or not. Research suggests that the accuracy and trustworthiness of the assessment will be higher for anonymous peer reviews [161]. Without anonymity, users may feel uncomfortable giving an honest evaluation. This is even more true when evaluations include not just performance but also value perspectives [161].

In small groups peer anonymity is hard to ascertain, as social processes outside the review procedure may easily reveal the identity of the evaluators. Moreover, weak evaluations may result in mistrust, and in high-stakes team situations (e.g. a military unit), openness may be the only way to prevent criticism from ruining the social dynamic. A possible solution is thus to involve users in the determination of the peer evaluation procedures [161]; this will guarantee that users will find these procedures acceptable and fair.

The effect of peer assessment may depend on the social dynamic of the online interaction. Gamification (mentioned earlier) can enable users to push themselves further [63, 107], but when used in a social environment, it can also turn into a competitive dynamic. This dynamic can be good or bad, depending on how it is implemented.

***Recommendation:** Create a selective social learning environment with a positive group dynamic*

A TLA implementation with a social component can have a very beneficial impact on learning outcomes, but it can also result in social overload. Interaction can be beneficial, but differences in communication styles can result in friction, as can incompatible social comparison styles. At a more formal level, peer assessment may have a positive or a negative impact, depending on the social dynamic that the system creates (competitive or cooperative). Based on the presented analysis, we can make the following recommendations:

- **Give users control over what to share**—FERPA prevents TLA from disclosing educational records to the public. In case of social sharing, TLA should thus refrain from sharing any learning outcomes with others by default. Rather, TLA should require an explicit decision from users before sharing learning outcomes with others.

- **Allow users to limit their social connections**—Users may not want to share their learning experience with all their social network contacts. Instead, a socially-capable TLA experience should allow users to limit their connections to those they deem relevant for each specific learning application. This prevents social overload.
- **Implement a "learning buddy" recommender**—To support users in selecting the best social connections to include in their learning experience, a socially-capable TLA should implement a recommender that can pair learners with similar communication styles to prevent unbalanced contributions. That recommendation can also pair learners with compatible social comparison styles, specifically, a person who prefers upward social comparison should be paired with a person who prefers downward social comparison.

- **Involve users in the determination of the peer evaluation procedures**—An important aspect of a socially-capable TLA environment is that it allows for mutual assessment. Users should have access to peer evaluations that are simple and allow for feedback. To prevent criticism from ruining the social dynamic, it is better to create a cooperative culture around peer evaluation. Training department managers can accomplish this by determining the peer evaluation procedures by consensus.

5.3 Research, promotion, and mission planning

The primary purpose of data collected by TLA-based systems is to allow the TLA Processors to provide personalized learning recommendations. The data can however also be used for research, for mission planning, and to decide on promotions. Privacy experts argue that secondary use of 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 [347].

The use of data for research, promotion, and mission planning is particularly sensitive because it involves employers collecting and using data about their workers. Researches have shown that "on the job monitoring/tracking" can have a deleterious effect on workers' performance. Specifically, whereas Nebeker and Tatum show that automated computer monitoring can lead to increased speed of work, they found no improvement in work quality, satisfaction, and stress [255]. Chalykoff and Kochan in fact showed that there are significant negative effects of monitoring [53]. They were able to resolve these negative effects for some (but not all) employees by giving them extensive feedback and performance appraisal.

Aside from that, there are laws and regulations surrounding research and employment-related practices that need to be adhered to, and even beyond these formal restrictions it would serve TLA developers to think about how TLA can safeguard the ethical treatment of research subjects and employees. This subsection discusses ethical guidelines that serve this purpose.

This subsection covers the use of data for research, promotion, and mission planning (see Table Table 20). We argue for the establishment of responsible practices regarding how data will be used for these purposes, as well as the clear communication of these practices to users.

Table 20: Recommendations regarding research, promotion, and mission planning

Let Users Know About Secondary Data Use
– Communicate secondary data use practices to users
– Indicate exactly which data was used and for what purpose
Follow IRB Guidelines for Research
– Anonymize research data
– Allow for the communication of incidental findings
Act Responsibly Regarding Placement and Promotion Decisions
– Establish guidelines surrounding conflicts between competence and preferences
– Make sure that promotion decisions are made in a non-discriminatory manner

IRBs require research to be conducted on anonymized data

Data captured and generated by the TLA for research purposes, allowing product developers and publishers to fine-tune their training experiences, as well as allowing education and training organizations to develop more accurate and fine-grained competency frameworks [310]. IRBs usually prefer such research to be conducted using de-identified data. However, in the era of powerful data analytics that can uncover very fine-grained insights, this raises its own ethical concerns. For example, in an analysis of big data ethics in the military, Schneider et al. ask themselves what if an analyst discovers a pattern of mental illness or suicidal thoughts in the data of a single user: "Does the military's prerogative to prevent suicides—arguably at any cost—override […] concerns about privacy and fairness?" [312].

Whether to report "incidental findings" to the user has been debated by ethicists, but there is no clear guidance on whether and how such findings should be disclosed [167]. As for how: one could keep a table linking participants to anonymized codes in a separate, protected location. This allows researchers to reach out to participants if an incidental finding indeed does occur.

For placement, competency and preferences may be at odds

TLA data can also be used to make deployment decisions, allowing supervisors to recruit teams with uniquely matched competencies, or, alternatively, train up existing teams to attain the competencies needed for a certain mission. Again, an ethical discussion would need to take place regarding the potentially conflicting roles of competency and user preferences: if a user is the only one in their division to have a certain language competency, but does not want to be deployed to the country where that language is spoken, should the user's competency or their preference take precedence in a planning officer's decision of whether to deploy the user? What if this decision not only affects the user, but also the rest of their unit?

As such decisions are likely going to depend on a complicated mix of factors, it is important that the procedures are supported by all employees involved. This means establishing and communicating the procedures beforehand, and possibly giving users a say in the development of these procedures.

Algorithmic promotion decisions could obscure unwanted biases

Finally, TLA data can be used to make promotion decisions. Schneider et al. highlight the important ethical considerations of using machine judgment for promotion decisions [312]. On the one hand, one may argue that data-driven promotion decisions are void of emotional and political biases, thereby increasing fairness. On the other hand, algorithmic decisions have been shown to incorporate biases themselves—and even worse, to obscure them [229]. One may also argue that it is important not to see a user as merely the sum of his/her competencies, as some qualities may be hard to quantify.

In the context of employment discrimination laws (Title VII, ADA, ADEA), it is advisable to check algorithmic promotion and placement decisions for inadvertent discriminatory biases. This can be done by hand, e.g. using the Delphi method [303], but automated solutions may be available: recent work has explored the use of "propensity scoring" to reduce availability biases in datasets [311]. Using the same method as a *post hoc* filtering technique can reduce unwanted biases in recommendation results as well.

***Recommendation:** Start discussing the ethics of learning data analytics*

Beyond existing laws and regulations, TLA developers should start discussing the ethics of learning data analytics in the context of research, promotion, and mission planning [312]. Based on the presented analysis, we can make the following recommendations:

- **Let users know about secondary data use**—Users may not expect that their data will be used for research, placement, and promotion decisions. Training department managers should communicate secondary data use practices to users. Since the TLA Data Core may contain a very wide variety of information, it is best to indicate exactly which data was used and for what purpose, and give users extensive performance appraisal.
- **Follow IRB guidelines for research**—Human subjects research is subject to review by an Institutional Review Board. IRBs usually require researchers to anonymize data as much as possible, but they allow researchers to keep a key list with research subjects' identities in an offline secure location to allow for the communication of incidental findings.
- **Act responsibly regarding placement and promotion decisions**—Training department managers should acknowledge the fact that TLA users are more than the sum of their training data. They must establish clear guidelines surrounding potential conflicts between competences and preferences when it comes to placement decisions. Moreover, they should make sure that algorithmic promotion decisions are made in a non-discriminatory manner.

6 Privacy support mechanisms

Problem: How can we support TLA users making privacy settings? Throughout this document we have demonstrated that TLA-based systems must deal with a lot of privacy-related issues. Even when they are designed and implemented with privacy in mind, it will be inevitable for these systems to have a wide array of privacy settings that allow users to customize their experience to fit their personal preferences when it comes to the unavoidable tradeoff between privacy and utility. How can we help users in making these privacy settings?

Current state of the art: There are problems with the current paradigms of "notice and control" and "privacy nudging". To help users with this tradeoff, many privacy experts recommend the practice of *notice and control*: giving users comprehensive control over their privacy, while at the same time providing them with more information about the implications of their decisions l [43, 52, 213, 305, 345, 397]. Notice and control are also at the heart of existing or planned regulatory schemes [90, 381]. However, research in the past few years has unveiled a fair number of "privacy paradoxes": situations or conditions in which transparency and control do *not* increase people's privacy, or even *decrease* it (see Section 1.1).

An alternative solution that has recently gained more traction is *privacy nudging*, an approach to privacy support that attempts to relieve some of the burden of privacy decision-making, by making it easier for people to make the right choice, without limiting their ability to choose freely [4, 8, 23, 375, 376]. Privacy nudging has also had only limited success, arguably because privacy nudges take a "one-size-fits-all" approach to privacy [336]: They assume that the "right" privacy decision is the same for every user, piece of information, and situation.

Solution: Introduce personalized privacy decision support with "user-tailored privacy". To overcome these shortcomings of transparency-and-control and privacy nudges, privacy scholars need to move beyond the "one-size-fits-all" approach that is embodied in both nudges and transparency and control. Because of the high variability and context-dependency of people's privacy decisions, nudges need to be *tailored* to the user and her context. The idea of *user-tailored privacy* is the latest development in the quest for more usable privacy support.

This section discusses existing techniques for privacy notices, control, and nudging. It also discusses their shortcomings. It then sets the stage for user-tailored privacy, which will be central to the next version (0.2) of this specification document. Key findings and recommendations are presented in Table 21.

Table 21: Key findings regarding data sharing

	Key Findings	Recommendations
Privacy notices (6.1)	– Privacy nutrition labels give quick overview information – Textured agreements connect overview to detail – Privacy comics can increase efficacy and motivation – Privacy notices may not always work	– Use textured, comic-based privacy nutrition labels – Choose simplicity over notices
Control mechanisms (6.2)	– Accessible privacy controls increase self-efficacy – Graphical designs can simplify access control matrices – Privacy control goes beyond disclosure – Users are not always motivated to take control	– Use accessible, graphical privacy controls – Choose simplicity over control
Privacy nudging (6.3)	– Justifications provide a shortcut to decision-making – Audience feedback makes users more aware of who sees their data – Defaults make it convenient to take the right action – Nudges may threaten autonomy	– Use nudges if there is a virtual consensus
User-tailored privacy (6.4)	– Privacy behaviors vary, but are predictable – Users' behaviors can be used to make adaptations	– Employ user-tailored privacy when possible

6.1 Privacy notices

Several designs solutions have been proposed to increase users' understanding of privacy-related information. This subsection covers these solutions, such as nutrition labels, textured notices, and comics, and critically appraises their effectiveness (see Table 22).

Table 22: Recommendations regarding privacy notices

Use Textured, Comic-Based Privacy Nutrition Labels
– Use privacy nutrition labels to give people a quick overview
– Make privacy notices textured to connect to the details
– Use comics to make privacy notices attractive and approachable
Choose Simplicity Over Notices
– Use notices sparingly
– Make privacy decisions simpler rather than relying on notices

"Privacy nutrition labels" give quick overview information

One realization about privacy notices is that the complexity of online privacy policies is ever-increasing [245]: they are often written in a legalistic and confusing manner, and require a college reading level to understand them [16, 50, 171, 238, 363]. Indeed, while many people claim to read online privacy policies [147, 244], many do not actually read them [9, 30, 31, 129, 154, 171, 322, 362], or do not read closely enough to understand them [274]. A lot of work has therefore gone into standardizing and summarizing privacy statements [111].

Textured agreements connect overview to detail

A problem with summarizing privacy notices is that they are often too simplistic to accurately represent the policies they reflect [258]. If users ignore the full policy, they may thus not have all the information they need to make a decision [165].

Textured agreements keep the original policy intact, but add layers of emphasis (e.g. headings, bullets, bold text, highlights, graphics) to make the text more readable [165, 350]. Textured agreements *increase* (rather than decrease) the amount of time people spend reading the agreements, primarily because more participants end up looking through the entire agreement.

Privacy comics can increase efficacy and motivation

Another design idea used to increase users' efficacy and motivation to read privacy policies is to use comics. Comic books are already used to e.g. increase literacy [42, 102] and provide health education [113, 235]. The visual aspects of comics are captivating, and can often be used to understand the story without having to read any text [86, 165]. This makes them particularly appropriate for increasing the motivation and self-efficacy of people with lower literacy levels and/or a visual learning style [85, 297]. Work on privacy comics is still in its infancy [178], but we argue that they are likely to be uniquely capable of instilling motivation and ability in users, who would normally forgo learning about privacy.

Privacy notices may not always work

Although the consensus is that users should be informed about the privacy decisions they are asked to make [87, 145, 243, 301, 396], the reality is that doing so often makes users more fearful or unwilling to come to a decision. For example:

- Marketers 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 [2, 33, 142].
- Privacy policies have been shown to incite privacy concerns rather than easing them [291].
- John et al. [155] demonstrate that even subtle privacy-minded designs and information may trigger users' privacy fears and thereby reduce disclosure and participation rather than increasing it.
- Adjerid et al. [8] show that the impact of privacy notices depends on their specific framing, and that distractions can easily nullify any effect of privacy notices.

The conclusion, then, is that transparency does not work well in practice, especially for systems that process large amounts of personal data, which is increasingly the case online [352]. Nissenbaum [258] postulates this as the Transparency Paradox: 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.

> ***Recommendation:*** *Where appropriate, use textured, comic-based privacy nutrition labels*

TLA-based implementations are likely to have a plethora of privacy settings. Users should not be expected to understand these settings without help. Based on the presented analysis, we can make the following recommendations to ADL and other TLA performers:


- **Use textured, comic-based privacy nutrition labels**—Users are likely not going to read full-length privacy statements. User-facing apps should therefore use privacy nutrition labels to give people a quick overview of the privacy implications of using the system. An overview is likely not enough to help users make fully informed decisions, hence applications should make privacy notices textured to connect to the details of the available privacy settings. Finally, applications should use comics to make privacy notices attractive and approachable to people at lower reading levels.
- **Choose simplicity over notices**—Privacy notices sometimes have an effect that is opposite from the intended effect. Hence, TLA-based applications should use notices sparingly, and work to make privacy decisions simpler rather than relying on notices to inform users about the decisions they are expected to make.

6.2 Control mechanisms

Like with privacy notices, several designs solutions have been proposed to give users more intuitive control over their privacy settings. This subsection covers these solutions—e.g., controls that are easily accessible, graphical, and go beyond information access—and critically appraises their effectiveness (see Table 23).

Table 23: Recommendations regarding control mechanisms

Use Accessible, Graphical Privacy Controls
– Make controls obvious and easily accessible
– Use graphical methods to provide control
– Provide controls that go beyond information access
Choose Simplicity Over Control
– Use a privacy setting that works for everyone (where possible)
– Not make control too detailed

Accessible privacy controls increase self-efficacy

Knijnenburg et al. investigated form auto-completion tools, which automatically fill out Web forms based on previously collected personal information [183]. They demonstrated that users of these tools typically fail to consider the perceived risk and relevance of each piece of potentially private information. Consequently, they developed two alternative design solutions to promote more explicit privacy decision making. The first design solution—the "remove" type auto-completion tool—improves upon traditional auto-completion by allowing users to remove an auto-completed entry by means of a button adjacent to each field. The second design solution—the "add" type auto-completion tool—leaves all fields blank by default and provides a button to add the pre-collected information to each field. Knijnenburg et al. argued that the process of disclosing personal information is more elaborate for users of the add/remove auto-completion tools than for users of the traditional auto-complete tool. In their experiment, they indeed demonstrate that due to the availability of buttons, users feel more able to take control over their disclosure, and hence become more deliberate about their decisions when using add/remove tools.

These results suggest that simple privacy controls can make users feel more in control—and, indeed, take control—over their privacy-settings. Even though the suggested buttons made removing/adding information only slightly easier, they significantly increased users' focus on the purpose of the requested information in deciding what to disclose.

Graphical designs can simplify access control matrices

In many cases, privacy control can be formulated as an "access control matrix", for example when deciding what to share with whom. Such a control matrix may for example be required to specify which User Facing Apps have access to what data in the TLA Data Core (see Section 4.1).

