International Journal of Mobile Human Computer Interaction
Volume 11 • Issue 4 • October-December 2019

Eliciting Design Guidelines for Privacy Notifications in mHealth Environments

Patrick Murmann, Karlstad University, Karlstad, Sweden

ABSTRACT

The possibilities of employing mobile health (mhealth) devices for the purpose of self-quantification and fitness tracking are increasing; yet few users of online mhealth services possess proven knowledge of how their personal data are processed once the data have been disclosed. Ex post transparency-enhancing tools (TETs) can provide such insight and guide users in making informed decisions with respect to intervening with the processing of their personal data. At present, however, there are no suitable guidelines that aid designers of TETs in implementing privacy notifications that reflect their recipients' needs in terms of what they want to be notified about and the level of guidance required to audit their data effectively. Based on an analysis of gaps related to TETs, the findings of a study on privacy notification preferences, and the findings on notifications and privacy notices discussed in the literature, this paper proposes a set of guidelines for the human-centred design of privacy notifications that facilitate ex post transparency.

KEYWORDS

Data Transparency, Human-Centred Design, Individualisation, Intervenability, Mobile Health (mHealth), Notification, Privacy, Transparency-Enhancing Tool (TET), Usability

1. INTRODUCTION

The number of users of mobile health (mhealth) devices is increasing (Statista, 2018), as is the spectrum of applications related to personal informatics (Knowles et al., 2018). However, few users of online services know how their personal data are processed by the data services they are relying on (Lau et al., 2018). This imbalance of knowledge, and hence power, between service providers and users is in stark contrast to the statutes of the EU General Data Protection Regulation (GDPR) (European Parliament and the Council of the European Union, 2016), which mandate transparency with respect to how personal data are processed. The Regulation considers data transparency a prerequisite for enabling data subjects to make informed decisions about intervening with the processing of their personal data, i. e. the right to access, rectification, to object to processing and profiling, and to have their data erased (GDPR Art. 12 et seq.). The deviation from the legal statutes is particularly remarkable because 'data concerning health' are considered special categories of data (GDPR Art. 9) whose processing warrants special care and responsibilities on the part of data controllers (Art. 29 Working Party, 2011).

One way of providing users of online data services with insight about the processing of their personal data is by means of *ex post transparency-enhancing tools* (TETs). Ex post TETs provide intelligible information about how their personal data *have been processed*. In this respect, ex post TETs differ from *ex ante* TETs, the latter of which communicate risk and potential outcome *before* users perform an action, such as before signing up for a data service or before installing an app. For

DOI: 10.4018/IJMHCI.2019100106

the sake of readability and unless the context in question requires clarification as to whether it refers to an ex ante or ex post scenario, the term 'TET' will be used in lieu of 'ex post TET' throughout this article to refer explicitly to *ex post* transparency-enhancing tools.

TETs retrospectively provide users of online data services with transparency about the processing of their personal data and guide them in making informed decisions with respect to managing the data they have disclosed previously. Hence, TETs can serve as indicators of facts that help users to hold data controllers accountable for how their personal data have been processed. The medium that facilitates ex post transparency discussed in this paper is privacy notifications, which notify users about events related to personal data processing deemed relevant for them. However, many TETs discussed in the literature are limited in terms of their usability in that their design does not systematically reflect the needs of their final users (Murmann and Fischer-Hübner, 2017a). This suggests research that addresses usable TETs specifically through the lens of human-centred design (International Organization for Standardization, 2010), and motivates the following research questions:

1. What kind of model is required to adequately describe the conceptual and functional nature of TETs that employ privacy notifications to enable intervenability?
2. What findings exist in the body of literature that lend themselves to conceptualising design guidelines for privacy notifications received on mobile devices?
3. What guidelines can be inferred from the model and the findings in the literature for the design of TETs that best reflect the individual needs of their users?

The context of use considered throughout this article is ex post transparency of personal data processed in mhealth environments. Data subjects disclose their health data to online services who process the data, and potentially share them with third parties. As users of the online mhealth services, data subjects take the role of both originators and auditors of their personal data, and achieve data transparency by obtaining retrospective information about how their data have been processed by means of a TET. In this regard, mhealth specifically pertains to scenarios of fitness tracking and self-quantification. Conceptually, this domain differs from clinical mhealth environments in that personal data generated therein are managed and audited by the data subjects themselves rather than by a professional caretaker or administrator. Hence, the final user of a TET is typically a lay person with respect to information privacy. The TET itself is considered to be a technological artefact in the form of an application running on a mobile device that is controlled by the user.

The main contribution of this paper is to analyse the findings established in the literature, to apply them to a model that reflects the asynchronous nature of privacy notifications, and to derive from them guidelines for the design of TETs that leverage usable privacy notifications as a means to facilitate ex post transparency in the context of mhealth environments.

The rest of this paper is structured as follows: Section 2 summarises related work, gaps detected in the literature regarding usable TETs, and findings on user preferences for privacy notifications. Section 3 presents a model that captures the phases users traverse during the interaction with a TET that relies on privacy notifications. Section 4 infers design guidelines for usable privacy notifications based on the findings obtained from the literature. Section 5 briefly discusses the model and guidelines, and Section 6 concludes the paper.

2. PREVIOUS WORK

To motivate the model and guidelines presented in Sections 3 and 4, respectively, this section draws upon three types of research: Section 2.1 briefly touches upon publications related to establishing models for privacy-enhancing technology, Section 2.2 reports on gaps with respect to usable TETs detected in the body of literature, and Section 2.3 reports on the findings of a recently conducted study on user preferences for privacy notifications.

2.1. Related Work

This paper is motivated by the concept of push and pull notices as discussed by the Article 29 Working Party (Data Protection Working Party, 2016) with respect to their guidelines on transparency. One way of facilitating transparency suggested by them is by means of just-in-time notices that provide users of data services *ad hoc* with 'privacy information' relevant for them. Similarly, Patrick and Kenny (Patrick and Kenny, 2003) discuss such notices as a means to inform users about details regarding data processing as needed, i. e. while a user navigates and 'clicks through' a particular transactional context, such as while signing up for a service. The rationale underlying such notices is that specific information on personal data processing are revealed one at a time rather than in the form of a monolithic privacy notice or policy. By doing so, incremental notices reduce the cognitive load imposed on the reader, while improving the contextual associativity. Both groups of authors aim to provide users of data services with meaningful information about future data processing, and hence enable ex ante transparency. Conversely, this paper seeks to apply the aforementioned principles to scenarios pertaining to ex post transparency, in which just-in-time notices take the shape of privacy notifications triggered in response to