Several solutions have been proposed to simplify such control matrices. One solution is to group recipients, cf. "Circles" in Google+ [160, 181, 377]. Another solution involves creating a graphical representation of the control matrix that is automatically sorted to show interesting patterns [257]. Finally, Raber et al. proposed "wedges" to combine two dimensions (recipient type and social distance) in a single intuitive interface. They found that users could make more accurate privacy decisions using the wedges-based interface, and liked the interface better [295].

Privacy control goes beyond disclosure

Selective information sharing is just one of many strategies SNS users may employ to alleviate privacy tensions [206, 233, 360]. This realization is in line with Altman's broader definition of privacy as “an interpersonal boundary process by which a person or group regulates interaction with others,” by altering the degree of openness of the self to others [12].

Likewise, privacy control can be provided in more diverse and more intuitive ways than a traditional “sharing matrix” in which users specify who gets to see what [136, 207, 251]. For example, Facebook’s privacy features support a variety privacy management behaviors that go beyond selective information sharing; 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) [162, 387]. Research has found that it is important to give users the privacy features they want, lest they experience reduced social connectedness, and miss out on social capital [388].

Users may not always be motivated to take control over their privacy

The Control Paradox states that while users claim to *want* full control over their data [5, 28, 39, 198, 343, 344, 356, 378, 393], they avoid the hassle of actually *exploiting* this control [61]. In combination with overly permissive defaults [37, 116], this leads to a predominance of over-sharing. For example:

- Larose and Rifon [210] find that privacy seals influence disclosure tendencies only for participants that are either motivated or have a high self-efficacy.
- Besmer et al. [32] find that participants were only influenced by social navigation cues if they already had a tendency to change their settings.
- Gross and Acquisti [116] show that only a small number of Facebook users change the default privacy preferences.

Like transparency, control does not work well in practice. Systems like Facebook that manage large amounts of personal user data have to resort to “labyrinthian” privacy controls [91]. As a result most Facebook users do not seem to know the implications of their own privacy settings [225, 339], and share postings in a manner that is often inconsistent with their own disclosure intentions [231].

Recommendation: *Where appropriate, use accessible, graphical privacy controls*

TLA-based implementations should have privacy settings interfaces that are easy to use. Based on the presented analysis, we can make the following recommendations to ADL and other TLA performers:

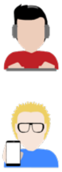

- **Use accessible, graphical privacy controls**—Privacy settings should not be buried deep inside a system's settings. Instead, TLA-based applications should make controls obvious and easily accessible. Doing so increases users' self-efficacy. Moreover, some privacy control features can be very complex. Apps should use graphical methods to provide control in these cases; this makes control more intuitive. Finally, TLA-based applications should make sure to provide controls that go beyond information access. This allows users to address relational, territorial, network, and interactional boundaries.
- **Choose simplicity over control**—Users say they want control over their privacy, but they rarely use it. Therefore, where possible, User Facing Apps should use a privacy setting that works for everyone (where possible), so that control is not needed. In any case, apps should not make control too detailed, that way users will not feel overwhelmed.

6.3 Privacy nudging

Privacy nudging, a recent approach to support privacy decisions, tries to overcome some of the problems with privacy notices and control. Nudges are subtle yet persuasive cues that makes people more likely to decide in one direction or the other [349]. Carefully designed nudges make it easier for people to make the right choice, without limiting their ability to choose freely. This subsection describes privacy nudges that have been tested (e.g., justifications, audience feedback, and defaults), and discusses their shortcomings (see Table 24).

Table 24: Recommendations regarding privacy nudging

Use Nudges if There is a Virtual Consensus
– Use justifications, audience feedback, and defaults when virtually all users agree on the optimal privacy setting
– Use nudges to provide users choice in the unlikely event that they want a different setting after all

Justifications provide a shortcut to privacy decision-making

The type of nudge that is most extensively implemented in real systems is *justifications*. A justification is a succinct reason to disclose or not disclose a certain piece of information. It differs from a privacy notice in its brevity and its purpose: rather than educating users about privacy, justifications make it easier to rationalize the decision [34, 321] and to minimize the regret associated with choosing the wrong option [62, 146]. Justifications include providing a reason for requesting the information [64], highlighting the benefits of disclosure [197, 373], and appealing to the social norm [7, 32, 278].

The effect of justifications seems to vary. Specifically:

- In a study by Kobsa and Teltzrow [197], users were about 8.3% more likely to disclose information when they knew the benefits of disclosing the information.

- In an experiment by Acquisti et al. [7], users were about 27% more likely to do disclose information when they learned that many others decided to disclose the same information.
- Besmer et al. [32] found 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.
- Patil et al. [278] rate social navigation cues as a secondary effect.
- Knijnenburg and Kobsa [182] test various different notifications, and find that while they all seem to be perceived as useful (except for the social justification), none of them seem to increase users' trust in or satisfaction with the system.

Another justification strategy is to provide a symbolic rather than textual privacy indicator, e.g. a "privacy seal". Again, such indicators have varied success:

- Egelman et al. [87] show that privacy indicators next to search results can entice users to pay a premium to vendors with higher privacy scores.
- Users of Xu et al.'s [396] location-based coupon service were more likely to disclose information when the site displayed either a TRUSTe seal or a legal statement.
- In Hui et al.'s [145] marketing survey, a privacy seal did not significantly increase disclosure.
- Studying an online CD retailer, Metzger [243] found that their seal had no effect.
- Rifon and Larose [301] show that warnings and seals at an online retailer website influence users in certain situations only.

Audience feedback makes users more aware of who sees their data

In social media, nudges often relate to the real or potential audience of a shared piece of information. Again, the effects of these nudges are mixed. For example:

- In location sharing services, researchers have experimented with giving users real-time feedback on who is requesting or viewing their location [153, 359]. Users appreciate the information, but find that it can easily become excessive and annoying.
- Wang et al. [375, 376] provide users with detailed feedback about the potential audience when posting a Facebook message. Some users consider this tool helpful, but they find no significant differences in posting behavior.

Somewhat related Wang et al. [375, 376] consider two other types of nudges as well: sentiment feedback (telling users whether the message they are about to post is likely to be perceived as positive or negative) and a post timer (delaying Facebook posts by 10 seconds, which allows users to change their mind). Some of the participants in their study seemed to like these tools, but others found them intrusive and annoying.

Defaults make it more convenient for users to take the right action


Another approach to nudging users' privacy decisions is to provide sensible defaults. Correctly chosen defaults make it easier to choose the right action, or may not even require any action at all. In this sense, defaults reduce physical [307, 349] or mental effort [364]. Defaults also provide an implicit normative cue, e.g., a default order communicates what the system thinks is most important, and a default value communicates what the system thinks you should do [240]. Finally, default values may work due to the 'endowment effect': 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 [6, 358].

Providing a certain default option nudges users in the direction of that default [349]. Therefore, while most existing work on default effects in the privacy field regards them as a nuisance, several researchers have recently suggested that they can also be used as nudges [4, 8, 23]. Empirical work on defaults as nudges is sparse:

- Knijnenburg et al. [181] showed that the odds of disclosure when social network information was shared with everyone by default were 3.9 times as high as in the private-by-default-condition, although this effect is smaller for participants with low interpersonal privacy concerns and when categories are ordered weaker ties first.
- Both Johnson et al. [157] and Lai and Hui independently showed that the sign-up rates for newsletters was about 25 percentage-points higher when sign-up was the default setting.

The order in which information requests are made can also be perceived as a default:

- Knijnenburg et al. [181] showed that the odds of disclosure when users were asked about weaker ties first were 1.8 times as high as when users were asked about stronger ties first. This "door in the face" effect confirms earlier findings by Acquisti et al. [7].
- Knijnenburg [175] found that the opposite order is more effective in sustaining disclosure when answering more questions is optional.

Nudges may threaten user autonomy

The privacy nudges evaluated in existing work show mixed results: they usually only worked for some users, and left others unaffected or even dissatisfied. Because of this, researchers argue for "smart nudges", such as smart default settings that match the average user's preferences [327, 332].

But what if the "average user's privacy preferences" do not exist? In Section 1 we have cited ample evidence that people vary extensively in their information disclosure behavior, and that even for the same person this decision depends on the context in which it is made. The current implementations of nudges, however, take a "one-size-fits-all" approach to privacy [336]: They assume that the "true cost" [155] of disclosure is roughly the same for every user, for every piece of information, in every situation. Since such nudges are rarely good for everyone, some researchers therefore argue that they may threaten consumer autonomy [327, 332].

***Recommendation:** only use nudges if there is a virtual consensus*

Nudges are an interesting way to help users make the right choice without limiting their decision freedom. In most privacy settings, the "right choice" is difficult to define, though, hence nudges will not be welcomed by every user. Based on the presented analysis, we can make the following recommendations to ADL and other TLA performers:

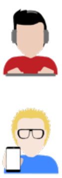

- **Use nudges if there is a virtual consensus**—TLA-based apps should use justifications, audience feedback, and defaults on when virtually all users agree on the optimal privacy setting. In those cases, apps can use nudges to provide users choice in the unlikely event that they want a different setting after all. More intelligent forms of nudges are discussed in the next subsection.

6.4 User-tailored privacy

How can we reconcile the need for extensive customizability with users' apparent lack of skills and motivation to manage their own privacy settings? This subsection discusses User-Tailored Privacy (UTP) as means to support users' privacy decision-making (see Table 25). With UTP, a system would first measure users' privacy-related characteristics and behaviors, use this as input to model their privacy preferences, and then adapt the system's privacy settings to these preferences. Figure 13 shows a schematic overview of UTP.

Table 25: Recommendations regarding user-tailored privacy

Employ User-Tailored Privacy to Support Users' Privacy Decision-Making
– Determine the TLA-specific modeling factors and clusters that can be used as input for privacy modeling
– Specify adaptation strategies that will be used to implement the privacy adaptations

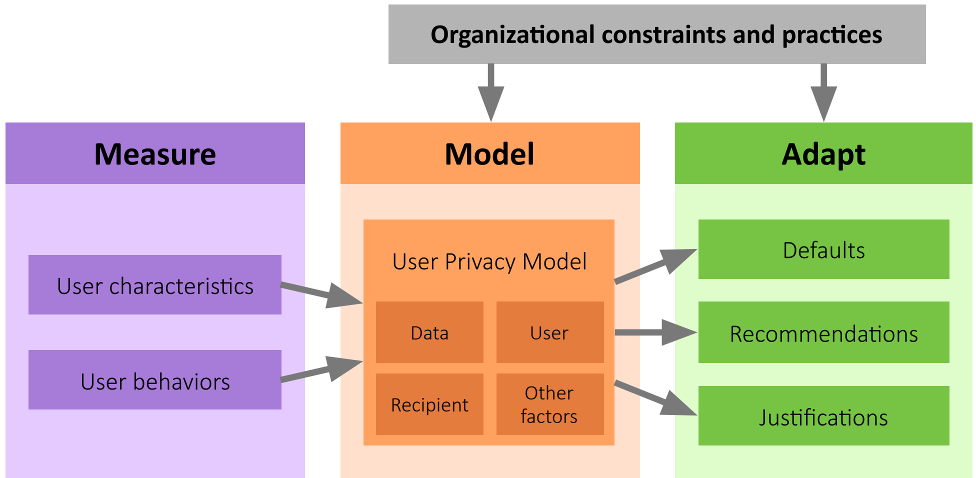

Figure 13: A schematic overview of User-Tailored Privacy (UTP)

Privacy behaviors vary, but are predictable

The User Privacy Model underlying UTP exploits the fact that users' privacy behaviors are predictable. For example, several researchers have found that people's privacy concerns are multi-dimensional, meaning that they have different preferences for different types of information [184, 199, 228, 266, 336]. In fact, research also shows that there exist distinct profiles of privacy behaviors among users [184, 266, 391]. For example, users' public disclosure of 15 types of Facebook profile items demonstrated a 4-dimensional structure: Facebook activity (e.g. wall posts, status), Location (e.g. city, state/province), Contact info (e.g. phone number, email address), and Interests (e.g. likes, groups). Users in this dataset were clustered into 5 distinct profiles [184].

Moreover, the recipient of the information seems to play an important role in users' disclosure decisions, both in commercial and social privacy settings [160, 181, 183, 214, 277, 377]. Again, certain "groupings" can be made. For example, for social network recipients, Knijnenburg et al. [186] found that five categories (Family members, Friends, Classmates, Colleagues, Acquaintances) resulted in the optimal solution in terms of privacy threat and usability.

Finally, 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 [28, 81, 392].

Users' behavioral patterns can be used to make privacy-related adaptations

UTP can subsequently use these patterns to provide privacy-related adaptations. These adaptations could take the form of a default setting or a recommendation, either with or without an accompanying justification:

- Knijnenburg and Kobsa [180] demonstrated the potential of "adaptive justifications", changing the presented justification based on the users' overall disclosure tendency and their gender. This method significantly helped users with their information disclosure decisions.
- Knijnenburg and Jin [179] used a user-tailored approach to simplify the sharing options in a location-sharing system. The study considered a hypothetical system that allowed users to "check in" to a location using one of eight sharing options. We found that reducing the number of options adaptively resulted in somewhat higher perceived decision help.
- Knijnenburg [175] studied adaptive request orders in a demographics-based health recommender system (Figure 14). The system asks demographics questions in a sequential order, and recommendations are adapted to the user's answers on the fly. The user can skip questions if they deem them too sensitive. A study tested several means of ordering the demographics questions. Request orders that automatically trade off the usefulness and sensitivity of the items to be disclosed improved the users' experience.

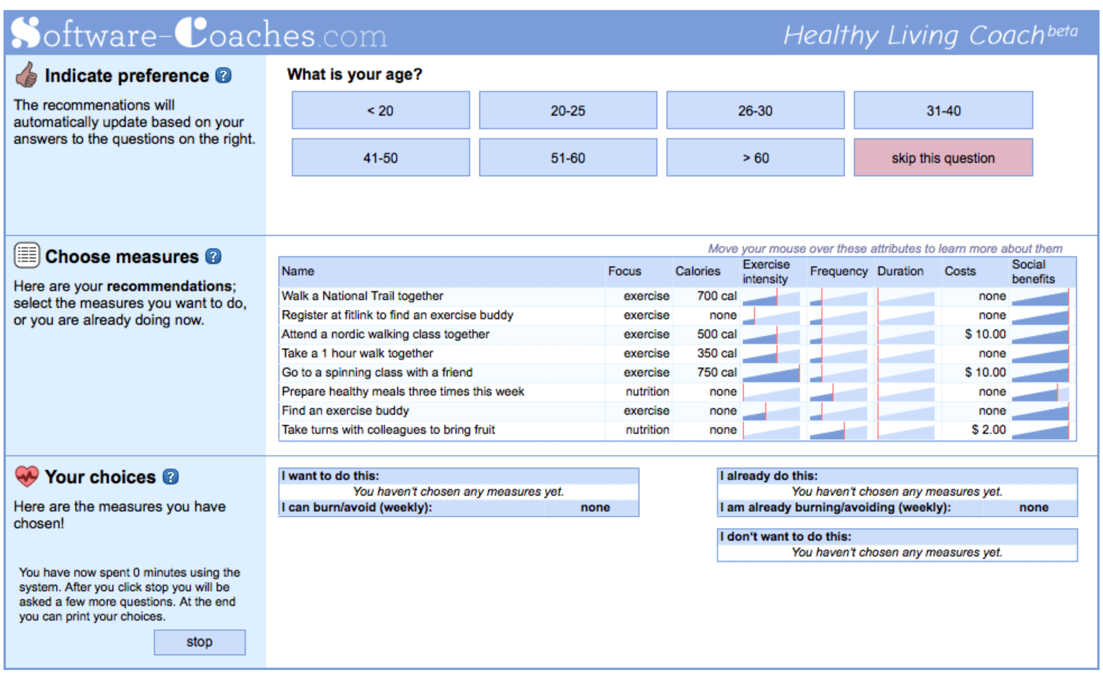

Figure 14: A demographics-based health recommender system that uses adaptive request orders to decide which demographics question to ask next

Recommendation: *Employ user-tailored privacy when possible*

The Idea of UTP fits very well within the extensive user modeling approach of the TLA. Moreover, given the complexity of the 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. A number of TLA-related examples may help illustrate UTP:

- The TLA normally tracks users' location (Data) in order to give context-relevant training exercises (Organizational practice). However, UTP knows that like many young mothers (User characteristic), Mary (User) does not want her location (Data) tracked outside work hours (Other factor). It therefore turns the location tracker off by default when Mary is not on the clock (Default).
- David needs to decide how to share his recent milestones—two certificates he has recently earned (Data)—within his organization (Recipient). Due to the rules of his employer (Organizational constraint), UTP requires him to share these milestones with his direct supervisor (Recipient). Moreover, from his previous interactions (User behaviors), the UTP knows that David keeps close ties to several other divisions. UTP therefore suggests (Recommendation) that he should share his new certifications with the heads of these divisions (Recipient) as well, arguing they are likely to be interested in exploiting his newly gained skills (Justification).

UTP aims to strike the balance between giving users no control over, or information about, their privacy at all (which will be insufficient in highly sensitive situations, and may deter privacy-minded individuals) and giving them full control and information (which makes setting one's privacy settings unmanageably complex). 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 [175]. Based on the presented analysis, we can make the following recommendations to ADL and other TLA performers:

- **Employ User-Tailored Privacy to support users' privacy decision-making**—Employing UTP within TLA consists of two steps. First, one should determine the TLA-specific modeling factors and clusters that can be used as input for privacy modeling. This addresses the input-side of UTP. Then, one should specify adaptation strategies that will be used to implement the privacy adaptations. This determines whether the adaptations will take the form of a user-tailored justification, a smart default, or an adaptive request order.

These steps will be undertaken in version 0.2 of this document.

Conclusion

In this document, we have made recommendations regarding the Operational Characteristics of TLA-based systems that impact users' privacy concerns. These recommendations will allow ADL and other TLA performers to select the characteristics that best alleviate users' concerns. The recommendations in the current version of the document are tentative; the specifics of selected solutions will be added after intensive discussion with ADL and other TLA performers during the development of version 1.0 of this document.

In the meanwhile, we suggest that TLA performers pay attention to the **User Characteristics** of TLA users, e.g., by tailoring to different privacy management strategies and communication styles in their system designs.

Moreover, rather than collecting users' personal information indiscriminately, TLA performers should consider **Input Data Characteristics**. They should make clear distinctions between the collection and use of various data types, and allow users to scrutinize and correct potential mistakes in system predictions.

TLA performers will want to present adaptations to users, and in doing this they should consider the **Output Characteristics** of such adaptations. For example, learning activity recommendations should be carefully planned and tailored in a way that prevents interrupting the user's current task or leaking sensitive information.

The collection of vast amounts of information also raises question about **Data Location and Ownership**. TLA performers should use client-side methods for context data, and to allow users to designate a "data steward" to manage their data in accordance with their privacy preferences. TLA performers also allow users to take their data with them as they move between employers.

Moving to social and organizational aspects, TLA performers should be careful regarding **Data Sharing**. Specifically, they should make users aware of what information collected about them is used and how, and act ethically and responsibly regarding research placement and promotion decisions.

Finally, TLA performers should think carefully about providing **Privacy Support Mechanisms**, especially since the traditional paradigms of "notice and control" and "privacy nudging" seem to have failed. We propose user-tailored privacy as a way to give users more accessible yet still customizable privacy controls. Given the complexity of privacy in advanced distributed learning systems, upcoming versions of this document will delve deeper into the idea of user-tailored privacy as a decision-support mechanism for TLA.

References

1. 6 new facts about Facebook: 2014. http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/. Accessed: 2014-04-25.
2. 12 Surprising A/B Test Results to Stop You Making Assumptions: 2012. http://unbounce.com/a-b-testing/shocking-results/. Accessed: 2014-02-16.
3. Ackerman, M.S., Cranor, L.F. and Reagle, J. 1999. Privacy in e-commerce: examining user scenarios and privacy preferences. Proceedings of the 1st ACM conference on electronic commerce (Denver, CO, 1999), 1–8. DOI= http://dx.doi.org/10.1145/336992.336995.
4. Acquisti, A. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. IEEE Security and Privacy. 7, (Nov. 2009), 82–85. DOI= http://dx.doi.org/10.1109/MSP.2009.163.
5. Acquisti, A. and Gross, R. 2006. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. Privacy Enhancing Technologies. G. Danezis and P. Golle, eds. Springer Berlin / Heidelberg. 36–58.
6. Acquisti, A., John, L. and Loewenstein, G. 2009. What is privacy worth? Proceedings of the Twenty First Workshop on Information Systems and Economics (Phoenix, AZ, Dec. 2009).
7. Acquisti, A., John, L.K. and Loewenstein, G. 2012. The Impact of Relative Standards on the Propensity to Disclose. Journal of Marketing Research. 49, 2 (2012), 160–174. DOI= http://dx.doi.org/10.1509/jmr.09.0215.
8. Adjerid, I., Acquisti, A., Brandimarte, L. and Loewenstein, G. 2013. Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. Proceedings of the Ninth Symposium on Usable Privacy and Security (New York, NY, USA, 2013), 9:1–9:11. DOI= http://dx.doi.org/10.1145/2501604.2501613.
9. Adkinson, W.F., Eisenach, J.A. and Lenard, T.M. 2002. Privacy Online: A Report on the Information Practices and Policies of Commercial Web Sites. Privacy & Freedom Foundation.
10. Adomavicius, G. and Tuzhilin, A. 2011. Context-Aware Recommender Systems. Recommender Systems Handbook. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer US. 217–253.
11. Alemdar, H. and Ersoy, C. 2010. Wireless sensor networks for healthcare: A survey. Computer Networks. 54, 15 (Oct. 2010), 2688–2710. DOI= http://dx.doi.org/10.1016/j.comnet.2010.05.003.
12. Altman, I. 1975. The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding. Brooks/Cole Publishing Company, Monterey, CA.
13. Amatriain, X., Pujol, J.M., Tintarev, N. and Oliver, N. 2009. Rate It Again: Increasing Recommendation Accuracy by User Re-rating. Proceedings of the Third ACM Conference on Recommender Systems (New York, NY, USA, 2009), 173–180. DOI= http://dx.doi.org/10.1145/1639714.1639744.
14. Andrade, E.B., Kaltcheva, V. and Weitz, B. 2002. Self-Disclosure on the Web: The Impact of Privacy Policy, Reward, and Company Reputation. Advances in Consumer Research. S.M. Broniarczyk and K. Nakamoto, eds. Association for Consumer Research. 350–353.
15. Angst, C.M. and Agarwal, R. 2009. Adoption of electronic health records in the presence of privacy concerns: the elaboration likelihood model and individual persuasion. MIS Quarterly. 33, 2 (Jun. 2009), 339–370.
16. Antón, A.I., Earp, J.B., He, Q., Stufflebeam, W., Bolchini, D. and Jensen, C. 2004. Financial privacy policies and the need for standardization. IEEE Security Privacy. 2, 2 (Mar. 2004), 36–45. DOI= http://dx.doi.org/10.1109/MSECP.2004.1281243.
17. Aoki, P. and Woodruff, A. 2005. Making space for stories. Proceedings of the SIGCHI conference on Human factors in computing systems. (2005), 181–190. DOI= http://dx.doi.org/10.1145/1054972.1054998.
18. Ardagna, C.A., Vimercati, S.D.C. di, Pedrini, E. and Samarati, P. 2011. Privacy-Aware Access Control System: Evaluation and Decision. Digital Privacy. J. Camenisch, R. Leenes, and D. Sommer, eds. Springer Berlin Heidelberg. 377–395.
19. Arlein, R.M., Jai, B., Jakobsson, M., Monrose, F. and Reiter, M.K. 2000. Privacy-Preserving Global Customization. 2nd ACM Conference on Electronic Commerce (Minneapolis, MN, 2000), 176–184.
20. Atallah, M.J. and Du, W. 2001. Secure Multi-party Computational Geometry. Algorithms and Data Structures (Aug. 2001), 165–179. DOI= http://dx.doi.org/10.1007/3-540-44634-6_16.

21. Awad, N.F. and Krishnan, M.S. 2006. The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled Online for Personalization. MIS Quarterly. 30, 1 (Mar. 2006), 13–28.
22. Bain, P. and Taylor, P. 2000. Entrapped by the "electronic panopticon"? Worker resistance in the call centre. New Technology, Work and Employment. 15, 1 (Mar. 2000), 2–18. DOI= http://dx.doi.org/10.1111/1468-005X.00061.
23. Balebako, R., Leon, P.G., Mugan, J., Acquisti, A., Cranor, L.F. and Sadeh, N. 2011. Nudging users towards privacy on mobile devices. CHI 2011 workshop on Persuasion, Influence, Nudge and Coercion Through Mobile Devices (Vancouver, Canada, 2011), 23–26.
24. Bansal, G., Zahedi, F. and Gefen, D. 2008. The Moderating Influence of Privacy Concern on the Efficacy of Privacy Assurance Mechanisms for Building Trust: A Multiple-Context Investigation. ICIS 2008 Proceedings (Paris, France, 2008).
25. Baraglia, R., Lucchese, C., Orlando, S., Serrano', M. and Silvestri, F. 2006. A Privacy Preserving Web Recommender System. Proceedings of the 2006 ACM Symposium on Applied Computing (New York, NY, USA, 2006), 559–563. DOI= http://dx.doi.org/10.1145/1141277.1141407.
26. Barcena, M.B., Wueest, C. and Lau, H. 2014. How safe is your quantified self? Tracking, monitoring, and wearable tech. Symantech.
27. Beckwith, R. and Mainwaring, S. 2005. Privacy: Personal information, threats, and technologies. Proceedings of the Proceedings 2005 IEEE International Symposium on Technology and Society (2005), 9–16.
28. Benisch, M., Kelley, P.G., Sadeh, N. and Cranor, L.F. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. Personal Ubiquitous Computing. 15, 7 (Oct. 2011), 679–694. DOI= http://dx.doi.org/10.1007/s00779-010-0346-0.
29. Benjamin, J. 2017. A hands-on look at Theater Mode for Apple Watch in watchOS 3.2 beta 1 [Video]. 9to5Mac.
30. Berendt, B., Günther, O. and Spiekermann, S. 2005. Privacy in E-Commerce: Stated Preferences vs. Actual Behavior. Communications of the ACM. 48, 4 (2005), 101–106. DOI= http://dx.doi.org/10.1145/1053291.1053295.
31. Bergmann, M. 2009. Testing Privacy Awareness. The Future of Identity in the Information Society. V. Matyáš, S. Fischer-Hübner, D. Cvrček, and P. Švenda, eds. Springer Berlin Heidelberg. 237–253.
32. Besmer, A., Watson, J. and Lipford, H.R. 2010. The impact of social navigation on privacy policy configuration. Proceedings of the Sixth Symposium on Usable Privacy and Security (Redmond, Washington, Jul. 2010), 7:1-7:10. DOI= http://dx.doi.org/10.1145/1837110.1837120.
33. Best Practice Gone Bad: 4 Shocking A/B Tests: 2012. http://www.getelastic.com/best-practice-gone-bad-4-shocking-ab-tests/. Accessed: 2013-01-02.
34. Bettman, J.R., Luce, M.F. and Payne, J.W. 1998. Constructive consumer choice processes. Journal of consumer research. 25, 3 (1998), 187–217.
35. Bhatnagar, A., Misra, S. and Rao, H.R. 2000. On risk, convenience, and Internet shopping behavior. Communications of the ACM. 43, 11 (Nov. 2000), 98–105. DOI= http://dx.doi.org/10.1145/353360.353371.
36. Birnholtz, J., Guillory, J., Hancock, J. and Bazarova, N. 2010. "on my way": deceptive texting and interpersonal awareness narratives. Proceedings of the 2010 ACM conference on Computer supported cooperative work (New York, NY, USA, 2010), 1–4. DOI= http://dx.doi.org/10.1145/1718918.1718920.
37. Bonneau, J. and Preibusch, S. 2010. The Privacy Jungle:On the Market for Data Protection in Social Networks. Economics of Information Security and Privacy. T. Moore, D. Pym, and C. Ioannidis, eds. Springer US. 121–167.
38. Breslow, L., Pritchard, D., DeBoer, J., Stump, G., Ho, A. and Seaton, D. 2013. Studying learning in the worldwide classroom: Research into edX's first MOOC. Research & Practice in Assessment. 8, (2013), 13–25.
39. Brodie, C., Karat, C.-M. and Karat, J. 2004. How Personalization of an E-Commerce Website Affects Consumer Trust. Designing Personalized User Experience for eCommerce. C.-M. Karat, J.O. Blom, and J. Karat, eds. Kluwer Academic Publishers. 185–206.
40. Brudy, F., Ledo, D., Greenberg, S. and Butz, A. 2014. Is Anyone Looking? Mitigating Shoulder Surfing on Public Displays Through Awareness and Protection. Proceedings of The International Symposium on Pervasive Displays (New York, NY, USA, 2014), 1:1–1:6. DOI= http://dx.doi.org/10.1145/2611009.2611028.

41. Brusilovsky, P. 2007. Adaptive Navigation Support. The Adaptive Web: Methods and Strategies of Web Personalization. Springer Berlin / Heidelberg. 263–290.
42. Bucher, K.T. and Manning, M.L. 2004. Bringing Graphic Novels into a School's Curriculum. The Clearing House. 78, 2 (Nov. 2004), 67–72.
43. Bulgurcu, B. 2012. Understanding the information privacy-related perceptions and behaviors of an online social network user. University of British Columbia.
44. Cacioppo, J.T., Petty, R.E., Kao, C.F. and Rodriguez, R. 1986. Central and peripheral routes to persuasion: An individual difference perspective. Journal of personality and social psychology. 51, 5 (1986), 1032.
45. Calandrino, J.A., Kilzer, A., Narayanan, A., Felten, E.W. and Shmatikov, V. 2011. You Might Also Like: Privacy Risks of Collaborative Filtering. Proceedings of the 2011 IEEE Symposium on Security and Privacy (May 2011), 231–246. DOI= http://dx.doi.org/10.1109/SP.2011.40.
46. Canny, J. 2002. Collaborative Filtering with Privacy via Factor Analysis. 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (Tampere, Finland, Aug. 2002), 238–245. DOI= http://dx.doi.org/10.1145/564376.564419.
47. Cantador, I., Fernández-Tobías, I., Berkovsky, S. and Cremonesi, P. 2015. Cross-Domain Recommender Systems. Recommender Systems Handbook. F. Ricci, L. Rokach, and B. Shapira, eds. Springer US. 919–959. DOI= http://dx.doi.org/10.1007/978-1-4899-7637-6_27.
48. Cao, J. and Everard, A. 2007. Influence of Culture on Attitude Towards Instant Messaging: Balance Between Awareness and Privacy. Human-Computer Interaction. Interaction Platforms and Techniques. J. Jacko, ed. Springer Berlin / Heidelberg. 236–240.
49. Cassel, L.N. and Wolz, U. 2001. Client Side Personalization. DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries (Dublin, Ireland, 2001), 8–12.
50. Cate, F.H. 2006. The Failure of Fair Information Practice Principles. Consumer Protection in the Age of the "Information Economy." J.K. Winn, ed. Ashgate Publishing Company.
51. Cavoukian, A. 2009. Privacy by Design. Information and Privacy Commissioner of Ontario, Canada.
52. Cavusoglu, H., Phan, T. and Cavusoglu, H. 2013. Privacy Controls and Content Sharing Patterns of Online Social Network Users: A Natural Experiment. ICIS 2013 Proceedings (Milan, Italy, 2013).
53. Chalykoff, J. and Kochan, T.A. 1989. Computer-Aided Monitoring: Its Influence on Employee Job Satisfaction and Turnover. Personnel Psychology. 42, 4 (Dec. 1989), 807–834. DOI= http://dx.doi.org/10.1111/j.1744-6570.1989.tb00676.x.
54. Chellappa, R.K. and Sin, R.G. 2005. Personalization versus privacy: An empirical examination of the online consumer's dilemma. Information Technology and Management. 6, 2 (2005), 181–202. DOI= http://dx.doi.org/10.1007/s10799-005-5879-y.
55. Chen, D., Fraiberger, S.P., Moakler, R. and Provost, F. 2015. Enhancing Transparency and Control when Drawing Data-Driven Inferences about Individuals.
56. Chen, L. and Pu, P. 2011. Critiquing-based recommenders: survey and emerging trends. User Modeling and User-Adapted Interaction. 22, 1–2 (Oct. 2011), 125–150. DOI= http://dx.doi.org/10.1007/s11257-011-9108-6.
57. Chen, Y.-F. 2008. Herd behavior in purchasing books online. Computers in Human Behavior. 24, 5 (Sep. 2008), 1977–1992. DOI= http://dx.doi.org/10.1016/j.chb.2007.08.004.
58. Cho, D., Kim, S. and Acquisti, A. 2012. Empirical analysis of online anonymity and user behaviors: the impact of real name policy. 2012 45th Hawaii International Conference on System Science (Jan. 2012), 3041–3050. DOI= http://dx.doi.org/10.1109/HICSS.2012.241.
59. Cockcroft, S. and Rekker, S. 2015. The relationship between culture and information privacy policy. Electronic Markets. (Jul. 2015), 1–18. DOI= http://dx.doi.org/10.1007/s12525-015-0195-9.
60. Coleman, G. 2012. Phreaks, hackers, and trolls: The politics of transgression and spectacle. The social media reader. M. Mandiberg, ed. 99–119.
61. Compañó, R. and Lusoli, W. 2010. The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas. Economics of Information Security and Privacy. T. Moore, D. Pym, and C. Ioannidis, eds. Springer US. 169–185.
62. Connolly, T. and Zeelenberg, M. 2002. Regret in decision making. Current directions in psychological science. 11, 6 (2002), 212–216.
63. Consolvo, S., McDonald, D.W., Toscos, T., Chen, M.Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., Smith, I. and Landay, J.A. 2008. Activity Sensing in the Wild: A Field Trial of Ubifit

Garden. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (New York, NY, USA, 2008), 1797–1806. DOI= http://dx.doi.org/10.1145/1357054.1357335.

64. Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J. and Powledge, P. 2005. Location disclosure to social relations: why, when, & what people want to share. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Portland, OR, 2005), 81–90. DOI= http://dx.doi.org/10.1145/1054972.1054985.
65. Cosley, D., Lam, S.K., Albert, I., Konstan, J.A. and Riedl, J. 2003. Is Seeing Believing?: How Recommender System Interfaces Affect Users' Opinions. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Ft. Lauderdale, FL, 2003), 585–592. DOI= http://dx.doi.org/10.1145/642611.642713.
66. Cramer, H., Evers, V., Ramlal, S., Someren, M., Rutledge, L., Stash, N., Aroyo, L. and Wielinga, B. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. User Modeling and User-Adapted Interaction. 18, 5 (Aug. 2008), 455–496. DOI= http://dx.doi.org/10.1007/s11257-008-9051-3.
67. Cremonesi, P., Garzotto, F. and Turrin, R. 2012. Investigating the Persuasion Potential of Recommender Systems from a Quality Perspective: An Empirical Study. ACM Transactions on Interactive Intelligent Systems. 2, 2 (Jun. 2012), 11:1–11:41. DOI= http://dx.doi.org/10.1145/2209310.2209314.
68. Cremonini, L. and Valeri, L. 2003. Benchmarking Security and Trust in Europe and the US. Technical Report #IST 2000-26746. RAND Europe.
69. Culnan, M.J. and Armstrong, P.K. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. Organization Science. 10, 1 (1999), 104–115. DOI= http://dx.doi.org/10.1287/orsc.10.1.104.
70. Culnan, M.J. and Bies, R.J. 2003. Consumer Privacy: Balancing Economic and Justice Considerations. Journal of Social Issues. 59, 2 (2003), 323–342. DOI= http://dx.doi.org/10.1111/1540-4560.00067.
71. Cutrell, E., Czerwinski, M. and Horvitz, E. 2001. Notification, Disruption, and Memory: Effects of Messaging Interruptions on Memory and Performance. INTERACT (2001), 263–269.
72. Cutrell, E.B., Czerwinski, M. and Horvitz, E. 2000. Effects of Instant Messaging Interruptions on Computing Tasks. CHI '00 Extended Abstracts on Human Factors in Computing Systems (New York, NY, USA, 2000), 99–100. DOI= http://dx.doi.org/10.1145/633292.633351.
73. Czarkowski, M. and Kay, J. 2000. Bringing Scrutability to Adaptive Hypertext Teaching. Intelligent Tutoring Systems 2000. G. Gauthier, C. Frasson, and K. VanLehn, eds. Springer. 423–433.
74. Czarkowski, M. and Kay, J. 2003. How to Give the User a Sense of Control Over the Personalization of AH? AH2003: Workshop on Adaptive Hypermedia and Adaptive Web-Based Systems (Budapest, Hungary; Johnstown, PA; Nottingham, England, 2003).
75. Dabbah, M.A., Woo, W.L. and Dlay, S.S. 2007. Secure Authentication for Face Recognition. 2007 IEEE Symposium on Computational Intelligence in Image and Signal Processing (Apr. 2007), 121–126. DOI= http://dx.doi.org/10.1109/CIISP.2007.369304.
76. Dass, S.C., Zhu, Y. and Jain, A.K. 2006. Validating a Biometric Authentication System: Sample Size Requirements. IEEE Transactions on Pattern Analysis and Machine Intelligence. 28, 12 (Dec. 2006), 1902–1319. DOI= http://dx.doi.org/10.1109/TPAMI.2006.255.
77. Deci, E.L., Koestner, R. and Ryan, R.M. 1999. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. Psychological Bulletin. 125, 6 (1999), 627–668. DOI= http://dx.doi.org/10.1037/0033-2909.125.6.627.
78. Deshpande, M. and Karypis, G. 2004. Item-based top-N Recommendation Algorithms. ACM Trans. Inf. Syst. 22, 1 (Jan. 2004), 143–177. DOI= http://dx.doi.org/10.1145/963770.963776.
79. Dinev, T., Bellotto, M., Hart, P., Russo, V., Serra, I. and Colautti, C. 2006. Privacy calculus model in e-commerce - a study of Italy and the United States. European Journal of Information Systems. 15, 4 (Aug. 2006), 389–402. DOI= http://dx.doi.org/http://dx.doi.org.janus.libr.tue.nl/10.1057/palgrave.ejis.3000590.
80. Dinev, T. and Hart, P. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. Information Systems Research. 17, 1 (Mar. 2006), 61–80. DOI= http://dx.doi.org/10.1287/isre.1060.0080.
81. Dong, C., Jin, H. and Knijnenburg, B.P. 2015. Predicting Privacy Behavior on Online Social Networks. Ninth International AAAI Conference on Web and Social Media (Apr. 2015), 91–100.
82. Donker, T., Petrie, K., Proudfoot, J., Clarke, J., Birch, M.-R. and Christensen, H. 2013. Smartphones for Smarter Delivery of Mental Health Programs: A Systematic Review. Journal of Medical Internet Research. 15, 11 (2013), e247. DOI= http://dx.doi.org/10.2196/jmir.2791.

83. Donley, M.B. 2007. Department of Defense Privacy Program. Technical Report #DoD 5400.11-R. Department of Defense.
84. Duhigg, C. 2012. How Companies Learn Your Secrets. The New York Times.
85. Duncan, R., Smith, M.J. and Levitz, P. 2009. The Power of Comics: History, Form and Culture. Bloomsbury Academic.
86. Eargle, D., Galletta, D., Kirwan, B. and Vance, T. 2015. Integrating Facial Threat Signals into Security Messages: An Extension of Media Naturalness Theory to an Information Security Context. Proc. IFIP Dewald Roode (Newark, DE, 2015).
87. Egelman, S., Tsai, J., Cranor, L.F. and Acquisti, A. 2009. Timing is everything?: the effects of timing and placement of online privacy indicators. Proceedings of the 27th international conference on Human factors in computing systems (2009), 319–328.
88. Ekstrand, M.D. and Willemsen, M.C. 2016. Behaviorism is Not Enough: Better Recommendations Through Listening to Users. Proceedings of the 10th ACM Conference on Recommender Systems (New York, NY, USA, 2016), 221–224. DOI= http://dx.doi.org/10.1145/2959100.2959179.
89. Ellison, N.B., Vitak, J., Steinfield, C., Gray, R. and Lampe, C. 2011. Negotiating privacy concerns and social capital needs in a social media environment. Privacy online. Springer. 19–32.
90. 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. Technical Report #2012/0010 (COD).
91. Facebook & your privacy: Who sees the data you share on the biggest social network? 2012. http://www.consumerreports.org/cro/magazine/2012/06/facebook-your-privacy. Accessed: 2012-05-13.
92. Fagin, R., Naor, M. and Winkler, P. 1996. Comparing Information Without Leaking It. Commun. ACM. 39, 5 (May 1996), 77–85. DOI= http://dx.doi.org/10.1145/229459.229469.
93. Farzan, R. and Brusilovsky, P. 2011. Encouraging user participation in a course recommender system: An impact on user behavior. Computers in Human Behavior. 27, 1 (Jan. 2011), 276–284. DOI= http://dx.doi.org/10.1016/j.chb.2010.08.005.
94. Farzan, R. and Brusilovsky, P. 2006. Social Navigation Support in a Course Recommendation System. Adaptive Hypermedia and Adaptive Web-Based Systems (Jun. 2006), 91–100. DOI= http://dx.doi.org/10.1007/11768012_11.
95. Featherman, M.S. and Pavlou, P.A. 2003. Predicting e-services adoption: a perceived risk facets perspective. International Journal of Human-Computer Studies. 59, 4 (Oct. 2003), 451–474. DOI= http://dx.doi.org/10.1016/S1071-5819(03)00111-3.
96. Felfernig, A. 2007. Knowledge-Based Recommender Technologies for Marketing and Sales. International Journal of Pattern Recognition and Artificial Intelligence. 21, 2 (2007), 333–354. DOI= http://dx.doi.org/10.1142/S0218001407005417.
97. Fishbein, M. and Ajzen, I. 1975. Belief, attitude, intention, and behavior: an introduction to theory and research. Addison-Wesley Pub. Co.
98. Folsom-Kovarik, J.T. and Raybourn, E.M. 2016. Total Learning Architecture (TLA) Enables Next-generation Learning via Meta-adaptation. Interservice/Industry Training, Simulation, and Education Conference Proceedings (Orlando, FL, Nov. 2016).
99. Fox, S. and Duggan, M. 2013. Health Online 2013. Pew Research Center.
100. Fox, T., Grunst, G. and Quast, K.-J. 1994. HyPLAN: A Context-Sensitive Hypermedia Help System. Adaptive User Support. R. Oppermann, ed. 126–193.
101. de Freitas, S. and Levene, M. 2003. Evaluating the development of wearable devices, personal data assistants and the use of other mobile devices in further and higher education institutions. JISC Technology and Standards Watch Report. TSW030 (Jun. 2003), 1–21.
102. Frey, N. and Fisher, D.B. 2008. Teaching Visual Literacy: Using Comic Books, Graphic Novels, Anime, Cartoons, and More to Develop Comprehension and Thinking Skills. Corwin.
103. FTC 2000. Privacy Online: Fair Information Practices in the Electronic Marketplace. A Report to Congress. Federal Trade Commission.
104. FTC 2012. Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers. Federal Trade Commission.

105. van de Garde-Perik, E., Markopoulos, P., de Ruyter, B., Eggen, B. and Ijsselsteijn, W. 2008. Investigating Privacy Attitudes and Behavior in Relation to Personalization. Social Science Computer Review. 26, 1 (Feb. 2008), 20–43. DOI= http://dx.doi.org/10.1177/0894439307307682.
106. Garg, A. 2015. Multi-device learning. Training & Development. 42, 4 (Aug. 2015), 10.
107. Gee, J.P. 2003. What Video Games Have to Teach Us About Learning and Literacy. Comput. Entertain. 1, 1 (Oct. 2003), 20–20. DOI= http://dx.doi.org/10.1145/950566.950595.
108. Gena, C., Brogi, R., Cena, F. and Vernero, F. 2011. The Impact of Rating Scales on User's Rating Behavior. User Modeling, Adaption and Personalization. J.A. Konstan, R. Conejo, J.L. Marzo, and N. Oliver, eds. Springer Berlin Heidelberg. 123–134.
109. Giles, D.C. and Newbold, J. 2011. Self- and Other-Diagnosis in User-Led Mental Health Online Communities. Qualitative Health Research. 21, 3 (Mar. 2011), 419–428. DOI= http://dx.doi.org/10.1177/1049732310381388.
110. Gomez-Uribe, C.A. and Hunt, N. 2015. The Netflix Recommender System: Algorithms, Business Value, and Innovation. ACM Trans. Manage. Inf. Syst. 6, 4 (Dec. 2015), 13:1–13:19. DOI= http://dx.doi.org/10.1145/2843948.
111. Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D. and Konstan, J. 2005. Stopping Spyware at the Gate: A User Study of Privacy, Notice and Spyware. Proceedings of the 2005 Symposium on Usable Privacy and Security (New York, NY, USA, 2005), 43–52. DOI= http://dx.doi.org/10.1145/1073001.1073006.
112. Gootman, S. 2016. OPM Hack: The Most Dangerous Threat to the Federal Government Today. Journal of Applied Security Research. 11, 4 (Oct. 2016), 517–525. DOI= http://dx.doi.org/10.1080/19361610.2016.1211876.
113. Green, M.J. and Myers, K.R. 2010. Graphic medicine: use of comics in medical education and patient care. BMJ. 340, (Jan. 2010), c863. DOI= http://dx.doi.org/10.1136/bmj.c863.
114. Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C. and Höllerer, T. 2010. SmallWorlds: Visualizing Social Recommendations. Computer Graphics Forum. 29, 3 (2010), 833–842. DOI= http://dx.doi.org/10.1111/j.1467-8659.2009.01679.x.
115. Gretzel, U. and Fesenmaier, D.R. 2006. Persuasion in Recommender Systems. International Journal of Electronic Commerce. 11, 2 (Dec. 2006), 81–100. DOI= http://dx.doi.org/10.2753/JEC1086-4415110204.
116. Gross, R. and Acquisti, A. 2005. Information Revelation and Privacy in Online Social Networks. Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society (Alexandria, VA, USA, 2005), 71–80. DOI= http://dx.doi.org/10.1145/1102199.1102214.
117. Grünewald, F., Meinel, C., Totschnig, M. and Willems, C. 2013. Designing MOOCs for the Support of Multiple Learning Styles. Scaling up Learning for Sustained Impact (Sep. 2013), 371–382. DOI= http://dx.doi.org/10.1007/978-3-642-40814-4_29.
118. Gueldenzoph, L.E. and May, G.L. 2002. Collaborative peer evaluation: Best practices for group member assessments. Business Communication Quarterly. 65, 1 (Mar. 2002), 9–21.
119. Guo, H., Chen, J., Wu, W. and Wang, W. 2009. Personalization as a service: the architecture and a case study. Proceedings of the first international workshop on Cloud data management (2009), 1–8.
120. Guy, I., Ronen, I. and Wilcox, E. 2009. Do you know?: recommending people to invite into your social network. Proceedings of the 14th international conference on Intelligent user interfaces (New York, NY, USA, 2009), 77–86. DOI= http://dx.doi.org/10.1145/1502650.1502664.
121. Haddadi, H. and Brown, I. 2014. Quantified self and the privacy challenge. Technology Law Futures. (2014).
122. Haddadi, H., Ofli, F., Mejova, Y., Weber, I. and Srivastava, J. 2015. 360-degree Quantified Self. 2015 International Conference on Healthcare Informatics (Oct. 2015), 587–592. DOI= http://dx.doi.org/10.1109/ICHI.2015.95.
123. Hagel, J. and Rayport, J.F. 1999. The Coming Battle for Customer Information. Creating value in the network economy. Harvard Business School Press. 159–171.
124. Hampton, K., Goulet, L.S., Marlow, C. and Rainie, L. 2012. Why most Facebook users get more than they give. Pew Internet & American Life Project.
125. Hancock, J., Birnholtz, J., Bazarova, N., Guillory, J., Perlin, J. and Amos, B. 2009. Butler lies: awareness, deception and design. Proceedings of the 27th international conference on Human factors in computing systems (Boston, MA, USA, 2009), 517–526. DOI= http://dx.doi.org/10.1145/1518701.1518782.

126. Hancock, J.T., Thom-Santelli, J. and Ritchie, T. 2004. Deception and design: the impact of communication technology on lying behavior. Proceedings of the SIGCHI conference on Human factors in computing systems (New York, NY, USA, 2004), 129–134. DOI= http://dx.doi.org/10.1145/985692.985709.
127. Hang, A., von Zezschwitz, E., De Luca, A. and Hussmann, H. 2012. Too Much Information!: User Attitudes Towards Smartphone Sharing. Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design (New York, NY, USA, 2012), 284–287. DOI= http://dx.doi.org/10.1145/2399016.2399061.
128. Hann, I.-H., Hui, K.-L., Lee, S.-Y. and Png, I. 2007. Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach. Journal of Management Information Systems. 24, 2 (Oct. 2007), 13–42. DOI= http://dx.doi.org/10.2753/MIS0742-1222240202.
129. Harris 2001. Privacy Notices Research: Final Results. Technical Report #Study No. 15338. Harris Interactive, Inc.
130. Harris Interactive inc. 2000. A Survey of Consumer Privacy Attitudes and Behaviors. Harris Interactive, Inc.
131. Harris, L., Westin, A.F. and associates 2003. Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits. Equifax Inc.
132. Harrison, C. and Hudson, S.E. 2011. A New Angle on Cheap LCDs: Making Positive Use of Optical Distortion. Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (New York, NY, USA, 2011), 537–540. DOI= http://dx.doi.org/10.1145/2047196.2047266.
133. Herlocker, J.L., Konstan, J.A. and Riedl, J. 2000. Explaining collaborative filtering recommendations. Proc. of the 2000 ACM conference on Computer supported cooperative work (Philadelphia, PA, 2000), 241–250. DOI= http://dx.doi.org/10.1145/358916.358995.
134. Hew, K.F. and Cheung, W.S. 2014. Students' and instructors' use of massive open online courses (MOOCs): Motivations and challenges. Educational Research Review. 12, (Jun. 2014), 45–58. DOI= http://dx.doi.org/10.1016/j.edurev.2014.05.001.
135. Ho, S.Y. and Tam, K. 2006. Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes. MIS Quarterly. 30, 4 (Dec. 2006), 865–890.
136. Hoepman, J.-H. 2014. Privacy Design Strategies. ICT Systems Security and Privacy Protection. N. Cuppens-Boulahia, F. Cuppens, S. Jajodia, A.A.E. Kalam, and T. Sans, eds. Springer Berlin Heidelberg. 446–459. DOI= http://dx.doi.org/10.1007/978-3-642-55415-5_38.
137. Hoffman, D.L., Novak, T.P. and Peralta, M. 1999. Building consumer trust online. Communications of the ACM. 42, 4 (Apr. 1999), 80–85. DOI= http://dx.doi.org/10.1145/299157.299175.
138. Hogg, M.A. 2000. Subjective Uncertainty Reduction through Self-categorization: A Motivational Theory of Social Identity Processes. European Review of Social Psychology. 11, 1 (Jan. 2000), 223–255. DOI= http://dx.doi.org/10.1080/14792772043000040.
139. Holzinger, A., Nischelwitzer, A. and Meisenberger, M. 2005. Mobile phones as a challenge for m-learning: examples for mobile interactive learning objects (MILOs). Third IEEE International Conference on Pervasive Computing and Communications Workshops (Mar. 2005), 307–311. DOI= http://dx.doi.org/10.1109/PERCOMW.2005.59.
140. Höök, K., Karlgren, J., Wærn, A., Dahlbäck, N., Jansson, C., Karlgren, K. and Lemaire, B. 1996. A glass box approach to adaptive hypermedia. User Modeling and User-Adapted Interaction. 6, 2–3 (Jul. 1996), 157–184. DOI= http://dx.doi.org/10.1007/BF00143966.
141. Horvitz, E., Koch, P. and Apacible, J. 2004. BusyBody: Creating and Fielding Personalized Models of the Cost of Interruption. Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (New York, NY, USA, 2004), 507–510. DOI= http://dx.doi.org/10.1145/1031607.1031690.
142. How Privacy Policy Affects Sign-Ups – Surprising Data From 4 A/B Tests: 2013. http://contentverve.com/sign-up-privacy-policy-tests/. Accessed: 2013-05-28.
143. Hsu, C. 2006. Privacy concerns, privacy practices and web site categories: Toward a situational paradigm. Online Information Review. 30, 5 (2006), 569–586. DOI= http://dx.doi.org/10.1108/14684520610706433.
144. Hui, K.-L., Tan, B.C.Y. and Goh, C.-Y. 2006. Online information disclosure: Motivators and measurements. ACM Transactions on Internet Technology. 6, 4 (Nov. 2006), 415–441. DOI= http://dx.doi.org/10.1145/1183463.1183467.
145. Hui, K.-L., Teo, H.H. and Lee, S.-Y.T. 2007. The Value of Privacy Assurance: An Exploratory Field Experiment. MIS Quarterly. 31, 1 (Mar. 2007), 19–33.

146. Inman, J.J. and Zeelenberg, M. 2002. Regret in repeat purchase versus switching decisions: The attenuating role of decision justifiability. Journal of Consumer Research. 29, 1 (2002), 116–128.
147. Internet Society 2012. Global Internet User Survey 2012. Internet Society.
148. Jacoby, J. and Kaplan, L.B. 1972. The Components of Perceived Risk. Proceedings of the Third Annual Conference of the Association for Consumer Research (Chicago, IL, 1972), 382–393.
149. Jain, A.K., Hong, L., Pankanti, S. and Bolle, R. 1997. An identity-authentication system using fingerprints. Proceedings of the IEEE. 85, 9 (Sep. 1997), 1365–1388. DOI= http://dx.doi.org/10.1109/5.628674.
150. Jakobsson, M. and Yung, M. 1996. Proving Without Knowing: On Oblivious, Agnostic and Blindolded Provers. Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology (London, UK, UK, 1996), 186–200.
151. Jameson, A., Willemsen, M.C., Felfernig, A., de Gemmis, M., Lops, P., Semeraro, G. and Chen, L. 2015. Human Decision Making and Recommender Systems. Recommender Systems Handbook, 2nd edition.
152. Jannach, D. and Adomavicius, G. 2016. Recommendations with a Purpose. Proceedings of the 10th ACM Conference on Recommender Systems (New York, NY, USA, 2016), 7–10. DOI= http://dx.doi.org/10.1145/2959100.2959186.
153. Jedrzejczyk, L., Price, B.A., Bandara, A.K. and Nuseibeh, B. 2010. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. Proceedings of the Sixth Symposium on Usable Privacy and Security (Redmond, Washington, 2010), 14:1-14:12. DOI= http://dx.doi.org/10.1145/1837110.1837129.
154. Jensen, C., Potts, C. and Jensen, C. 2005. Privacy Practices of Internet Users: Self-Reports versus Observed Behavior. International Journal of Human-Computer Studies. 63, 1–2 (2005), 203–227. DOI= http://dx.doi.org/10.1016/j.ijhcs.2005.04.019.
155. John, L.K., Acquisti, A. and Loewenstein, G. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. Journal of consumer research. 37, 5 (Feb. 2011), 858–873. DOI= http://dx.doi.org/10.1086/656423.
156. Johnson, D.W. |Johnson 1994. Learning Together and Alone. Cooperative, Competitive, and Individualistic Learning. Fourth Edition. Allyn and Bacon, 160 Gould Street, Needham Heights, MA 02194.
157. Johnson, E.J., Bellman, S. and Lohse, G.L. 2002. Defaults, Framing and Privacy: Why Opting In ≠ Opting Out. Marketing Letters. 13, 1 (2002), 5–15. DOI= http://dx.doi.org/10.1023/A:1015044207315.
158. Johnson, M., Egelman, S. and Bellovin, S.M. 2012. Facebook and privacy: it's complicated. Proceedings of the 8th Symposium on Usable Privacy and Security (Pittsburgh, PA, 2012), 9:1-9:15. DOI= http://dx.doi.org/10.1145/2335356.2335369.
159. Juels, A. 2001. Targeted Advertising ... and Privacy Too. Topics in Cryptology — CT-RSA 2001. D. Naccache, ed. Springer. 408–424.
160. Kairam, S., Brzozowski, M., Huffaker, D. and Chi, E. 2012. Talking in circles: Selective Sharing in Google+. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Austin, TX, 2012), 1065–1074. DOI= http://dx.doi.org/10.1145/2207676.2208552.
161. Kali, Y. and Ronen, M. 2005. Design Principles for Online Peer-evaluation: Fostering Objectivity. Proceedings of th 2005 Conference on Computer Support for Collaborative Learning: Learning 2005: The Next 10 Years! (Taipei, Taiwan, 2005), 247–251.
162. Karr-Wisniewski, P., Lipford, H. and Wilson, D. 2011. A New Social Order: Mechanisms for Social Network Site Boundary Regulation. (2011).
163. Kay, J. and Kummerfeld, B. 2013. Creating personalized systems that people can scrutinize and control: Drivers, principles and experience. ACM Transactions on Interactive Intelligent Systems. 2, 4 (Jan. 2013), 24:1–24:42. DOI= http://dx.doi.org/10.1145/2395123.2395129.
164. Kay, J. and Lum, A. 2005. Ontology-based user modelling for the Semantic Web. Workshop on Personalisation on the Semantic Web (Edinburgh, UK, 2005), 15–23.
165. Kay, M. and Terry, M. 2010. Textured Agreements: Re-envisioning Electronic Consent. Proceedings of the Sixth Symposium on Usable Privacy and Security (New York, NY, USA, 2010), 13:1–13:13. DOI= http://dx.doi.org/10.1145/1837110.1837127.
166. Kaye, A. 1992. Learning Together Apart. Collaborative Learning Through Computer Conferencing. A.R. Kaye, ed. Springer Berlin Heidelberg. 1–24. DOI= http://dx.doi.org/10.1007/978-3-642-77684-7_1.

167. Keane, M.A. 2008. Institutional Review Board Approaches to the Incidental Findings Problem. The Journal of Law, Medicine & Ethics. 36, 2 (Jun. 2008), 352–355. DOI= http://dx.doi.org/10.1111/j.1748-720X.2008.00279.x.
168. Kehr, F., Kowatsch, T., Wentzel, D. and Fleisch, E. 2015. Thinking Styles and Privacy Decisions: Need for Cognition, Faith into Intuition, and the Privacy Calculus. 12th International Conference on Wirtschaftsinformatik (Osnabrück, Germany, 2015).
169. Kehr, F., Wentzel, D. and Mayer, P. 2013. Rethinking the Privacy Calculus: On the Role of Dispositional Factors and Affect. ICIS 2013 Proceedings (Milan, Italy, 2013).
170. Keith, M.J., Babb, J.S., Lowry, P.B., Furner, C.P. and Abdullat, A. 2011. The Roles of Privacy Assurance, Network Effects, and Information Cascades in the Adoption of and Willingness to Pay for Location-Based Services with Mobile Applications. 2011 Dewald Roode Information Security Workshop (Blacksburg, VA, Sep. 2011).
171. Kelley, P.G., Cesca, L., Bresee, J. and Cranor, L.F. 2010. Standardizing privacy notices: an online study of the nutrition label approach. Proceedings of the 28th International Conference on Human Factors in Computing Systems, CHI 2010 (Atlanta, Georgia, 2010), 1573–1582. DOI= http://dx.doi.org/10.1145/1753326.1753561.
172. Kim, D.J., Ferrin, D.L. and Rao, H.R. 2008. A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents. Decision Support Systems. 44, 2 (Jan. 2008), 544–564. DOI= http://dx.doi.org/10.1016/j.dss.2007.07.001.
173. Kluver, D., Nguyen, T.T., Ekstrand, M., Sen, S. and Riedl, J. 2012. How Many Bits Per Rating? Proceedings of the Sixth ACM Conference on Recommender Systems (New York, NY, USA, 2012), 99–106. DOI= http://dx.doi.org/10.1145/2365952.2365974.
174. Knight, L. 2010. Social experiment:online privacy vs. personalization paradox. Upshot.
175. Knijnenburg, B.P. 2015. A user-tailored approach to privacy decision support. University of California, Irvine.
176. Knijnenburg, B.P. 2009. Adaptive advice: adapting a recommender system for energy-saving behaviors to personal differences in decision-making. Eindhoven University of Technology.
177. Knijnenburg, B.P., Bostandjiev, S., O'Donovan, J. and Kobsa, A. 2012. Inspectability and control in social recommenders. Proceedings of the sixth ACM conference on Recommender systems (New York, NY, USA, 2012), 43–50. DOI= http://dx.doi.org/10.1145/2365952.2365966.
178. Knijnenburg, B.P. and Cherry, D. 2016. Comics as a Medium for Privacy Notices. SOUPS 2016 workshop on the Future of Privacy Notices and Indicators (Denver, CO, Jun. 2016).
179. Knijnenburg, B.P. and Jin, H. 2013. The Persuasive Effect of Privacy Recommendations. Twelfth Annual Workshop on HCI Research in MIS (Milan, Italy, 2013), 16:1-16:5.
180. Knijnenburg, B.P. and Kobsa, A. 2013. Helping users with information disclosure decisions: potential for adaptation. Proceedings of the 2013 ACM international conference on Intelligent User Interfaces (Santa Monica, CA, Mar. 2013), 407–416. DOI= http://dx.doi.org/10.1145/2449396.2449448.
181. Knijnenburg, B.P. and Kobsa, A. 2014. Increasing Sharing Tendency Without Reducing Satisfaction: Finding the Best Privacy-Settings User Interface for Social Networks. ICIS 2014 Proceedings (Auckland, New Zealand, 2014).
182. Knijnenburg, B.P. and Kobsa, A. 2013. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. ACM Transactions on Interactive Intelligent Systems. 3, 3 (2013), 20:1-20:23. DOI= http://dx.doi.org/10.1145/2499670.
183. Knijnenburg, B.P., Kobsa, A. and Jin, H. 2013. Counteracting the Negative Effect of Form Auto-completion on the Privacy Calculus. ICIS 2013 Proceedings (Milan, Italy, 2013).
184. Knijnenburg, B.P., Kobsa, A. and Jin, H. 2013. Dimensionality of information disclosure behavior. International Journal of Human-Computer Studies. 71, 12 (2013), 1144–1162. DOI= http://dx.doi.org/10.1016/j.ijhcs.2013.06.003.
185. Knijnenburg, B.P., Kobsa, A. and Jin, H. 2013. Preference-based location sharing: are more privacy options really better? Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Paris, France, 2013), 2667–2676. DOI= http://dx.doi.org/10.1145/2470654.2481369.
186. Knijnenburg, B.P., Kobsa, A. and Jin, H. 2014. Segmenting the Recipients of Personal Information.
187. Knijnenburg, B.P., Reijmer, N.J.M. and Willemsen, M.C. 2011. Each to his own: how different users call for different interaction methods in recommender systems. Proceedings of the fifth ACM conference on Recommender systems (Chicago, IL, 2011), 141–148. DOI= http://dx.doi.org/10.1145/2043932.2043960.

188. Knijnenburg, B.P., Sivakumar, S. and Wilkinson, D. 2016. Recommender Systems for Self-Actualization. Proceedings of the 10th ACM Conference on Recommender Systems (New York, NY, USA, 2016), 11–14. DOI= http://dx.doi.org/10.1145/2959100.2959189.
189. Knijnenburg, B.P. and Willemsen, M.C. 2009. Understanding the effect of adaptive preference elicitation methods on user satisfaction of a recommender system. Proceedings of the third ACM conference on Recommender systems (New York, NY, 2009), 381–384. DOI= http://dx.doi.org/10.1145/1639714.1639793.
190. Knijnenburg, B.P., Willemsen, M.C. and Broeders, R. 2014. Smart Sustainability through System Satisfaction: Tailored Preference Elicitation for Energy-saving Recommenders. AMCIS 2014 proceedings (Savannah, GA, 2014).
191. Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H. and Newell, C. 2012. Explaining the user experience of recommender systems. User Modeling and User-Adapted Interaction. 22, 4–5 (2012), 441–504. DOI= http://dx.doi.org/10.1007/s11257-011-9118-4.
192. Knijnenburg, B.P., Willemsen, M.C. and Hirtbach, S. 2010. Receiving Recommendations and Providing Feedback: The User-Experience of a Recommender System. E-Commerce and Web Technologies. F. Buccafurri and G. Semeraro, eds. Springer. 207–216.
193. Kobsa, A. 2007. Privacy-Enhanced Personalization. Communications of the ACM. 50, 8 (2007), 24–33.
194. Kobsa, A., Cho, H. and Knijnenburg, B.P. 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. (Feb. 2016). DOI= http://dx.doi.org/10.1002/asi.23629.
195. Kobsa, A., Knijnenburg, B.P. and Livshits, B. 2014. Let's Do It at My Place Instead?: Attitudinal and Behavioral Study of Privacy in Client-side Personalization. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Toronto, Canada, 2014), 81–90. DOI= http://dx.doi.org/10.1145/2556288.2557102.
196. Kobsa, A. and Schreck, J. 2003. Privacy through pseudonymity in user-adaptive systems. ACM Transactions on Internet Technology. 3, 2 (May 2003), 149–183. DOI= http://dx.doi.org/10.1145/767193.767196.
197. Kobsa, A. and Teltzrow, M. 2005. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. Privacy Enhancing Technologies: Revised Selected Papers of the 4th International Workshop, PET 2004, Toronto, Canada, May 26-28, 2004. D. Martin and A. Serjantov, eds. Springer Berlin Heidelberg. 329–343.
198. Kolter, J. and Pernul, G. 2009. Generating User-Understandable Privacy Preferences. Conf. on Availability, Reliability and Security (Fukuoka, Japan, 2009), 299–306. DOI= http://dx.doi.org/10.1109/ARES.2009.89.
199. Koshimizu, T., Toriyama, T. and Babaguchi, N. 2006. Factors on the sense of privacy in video surveillance. Proceedings of the 3rd ACM workshop on Continuous archival and retrival of personal experences (New York, NY, USA, 2006), 35–44. DOI= http://dx.doi.org/10.1145/1178657.1178665.
200. Kosinski, M., Stillwell, D. and Graepel, T. 2013. Private traits and attributes are predictable from digital records of human behavior. Proceedings of the National Academy of Sciences. 110, 15 (Apr. 2013), 5802–5805. DOI= http://dx.doi.org/10.1073/pnas.1218772110.
201. Koutropoulos, A., Gallagher, M.S., Abajian, S.C., de Waard, I., Hogue, R.J., Keskin, N.O. and Rodriguez, C.O. 2011. Emotive Vocabulary in MOOCs: Context & Participant Retention. European Journal of Open, Distance and E-Learning. 1, (2011).
202. Krishnamurthy, B. and Wills, C. 2009. Privacy diffusion on the web: a longitudinal perspective. Proceedings of the 18th international conference on World Wide Web (New York, NY, USA, 2009), 541–550. DOI= http://dx.doi.org/10.1145/1526709.1526782.
203. Kushlev, K., Proulx, J. and Dunn, E.W. 2016. "Silence Your Phones": Smartphone Notifications Increase Inattention and Hyperactivity Symptoms. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (New York, NY, USA, 2016), 1011–1020. DOI= http://dx.doi.org/10.1145/2858036.2858359.
204. Lai, Y.-L. and Hui, K.-L. 2006. Internet Opt-In and Opt-Out: Investigating the Roles of Frames, Defaults and Privacy Concerns. Proceedings of the 2006 ACM SIGMIS Conference on Computer Personnel Research (Claremont, CA, 2006), 253–263. DOI= http://dx.doi.org/10.1145/1125170.1125230.
205. Lai, Y.-L. and Hui, K.-L. 2004. Opting-in or opting-out on the Internet: Does it Really Matter? ICIS 2004: Twenty-Fifth International Conference on Information Systems (Washington, D.C., 2004), 781–792.

206. Lampinen, A., Lehtinen, V., Lehmuskallio, A. and Tamminen, S. 2011. We're in it together: interpersonal management of disclosure in social network services. Proceedings of the SIGCHI conference on human factors in computing systems (2011), 3217–3226.
207. Langheinrich, M. 2001. Privacy by Design: Principles of Privacy-Aware Ubiquitous Systems. Ubicomp 2001. G.D. Abowd, B. Brumitt, and S.A.N. Shafer, eds. Springer-Verlag. 273–291.
208. Lanier, J. 2010. You Are Not a Gadget: A Manifesto. Thorndike Press.
209. Lao, E. and Kobsa, A. 2005. Privacy Attitudes of Internet Users in the U.S. and Europe. Technical Report #Internal Report. University of California, Irvine.
210. Larose, R. and Rifon, N.J. 2007. Promoting i-Safety: Effects of Privacy Warnings and Privacy Seals on Risk Assessment and Online Privacy Behavior. Journal of Consumer Affairs. 41, 1 (Jun. 2007), 127–149. DOI= http://dx.doi.org/10.1111/j.1745-6606.2006.00071.x.
211. Laufer, R.S., Proshansky, H.M. and Wolfe, M. 1973. Some Analytic Dimensions of Privacy. Proceedings of the Lund Conference on Architectural Psychology (Lund, Sweden, 1973).
212. Laufer, R.S. and Wolfe, M. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. Journal of Social Issues. 33, 3 (1977), 22–42. DOI= http://dx.doi.org/10.1111/j.1540-4560.1977.tb01880.x.
213. Lederer, S., Hong, J.I., Dey, A.K. and Landay, J.A. 2004. Personal privacy through understanding and action: five pitfalls for designers. Personal and Ubiquitous Computing. 8, 6 (Nov. 2004), 440–454. DOI= http://dx.doi.org/10.1007/s00779-004-0304-9.
214. Lederer, S., Mankoff, J. and Dey, A.K. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Ft. Lauderdale, FL, 2003), 724–725. DOI= http://dx.doi.org/10.1145/765891.765952.
215. Lee, V. 2013. The Quantified Self (QS) Movement and Some Emerging Opportunities for the Educational Technology Field. Educational Technology. November-December 2013 (Oct. 2013), 39–42.
216. Leonardi, P.M. 2014. Social Media, Knowledge Sharing, and Innovation: Toward a Theory of Communication Visibility. Information Systems Research. 25, 4 (Oct. 2014), 796–816. DOI= http://dx.doi.org/10.1287/isre.2014.0536.
217. Leonardi, P.M., Huysman, M. and Steinfield, C. 2013. Enterprise Social Media: Definition, History, and Prospects for the Study of Social Technologies in Organizations. Journal of Computer-Mediated Communication. 19, 1 (Oct. 2013), 1–19. DOI= http://dx.doi.org/10.1111/jcc4.12029.
218. Leventhal, J.C., Cummins, J.A., Schwartz, P.H., Martin, D.K. and Tierney, W.M. 2015. Designing a System for Patients Controlling Providers' Access to their Electronic Health Records: Organizational and Technical Challenges. Journal of General Internal Medicine. 30, 1 (Jan. 2015), 17–24. DOI= http://dx.doi.org/10.1007/s11606-014-3055-y.
219. Li, H., Sarathy, R. and Xu, H. 2011. The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. Decision Support Systems. 51, 3 (Jun. 2011), 434–445. DOI= http://dx.doi.org/10.1016/j.dss.2011.01.017.
220. Li, H., Sarathy, R. and Xu, H. 2010. Understanding situational online information disclosure as a privacy calculus. Journal of Computer Information Systems. 51, 1 (2010), 62–71.
221. Li, X. and Santhanam, R. 2008. Will it be Disclosure or Fabrication of Personal Information? An Examination of Persuasion Strategies on Prospective Employees. International Journal of Information Security and Privacy. 2, 4 (34 2008), 91–109. DOI= http://dx.doi.org/10.4018/jisp.2008100105.
222. Li, Y. 2014. The impact of disposition to privacy, website reputation and website familiarity on information privacy concerns. Decision Support Systems. 57, (Jan. 2014), 343–354. DOI= http://dx.doi.org/10.1016/j.dss.2013.09.018.
223. Li, Y. 2012. Theories in online information privacy research: A critical review and an integrated framework. Decision Support Systems. 54, 1 (Dec. 2012), 471–481. DOI= http://dx.doi.org/10.1016/j.dss.2012.06.010.
224. Linden, G., Smith, B. and York, J. 2003. Amazon.com recommendations: item-to-item collaborative filtering. IEEE Internet Computing. 7, 1 (Jan. 2003), 76–80. DOI= http://dx.doi.org/10.1109/MIC.2003.1167344.
225. Liu, Y., Gummadi, K.P., Krishnamurthy, B. and Mislove, A. 2011. Analyzing Facebook privacy settings: user expectations vs. reality. Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference (Berlin, Germany, 2011), 61–70. DOI= http://dx.doi.org/10.1145/2068816.2068823.

226. Lord, K.R., Lee, M.-S. and Sauer, P.L. 1995. The combined influence hypothesis: Central and peripheral antecedents of attitude toward the ad. Journal of Advertising. 24, 1 (1995), 73–85. DOI= http://dx.doi.org/10.1093/comjnl/bxs103.
227. Lowry, P.B., Moody, G., Vance, A., Jensen, M., Jenkins, J. and Wells, T. 2012. Using an elaboration likelihood approach to better understand the persuasiveness of website privacy assurance cues for online consumers. Journal of the American Society for Information Science and Technology. 63, 4 (Apr. 2012), 755–776. DOI= http://dx.doi.org/10.1002/asi.21705.
228. Lusoli, W., Bacigalupo, M., Lupiáñez-Villanueva, F., Andrade, N., Monteleone, S. and Maghiros, I. 2012. Pan-European Survey of Practices, Attitudes and Policy Preferences as Regards Personal Identity Data Management. Technical Report #ID 2086579. Social Science Research Network.
229. Lustig, C., Pine, K., Nardi, B., Irani, L., Lee, M.K., Nafus, D. and Sandvig, C. 2016. Algorithmic Authority: The Ethics, Politics, and Economics of Algorithms That Interpret, Decide, and Manage. Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (New York, NY, USA, 2016), 1057–1062. DOI= http://dx.doi.org/10.1145/2851581.2886426.
230. Machanavajjhala, A., Korolova, A. and Sarma, A.D. 2011. Personalized Social Recommendations: Accurate or Private. Proceedings of the VLDB Endowment. 4, 7 (Apr. 2011), 440–450. DOI= http://dx.doi.org/10.14778/1988776.1988780.
231. Madejski, M., Johnson, M. and Bellovin, S.M. 2012. A study of privacy settings errors in an online social network. Fourth International Workshop on SECurity and SOCial Networking (Lugano, Switzerland, 2012), 340–345. DOI= http://dx.doi.org/10.1109/PerComW.2012.6197507.
232. Malhotra, N.K., Kim, S.S. and Agarwal, J. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Nomological Framework. Information Systems Research. 15, 4 (2004), 336–355. DOI= http://dx.doi.org/10.1287/isre.1040.0032.
233. Marwick, A.E. and Boyd, D. 2014. Networked privacy: How teenagers negotiate context in social media. New Media & Society. (Jul. 2014), 1461444814543995. DOI= http://dx.doi.org/10.1177/1461444814543995.
234. Masiello, B. and Whitten, A. 2010. Engineering Privacy in an Age of Information Abundance. AAAI Spring Symposium: Intelligent Information Privacy Management (2010).
235. McAllister, M.P. 1992. Comic Books and AIDS. The Journal of Popular Culture. 26, 2 (Sep. 1992), 1–24. DOI= http://dx.doi.org/10.1111/j.0022-3840.1992.26021.x.
236. McCorry, D.G. 2014. With Cloud Technology, Who Owns Your Data. Federal Courts Law Review. 8, (2015 2014), 125–146.
237. McCune, J.C. 1998. Data, data, everywhere. Management Review. 87, 10 (1998), 10.
238. McDonald, A., Reeder, R., Kelley, P. and Cranor, L. 2009. A Comparative Study of Online Privacy Policies and Formats. Privacy Enhancing Technologies. Springer Berlin / Heidelberg. 37-55–55.
239. McGinty, L. and Smyth, B. 2006. Adaptive Selection: An Analysis of Critiquing and Preference-Based Feedback in Conversational Recommender Systems. International Journal of Electronic Commerce. 11, 2 (Dec. 2006), 35–57.
240. McKenzie, C.R.M., Liersch, M.J. and Finkelstein, S.R. 2006. Recommendations Implicit in Policy Defaults. Psychological Science. 17, 5 (May 2006), 414–420. DOI= http://dx.doi.org/10.1111/j.1467-9280.2006.01721.x.
241. McSherry, F. and Mironov, I. 2009. Differentially Private Recommender Systems: Building Privacy into the Net. Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (New York, NY, USA, 2009), 627–636. DOI= http://dx.doi.org/10.1145/1557019.1557090.
242. Metzger, M.J. 2007. Communication Privacy Management in Electronic Commerce. Journal of Computer-Mediated Communication. 12, 2 (2007), 335–361. DOI= http://dx.doi.org/10.1111/j.1083-6101.2007.00328.x.
243. Metzger, M.J. 2006. Effects of Site, Vendor, and Consumer Characteristics on Web Site Trust and Disclosure. Communication Research. 33, 3 (2006), 155–179.
244. Milne, G.R. and Culnan, M.J. 2004. Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don't Read) Online Privacy Notices. Journal of Interactive Marketing. 18, 3 (2004), 15–29.
245. Milne, G.R., Culnan, M.J. and Greene, H. 2006. A Longitudinal Assessment of Online Privacy Notice Readability. Journal of Public Policy & Marketing. 25, 2 (2006), 238–249.

246. Milne, G.R. and Gordon, M.E. 1993. Direct Mail Privacy-Efficiency Trade-offs within an Implied Social Contract Framework. Journal of Public Policy & Marketing. 12, 2 (Oct. 1993), 206–215. DOI= http://dx.doi.org/10.2307/30000091.
247. Moi, S.H., Rahim, N.B.A., Saad, P., Sim, P.L., Zakaria, Z. and Ibrahim, S. 2009. Iris Biometric Cryptography for Identity Document. 2009 International Conference of Soft Computing and Pattern Recognition (Dec. 2009), 736–741. DOI= http://dx.doi.org/10.1109/SoCPaR.2009.149.
248. Morik, K. 1989. User Models and Conversational Settings: Modeling the User's Wants. User Models in Dialog Systems. A.K. Wahlster and W, eds. 364–385.
249. Mosa, A.S.M., Yoo, I. and Sheets, L. 2012. A Systematic Review of Healthcare Applications for Smartphones. BMC Medical Informatics and Decision Making. 12, (2012), 67. DOI= http://dx.doi.org/10.1186/1472-6947-12-67.
250. Motti, V.G. and Caine, K. 2015. Micro interactions and Multi dimensional Graphical User Interfaces in the Design of Wrist Worn Wearables. Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Sep. 2015), 1712–1716. DOI= http://dx.doi.org/10.1177/1541931215591370.
251. Mulligan, D. and King, J. 2012. Bridging the Gap Between Privacy and Design. University of Pennsylvania Journal of Constitutional Law. 14, 4 (Mar. 2012), 989.
252. Mulligan, D. and Schwartz, A. 2000. Your Place or Mine?: Privacy Concerns and Solutions for Server and Client-Side Storage of Personal Information. Tenth conference on Computers, Freedom and Privacy (Toronto, Ontario, 2000), 81–84.
253. Narayanan, A. and Shmatikov, V. 2009. De-anonymizing Social Networks. 2009 30th IEEE Symposium on Security and Privacy (May 2009), 173–187. DOI= http://dx.doi.org/10.1109/SP.2009.22.
254. Narayanan, A. and Shmatikov, V. 2008. Robust De-anonymization of Large Sparse Datasets. 2008 IEEE Symposium on Security and Privacy (May 2008), 111–125. DOI= http://dx.doi.org/10.1109/SP.2008.33.
255. Nebeker, D.M. and Tatum, B.C. 1993. The Effects of Computer Monitoring, Standards, and Rewards on Work Performance, Job Satisfaction, and Stress1. Journal of Applied Social Psychology. 23, 7 (Apr. 1993), 508–536. DOI= http://dx.doi.org/10.1111/j.1559-1816.1993.tb01101.x.
256. Netflix Spilled Your Brokeback Mountain Secret, Lawsuit Claims: 2009. http://www.wired.com/2009/12/netflix-privacy-lawsuit/?+wired27b+%2528Blog+-+27B+Stroke+6+%2528Threat+Level%2529%2529/. Accessed: 2015-11-17.
257. Netter, M., Weber, M., Diener, M. and Pernul, G. 2014. Visualizing social roles - design and evaluation of a bird's-eye view of social network privacy settings. ECIS 2014 Proceedings (Jun. 2014).
258. Nissenbaum, H. 2011. A Contextual Approach to Privacy Online. Daedalus. 140, 4 (Oct. 2011), 32–48. DOI= http://dx.doi.org/10.1162/DAED_a_00113.
259. Nissenbaum, H. 2004. Privacy as Contextual Integrity. Washington Law Review. 79, (2004), 119–157.
260. Nissenbaum, H.F. 2009. Privacy in context : technology, policy, and the integrity of social life. Stanford Law Books.
261. Niu, Y., Shi, E., Chow, R., Golle, P. and Jakobsson, M. 2010. One Experience Collecting Sensitive Mobile Data. SOUPS 2010 Usable Security Experiment Reports (USER) Workshop (2010).
262. Norberg, P.A., Horne, D.R. and Horne, D.A. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. Journal of Consumer Affairs. 41, 1 (2007), 100–126. DOI= http://dx.doi.org/10.1111/j.1745-6606.2006.00070.x.
263. Notifications on your Apple Watch: 2016. https://support.apple.com/en-gb/HT204791. Accessed: 2017-02-05.
264. Ohm, P. 2010. Broken promises of privacy: Responding to the surprising failure of anonymization. UCLA law review. 57, (2010), 1701.
265. Olivero, N. and Lunt, P. 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 (2004), 243–262. DOI= http://dx.doi.org/10.1016/S0167-4870(02)00172-1.
266. Olson, J.S., Grudin, J. and Horvitz, E. 2005. A study of preferences for sharing and privacy. CHI '05 Extended Abstracts (Portland, OR, 2005), 1985–1988. DOI= http://dx.doi.org/10.1145/1056808.1057073.
267. Orwell, G. 1949. Nineteen Eighty-Four. A novel. Secker & Warburg.

268. Osch, W. v, Steinfield, C.W. and Balogh, B.A. 2015. Enterprise Social Media: Challenges and Opportunities for Organizational Communication and Collaboration. 2015 48th Hawaii International Conference on System Sciences (Jan. 2015), 763–772. DOI= http://dx.doi.org/10.1109/HICSS.2015.97.
269. van Osch, W., Bulgurcu, B. and Kane, G. 2016. Classifying Enterprise Social Media Users: A Mixed-Method Study of Organizational Social Media Use. ICIS 2016 Proceedings (Dec. 2016).
270. Page, X., Knijnenburg, B.P. and Kobsa, A. 2013. FYI: communication style preferences underlie differences in location-sharing adoption and usage. Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing (New York, NY, USA, 2013), 153–162. DOI= http://dx.doi.org/10.1145/2493432.2493487.
271. Page, X., Knijnenburg, B.P. and Kobsa, A. 2013. What a Tangled Web We Weave: Lying Backfires in Location-sharing Social Media. Proceedings of the 2013 Conference on Computer Supported Cooperative Work (San Antonio, TX, 2013), 273–284. DOI= http://dx.doi.org/10.1145/2441776.2441808.
272. Page, X., Kobsa, A. and Knijnenburg, B.P. 2012. Don’t Disturb My Circles! Boundary Preservation Is at the Center of Location-Sharing Concerns. Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media (Dublin, Ireland, May 2012), 266–273.
273. Page, X.W. 2014. Factors that Influence Adoption and Use of Location-Sharing Social Media. University of California, Irvine.
274. Pan, Y. and Zinkhan, G.M. 2006. Exploring the impact of online privacy disclosures on consumer trust. Journal of Retailing. 82, 4 (2006), 331–338. DOI= http://dx.doi.org/10.1016/j.jretai.2006.08.006.
275. Pariser, E. 2012. The filter bubble: how the new personalized Web is changing what we read and how we think. Penguin Books.
276. Patil, S. and Kobsa, A. 2005. Uncovering Privacy Attitudes in Instant Messaging. Proceedings of the 5th ACM Conference on Supporting Group Work (Sanibel Island, FL, 2005), 101–104.
277. Patil, S. and Lai, J. 2005. Who Gets to Know What when: Configuring Privacy Permissions in an Awareness Application. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Portland, OR, 2005), 101–110. DOI= http://dx.doi.org/10.1145/1054972.1054987.
278. Patil, S., Page, X. and Kobsa, A. 2011. With a little help from my friends: can social navigation inform interpersonal privacy preferences? Proceedings of the ACM 2011 conference on Computer supported cooperative work (Hangzhou, China, Mar. 2011), 391–394. DOI= http://dx.doi.org/10.1145/1958824.1958885.
279. Pavlou, P.A. 2003. Consumer Acceptance of Electronic Commerce: Integrating Trust and Risk with the Technology Acceptance Model. International Journal of Electronic Commerce. 7, 3 (2003), 101–134.
280. Pavlou, P.A. 2011. State of the Information Privacy Literature: Where Are We Now and Where Should We Go. MIS Quarterly. 35, 4 (2011), 977–988.
281. Petronio, S. 2002. Boundaries of Privacy: Dialectics of Disclosure. State University of New York Press.
282. Petronio, S. 1991. Communication Boundary Management: A Theoretical Model of Managing Disclosure of Private Information Between Marital Couples. Communication Theory. 1, 4 (Nov. 1991), 311–335. DOI= http://dx.doi.org/10.1111/j.1468-2885.1991.tb00023.x.
283. Petronio, S. 2010. Communication Privacy Management Theory: What Do We Know About Family Privacy Regulation? Journal of Family Theory & Review. 2, 3 (Sep. 2010), 175–196. DOI= http://dx.doi.org/10.1111/j.1756-2589.2010.00052.x.
284. Petry, N.M. and O’Brien, C.P. 2013. Internet gaming disorder and the DSM-5. Addiction. 108, 7 (Jul. 2013), 1186–1187. DOI= http://dx.doi.org/10.1111/add.12162.
285. Petty, R.E. and Cacioppo, J.T. 1986. The Elaboration Likelihood Model of Persuasion. Advances in Experimental Social Psychology. Leonard Berkowitz, ed. Academic Press. 123–205.
286. Petty, R.E., Cacioppo, J.T. and Schumann, D. 1983. Central and peripheral routes to advertising effectiveness: The moderating role of involvement. Journal of consumer research. (1983), 135–146.
287. Petty, R.E. and Wegener, D.T. 1999. The elaboration likelihood model: Current status and controversies. Dual-process theories in social psychology. S. Chaiken and Y. Trope, eds. Guilford Press. 37–72.
288. Phelan, C., Lampe, C. and Resnick, P. 2016. It’s Creepy, But It Doesn’T Bother Me. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (New York, NY, USA, 2016), 5240–5251. DOI= http://dx.doi.org/10.1145/2858036.2858381.

289. Phelps, J., Nowak, G. and Ferrell, E. 2000. Privacy Concerns and Consumer Willingness to Provide Personal Information. Journal of Public Policy & Marketing. 19, 1 (Mar. 2000), 27–41. DOI= http://dx.doi.org/10.1509/jppm.19.1.27.16941.
290. Pielot, M., Church, K. and de Oliveira, R. 2014. An In-situ Study of Mobile Phone Notifications. Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (New York, NY, USA, 2014), 233–242. DOI= http://dx.doi.org/10.1145/2628363.2628364.
291. Pollach, I. 2007. What's Wrong with Online Privacy Policies? Communications of the ACM. 50, 9 (Sep. 2007), 103–108. DOI= http://dx.doi.org/10.1145/1284621.1284627.
292. Poushter, J. 2016. Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies. Pew Research Center.
293. Pu, P. and Chen, L. 2006. Trust building with explanation interfaces. Proceedings of the 11th international conference on Intelligent user interfaces (2006), 93–100.
294. Pu, P., Faltings, B., Chen, L., Zhang, J. and Viappiani, P. 2011. Usability Guidelines for Product Recommenders Based on Example Critiquing Research. Recommender Systems Handbook. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer US. 511–545.
295. Raber, F., Luca, A.D. and Graus, M. 2016. Privacy Wedges: Area-Based Audience Selection for Social Network Posts. Twelfth Symposium on Usable Privacy and Security (SOUPS 2016) (Denver, CO, Jun. 2016).
296. Raybourn, E.M., Fabian, N., Davis, W., Parks, R.C., McClain, J., Trumbo, D., Regan, D. and Durlach, P. 2015. Data Privacy and Security Considerations for Personal Assistants for Learning (PAL). Proceedings of the 20th International Conference on Intelligent User Interfaces Companion (2015), 69–72. DOI= http://dx.doi.org/10.1145/2732158.2732195.
297. Reid, J.M. 1995. Learning Styles in the ESL/EFL Classroom. Heinle & Heinle Publishers.
298. Reiman, J.H. 1995. Driving to the Panopticon: A Philosophical Exploration of the Risks to Privacy Posed by the Highway Technology of the Future. Santa Clara Computer and High-Technology Law Journal. 11, (1995), 27.
299. Rendle, S. and Schmidt-Thieme, L. 2008. Online-updating Regularized Kernel Matrix Factorization Models for Large-scale Recommender Systems. Proceedings of the 2008 ACM Conference on Recommender Systems (New York, NY, USA, 2008), 251–258. DOI= http://dx.doi.org/10.1145/1454008.1454047.
300. Riboni, D. and Bettini, C. 2012. Private context-aware recommendation of points of interest: An initial investigation. Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops (Los Alamitos, CA, USA, 2012), 584–589. DOI= http://dx.doi.org/10.1109/PerComW.2012.6197582.
301. Rifon, N.J., LaRose, R. and Choi, S.M. 2005. Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures. Journal of Consumer Affairs. 39, 2 (2005), 339–360. DOI= http://dx.doi.org/10.1111/j.1745-6606.2005.00018.x.
302. Rose, E. 2005. Data Users versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information? Proceedings of the 38th Annual Hawaii International Conference on System Sciences, 2005. HICSS '05 (2005), 180c–180c. DOI= http://dx.doi.org/10.1109/HICSS.2005.184.
303. Rowe, G. and Wright, G. 1999. The Delphi technique as a forecasting tool: issues and analysis. International Journal of Forecasting. 15, 4 (Oct. 1999), 353–375. DOI= http://dx.doi.org/10.1016/S0169-2070(99)00018-7.
304. Rust, R.T., Kannan, P.K. and Peng, N. 2002. The Customer Economics of Internet Privacy. Journal of the Academy of Marketing Science. 30, 4 (Oct. 2002), 455–464. DOI= http://dx.doi.org/10.1177/009207002236917.
305. Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M. and Rao, J. 2009. Understanding and capturing people's privacy policies in a mobile social networking application. Personal and Ubiquitous Computing. 13, 6 (2009), 401–412. DOI= http://dx.doi.org/10.1007/s00779-008-0214-3.
306. Sahami Shirazi, A., Henze, N., Dingler, T., Pielot, M., Weber, D. and Schmidt, A. 2014. Large-scale Assessment of Mobile Notifications. Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (New York, NY, USA, 2014), 3055–3064. DOI= http://dx.doi.org/10.1145/2556288.2557189.
307. Samuelson, W. and Zeckhauser, R. 1988. Status quo bias in decision making. Journal of Risk and Uncertainty. 1, 1 (Mar. 1988), 7–59. DOI= http://dx.doi.org/10.1007/BF00055564.
308. Schaar, P. 2010. Privacy by Design. Identity in the Information Society. 3, 2 (Aug. 2010), 267–274. DOI= http://dx.doi.org/10.1007/s12394-010-0055-x.

309. Schafer, J.B., Konstan, J.A. and Riedl, J. 2001. E-Commerce Recommendation Applications. Data Mining and Knowledge Discovery. 5, (2001), 115–153.
310. Schatz, S. 2016. The Total Learning Architecture: Rationale and Design. Advanced Distributed Learning Initiative.
311. Schnabel, T., Swaminathan, A., Singh, A., Chandak, N. and Joachims, T. 2016. Recommendations as Treatments: Debiasing Learning and Evaluation. (2016), 1670–1679.
312. Schneider, K.F., Lyle, D.S. and Murphy, F.X. 2015. Framing the Big Data Ethics Debate for the Military. National Defense University Press.
313. Shamdasani, P.N., Stanaland, A.J. and Tan, J. 2001. Location, location, location: Insights for advertising placement on the web. Journal of Advertising Research. 41, 4 (2001), 7–21.
314. Shapiro, S.S. 2009. Privacy by design: moving from art to practice. Commun. ACM. 53, (Jun. 2009), 27–29. DOI= http://dx.doi.org/10.1145/1743546.1743559.
315. Sheehan, K.B. and Hoy, M.G. 2000. Dimensions of Privacy Concern Among Online Consumers. Journal of Public Policy & Marketing. 19, 1 (2000), 62–73.
316. Shelton, M., Lo, K. and Nardi, B. 2015. Online Media Forums as Separate Social Lives: A Qualitative Study of Disclosure Within and Beyond Reddit. iConference 2015 Proceedings (Newport Beach, CA, Mar. 2015).
317. Sheng, H., Nah, F.F.-H. and Siau, K. 2008. An Experimental Study on Ubiquitous commerce Adoption: Impact of Personalization and Privacy Concerns. Journal of the Association for Information Systems. 9, 6 (Jun. 2008), 344–376.
318. Shklovski, I., Mainwaring, S.D., Skúladóttir, H.H. and Borgthorsson, H. 2014. Leakiness and Creepiness in App Space: Perceptions of Privacy and Mobile App Use. Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (New York, NY, USA, 2014), 2347–2356. DOI= http://dx.doi.org/10.1145/2556288.2557421.
319. Shokri, R. and Shmatikov, V. 2015. Privacy-Preserving Deep Learning. (Oct. 2015), 1310–1321. DOI= http://dx.doi.org/10.1145/2810103.2813687.
320. Simon, H.A. 1959. Theories of Decision-Making in Economics and Behavioral Science. The American Economic Review. 49, 3 (Jun. 1959), 253–283. DOI= http://dx.doi.org/10.2307/1809901.
321. Simonson, I. 1989. Choice Based on Reasons: The Case of Attraction and Compromise Effects. Journal of Consumer Research. 16, 2 (Sep. 1989), 158–174.
322. Singleton, S.M. and Harper, J. 2002. With A Grain of Salt: What Consumer Privacy Surveys Don't Tell Us. Technical Report #ID 299930. Social Science Research Network.
323. Sinha, R.R. and Swearingen, K. 2001. Comparing Recommendations Made by Online Systems and Friends. DELOS Workshop on Personalisation and Recommender Systems in Digital Libraries (2001), 64–67.
324. Slovic, P., Finucane, M.L., Peters, E. and MacGregor, D.G. 2004. Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. Risk analysis. 24, 2 (2004), 311–322. DOI= http://dx.doi.org/10.1111/j.0272-4332.2004.00433.x.
325. Smith, H.J., Dinev, T. and Xu, H. 2011. Information Privacy Research: An Interdisciplinary Review. MIS Quarterly. 35, 4 (Dec. 2011), 989–1016.
326. Smith, H.J., Milberg, S.J. and Burke, S.J. 1996. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. MIS Quarterly. 20, 2 (1996), 167–196. DOI= http://dx.doi.org/10.2307/249477.
327. Smith, N.C., Goldstein, D.G. and Johnson, E.J. 2013. Choice Without Awareness: Ethical and Policy Implications of Defaults. Journal of Public Policy & Marketing. 32, 2 (Fall 2013), 159–172. DOI= http://dx.doi.org/10.1509/jppm.10.114.
328. Smyth, B. 2007. Case-Based Recommendation. The Adaptive Web: Methods and Strategies of Web Personalization. P. Brusilovsky, A. Kobsa, and W. Nejdl, eds. Springer Verlag. 342–376.
329. SoarTech 2016. TLA Design Document.
330. Solove, D. 2008. Understanding Privacy. Harvard University Press.
331. Solove, D.J. 2006. A Taxonomy of Privacy. University of Pennsylvania Law Review. 154, 3 (2006), 477–564. DOI= http://dx.doi.org/10.2307/40041279.
332. Solove, D.J. 2013. Privacy Self-Management and the Consent Dilemma. Harvard Law Review. 126, (2013), 1880–1903.
333. Sparling, E.I. and Sen, S. 2011. Rating: How Difficult is It? Proceedings of the Fifth ACM Conference on Recommender Systems (Chicago, IL, 2011), 149–156. DOI= http://dx.doi.org/10.1145/2043932.2043961.

334. Spekman, M.L.C., Konijn, E.A., Roelofsma, P.H.M.P. and Griffiths, M.D. 2013. Gaming addiction, definition and measurement: A large-scale empirical study. Computers in Human Behavior. 29, 6 (Nov. 2013), 2150–2155. DOI= http://dx.doi.org/10.1016/j.chb.2013.05.015.
335. Spiekermann, S. 2012. The Challenges of Privacy by Design. Commun. ACM. 55, 7 (Jul. 2012), 38–40. DOI= http://dx.doi.org/10.1145/2209249.2209263.
336. Spiekermann, S., Grossklags, J. and Berendt, B. 2001. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. Proceedings of the 3rd ACM conference on Electronic Commerce (Tampa, FL, 2001), 38–47.
337. Stone, D.L. 1981. The effects of the valence of outcomes for providing data and the perceived relevance of the data requested on privacy-related behaviors, beliefs, and attitudes. Doctoral Thesis #Thesis 31634 PhD. Purdue University.
338. Stone, E.F. and Stone, D.L. 1990. Privacy in Organizations: Theoretical Issues, Research Findings, and Protection Mechanisms. Research in Personnel and Human Resources Management. 8, (1990), 349–411.
339. Strater, K. and Lipford, H.R. 2008. Strategies and struggles with privacy in an online social networking community. Proceedings of the 22nd British HCI Group Annual Conference on People and Computers (Swinton, UK, 2008), 111–119.
340. Suen, H.K. 2014. Peer assessment for massive open online courses (MOOCs). The International Review of Research in Open and Distributed Learning. 16, 3 (2014).
341. Suler, J. 2004. The Online Disinhibition Effect. CyberPsychology & Behavior. 7, 3 (Jun. 2004), 321–326. DOI= http://dx.doi.org/10.1089/1094931041291295.
342. Sutanto, J., Palme, E., Tan, C.-H. and Phang, C.W. 2013. Addressing the Personalization-Privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users. MIS Quarterly. 37, 4 (2013), 1141–1164.
343. Tang, K., Hong, J. and Siewiorek, D. 2012. The implications of offering more disclosure choices for social location sharing. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Austin, TX, 2012), 391–394. DOI= http://dx.doi.org/10.1145/2207676.2207730.
344. Tang, K., Lin, J., Hong, J., Siewiorek, D. and Sadeh, N. 2010. Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing. Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing (Copenhagen, Denmark, 2010), 85–94. DOI= http://dx.doi.org/10.1145/1864349.1864363.
345. Taylor, D., Davis, D. and Jillapalli, R. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. Electronic Commerce Research. 9, 3 (Sep. 2009), 203–223. DOI= http://dx.doi.org/10.1007/s10660-009-9036-2.
346. Taylor, S.E. and Lobel, M. 1989. Social comparison activity under threat: Downward evaluation and upward contacts. Psychological Review. 96, 4 (1989), 569–575. DOI= http://dx.doi.org/10.1037/0033-295X.96.4.569.
347. Teltzrow, M. and Kobsa, A. 2004. Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study. Designing Personalized User Experiences for eCommerce. C.-M. Karat, J. Blom, and J. Karat, eds. Kluwer Academic Publishers. 315–332.
348. Tene, O. and Polonetsky, J. 2013. A Theory of Creepy: Technology, Privacy and Shifting Social Norms. Yale Journal of Law and Technology. 16, (2013), 59–102.
349. Thaler, R.H. and Sunstein, C. 2008. Nudge : improving decisions about health, wealth, and happiness. Yale University Press.
350. The Center for Information Policy Leadership 2005. Multi-Layered Notices Explained. Center for Information Policy Leadership, Hunton & Williams.
351. The Social Web: Who owns your data? 2012. http://www.zdnet.com/article/the-social-web-who-owns-your-data/. Accessed: 2017-01-27.
352. Three-quarters of smartphone owners use location-based services: 2012. http://pewinternet.org/~/media//Files/Reports/2012/PIP_Location_based_services_2012_Report.pdf.
353. Tintarev, N. and Masthoff, J. 2007. A Survey of Explanations in Recommender Systems. Data Engineering Workshop (Istanbul, Turkey, Apr. 2007), 801–810. DOI= http://dx.doi.org/10.1109/ICDEW.2007.4401070.
354. Tintarev, N. and Masthoff, J. 2011. Designing and Evaluating Explanations for Recommender Systems. Recommender Systems Handbook. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer US. 479–510.

355. Tintarev, N. and Masthoff, J. 2012. Evaluating the effectiveness of explanations for recommender systems. User Modeling and User-Adapted Interaction. 22, 4–5 (Feb. 2012), 399–439. DOI= http://dx.doi.org/10.1007/s11257-011-9117-5.
356. Toch, E., Cranshaw, J., Drielsma, P.H., Tsai, J.Y., Kelley, P.G., Springfield, J., Cranor, L., Hong, J. and Sadeh, N. 2010. Empirical models of privacy in location sharing. Proceedings of the 12th ACM international conference on Ubiquitous computing (Copenhagen, Denmark, 2010), 129–138. DOI= http://dx.doi.org/10.1145/1864349.1864364.
357. Treiblmaier, H. and Pollach, I. 2007. Users' Perceptions of Benefits and Costs of Personalization. ICIS 2007 Proceedings (2007).
358. Tsai, J.Y., Egelman, S., Cranor, L. and Acquisti, A. 2010. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. Information Systems Research. (Feb. 2010). DOI= http://dx.doi.org/10.1287/isre.1090.0260.
359. Tsai, J.Y., Kelley, P., Drielsma, P., Cranor, L.F., Hong, J. and Sadeh, N. 2009. Who's viewed you?: the impact of feedback in a mobile location-sharing application. Proceedings of the 27th international conference on Human factors in computing systems (Boston, MA, USA, 2009), 2003–2012. DOI= http://dx.doi.org/10.1145/1518701.1519005.
360. Tufekci, Z. 2008. Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites. Bulletin of Science, Technology & Society. 28, 1 (Feb. 2008), 20–36. DOI= http://dx.doi.org/10.1177/0270467607311484.
361. Turner, L.D., Allen, S.M. and Whitaker, R.M. 2015. Push or Delay? Decomposing Smartphone Notification Response Behaviour. Human Behavior Understanding. A.A. Salah, B.J.A. Kröse, and D.J. Cook, eds. Springer International Publishing. 69–83. DOI= http://dx.doi.org/10.1007/978-3-319-24195-1_6.
362. Turner, M.A. and Varghese, R. 2002. Making sense of the privacy debate: a comparative analysis of leading consumer privacy surveys. Privacy & American Business.
363. Turow, J., Feldman, L. and Meltzer, K. 2005. Open to Exploitation: America's Shoppers Online and Offline. Annenberg Public Policy Center, University of Pennsylvania.
364. Tversky, A. and Kahneman, D. 1974. Judgment under Uncertainty: Heuristics and Biases. Science. 185, 4157 (Sep. 1974), 1124–1131. DOI= http://dx.doi.org/10.1126/science.185.4157.1124.
365. US Air Force 2013. Nuclear Suerty Tamper Control and Detection Programs. Technical Report #Air Force Instruction 91-104.
366. U.S. Consumers Want More Personalized Retail Experience and Control Over Personal Information, Accenture Survey Shows: 2015. http://newsroom.accenture.com/news/us-consumers-want-more-personalized-retail-experience-and-control-over-personal-information-accenture-survey-shows.htm. Accessed: 2015-03-17.
367. Vallet, D., Friedman, A. and Berkovsky, S. 2014. Matrix Factorization without User Data Retention. Advances in Knowledge Discovery and Data Mining. V.S. Tseng, T.B. Ho, Z.-H. Zhou, A.L.P. Chen, and H.-Y. Kao, eds. Springer International Publishing. 569–580.
368. Van Osch, W. and Coursaris, C. 2012. The Duality of Social Media: Structuration and Socialization through Organizational Communication. SIGHCI 2012 Proceedings (Dec. 2012).
369. Van Slyke, C., Shim, J.T., Johnson, R. and Jiang, J.J. 2006. Concern for Information Privacy and Online Consumer Purchasing. Journal of the Association for Information Systems. 7, 1 (Jun. 2006).
370. Viano, R. 2015. OSLM. ADL Net.
371. Vickery, J.R. 2014. The curious case of Confession Bear: the reappropriation of online macro-image memes. Information, Communication & Society. 17, 3 (Mar. 2014), 301–325. DOI= http://dx.doi.org/10.1080/1369118X.2013.871056.
372. Visser, S.L. The Soldier and Autonomy. Military Medical Ethics. DIANE Publishing. 251–266.
373. Wang, W. and Benbasat, I. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. Journal of Management Information Systems. 23, 4 (2007), 217–246. DOI= http://dx.doi.org/10.2753/MIS0742-1222230410.
374. Wang, Y. and Kobsa, A. 2013. A PLA-based privacy-enhancing user modeling framework and its evaluation. User Modeling and User-Adapted Interaction: The Journal of Personalization Research. 1, 23 (2013), 41–82. DOI= http://dx.doi.org/10.1007/s11257-011-9114-8.

375. Wang, Y., Leon, P.G., Acquisti, A., Cranor, L.F., Forget, A. and Sadeh, N. 2014. A Field Trial of Privacy Nudges for Facebook. Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems (Toronto, Canada, 2014), 2367–2376. DOI= http://dx.doi.org/10.1145/2556288.2557413.
376. Wang, Y., Leon, P.G., Scott, K., Chen, X., Acquisti, A. and Cranor, L.F. 2013. Privacy Nudges for Social Media: An Exploratory Facebook Study. Second International Workshop on Privacy and Security in Online Social Media (Rio De Janeiro, Brazil, 2013), 763–770.
377. Watson, J., Besmer, A. and Lipford, H.R. 2012. +Your circles: sharing behavior on Google+. Proceedings of the 8th Symposium on Usable Privacy and Security (Pittsburgh, PA, 2012), 12:1-12:10. DOI= http://dx.doi.org/10.1145/2335356.2335373.
378. Wenning, R. and Schunter, M. 2006. The Platform for Privacy Preferences 1.1 (P3P1.1) Specification. W3C Working Group Note.
379. Westin, A.F., Harris, L. and associates 1981. The Dimensions of privacy : a national opinion research survey of attitudes toward privacy. Garland Publishing.
380. Westin, A.F. and Maurici, D. 1998. E-Commerce & Privacy: What the Net Users Want. Privacy & American Business, and PricewaterhouseCoopers LLP.
381. White House 2012. Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Economy. White House.
382. White, M., Hill, S., McGovern, P., Mills, C. and Smeaton, D. 2003. "High-performance" Management Practices, Working Hours and Work–Life Balance. British Journal of Industrial Relations. 41, 2 (Jun. 2003), 175–195. DOI= http://dx.doi.org/10.1111/1467-8543.00268.
383. White, T.B. 2004. Consumer Disclosure and Disclosure Avoidance: A Motivational Framework. Journal of Consumer Psychology. 14, 1&2 (2004), 41–51.
384. Why big data evangelists need to be reprogrammed: 2014. http://www.zdnet.com/article/why-big-data-evangelists-need-to-be-reprogrammed/. Accessed: 2015-10-28.
385. Wilkinson, D., Sivakumar, S., Cherry, D., Knijnenburg, B.P., Raybourn, E.M., Wisniewski, P. and Sloan, H. 2017. User-Tailored Privacy by Design. Proceedings of the Usable Security Mini Conference 2017 (San Diego, CA, 2017). DOI= http://dx.doi.org/http://dx.doi.org/10.14722/usec.2017.23007.
386. Wilson, D. and Valacich, J. 2012. Unpacking the Privacy Paradox: Irrational Decision-Making within the Privacy Calculus. ICIS 2012 Proceedings (Orlando, FL, Dec. 2012).
387. Wisniewski, P., Islam, A.K.M., Lipford, H.R. and Wilson, D. 2016. Framing and Measuring Multi-dimensional Interpersonal Privacy Preferences of Social Networking Site Users. Communications of the Association for Information Systems. 38, 1 (Jan. 2016).
388. Wisniewski, P., Islam, A.K.M.N., Knijnenburg, B.P. and Patil, S. 2015. Give Social Network Users the Privacy They Want. Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (Vancouver, Canada, 2015), 1427–1441. DOI= http://dx.doi.org/10.1145/2675133.2675256.
389. Wisniewski, P., Knijnenburg, B.P. and Richter Lipford, H. 2014. Profiling Facebook Users' Privacy Behaviors. SOUPS2014 Workshop on Privacy Personas and Segmentation (Menlo Park, CA, 2014).
390. Wisniewski, P., Lipford, H. and Wilson, D. 2012. Fighting for my space: coping mechanisms for SNS boundary regulation. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (New York, NY, USA, 2012), 609–618. DOI= http://dx.doi.org/10.1145/2207676.2207761.
391. Wisniewski, P.J., Knijnenburg, B.P. and Lipford, H.R. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. International Journal of Human-Computer Studies. 98, (Feb. 2017), 95–108. DOI= http://dx.doi.org/10.1016/j.ijhcs.2016.09.006.
392. Xie, J., Knijnenburg, B.P. and Jin, H. 2014. Location Sharing Privacy Preference: Analysis and Personalized Recommendation. Proceedings of the 19th International Conference on Intelligent User Interfaces (New York, NY, USA, 2014), 189–198. DOI= http://dx.doi.org/10.1145/2557500.2557504.
393. Xu, H. 2007. The effects of self-construal and perceived control on privacy concerns. ICIS 2007 Proceedings (2007), paper 125.
394. Xu, H., Luo, X. (Robert), Carroll, J.M. and Rosson, M.B. 2011. The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. Decision Support Systems. 51, 1 (Apr. 2011), 42–52. DOI= http://dx.doi.org/10.1016/j.dss.2010.11.017.

395. Xu, H., Teo, H.-H. and Tan, B.C.Y. 2005. Predicting the Adoption of Location-Based Services: The Role of Trust and Perceived Privacy Risk. Proceedings of the International Conference on Information Systems (Las Vegas, NV, Dec. 2005), 861–874.
396. Xu, H., Teo, H.-H., Tan, B.C.Y. and Agarwal, R. 2009. The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services. Journal of Management Information Systems. 26, 3 (Dec. 2009), 135–174. DOI= http://dx.doi.org/10.2753/MIS0742-1222260305.
397. Xu, H., Wang, N. and Grossklags, J. 2012. Privacy-by-ReDesign: Alleviating Privacy Concerns for Third-Party Applications. ICIS 2012 Proceedings (Orlando, FL, 2012).
398. Yang, S.-C., Hung, W.-C., Sung, K. and Farn, C.-K. 2006. Investigating initial trust toward e-tailers from the elaboration likelihood model perspective. Psychology and Marketing. 23, 5 (May 2006), 429–445. DOI= http://dx.doi.org/10.1002/mar.20120.
399. Yang, S.J.H. 2006. Context Aware Ubiquitous Learning Environments for Peer-to-Peer Collaborative Learning. Educational Technology & Society. 2006, (2006), 188–201.
400. Yar, M. 2013. Cybercrime and Society. SAGE Publications Ltd.
401. You Don’t Own Your Data: 2014. http://lifehacker.com/you-dont-own-your-data-1556088120. Accessed: 2017-01-27.
402. Zaslow, J. 2002. If TiVo Thinks You Are Gay, Here’s How to Set It Straight. Wall Street Journal (Eastern Edition).
403. Zhang, Y. 1996. Responses to humorous advertising: The moderating effect of need for cognition. Journal of Advertising. 25, 1 (1996), 15–32.
404. Zheleva, E. and Getoor, L. 2009. To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. Proceedings of the 18th international conference on World wide web (New York, NY, USA, 2009), 531–540. DOI= http://dx.doi.org/10.1145/1526709.1526781.
405. Zhou, B. and Pei, J. 2008. Preserving Privacy in Social Networks Against Neighborhood Attacks. 2008 IEEE 24th International Conference on Data Engineering (Apr. 2008), 506–515. DOI= http://dx.doi.org/10.1109/ICDE.2008.4497459.
406. Zhou, T. 2012. Examining Location-based Services Usage from the Perspectives of Unified Theory of Acceptance and Use of Technology and Privacy Risk. Journal of Electronic Commerce Research. 13, 2 (May 2012), 135–144.
407. Zhou, T. 2012. Understanding users’ initial trust in mobile banking: An elaboration likelihood perspective. Computers in Human Behavior. 28, 4 (Jul. 2012), 1518–1525. DOI= http://dx.doi.org/10.1016/j.chb.2012.03.021.
408. Zhu, T., Ren, Y., Zhou, W., Rong, J. and Xiong, P. 2014. An effective privacy preserving algorithm for neighborhood-based collaborative filtering. Future Generation Computer Systems. 36, (Jul. 2014), 142–155. DOI= http://dx.doi.org/10.1016/j.future.2013.07.019.
409. Facebook real-name policy controversy. Wikipedia, the free encyclopedia.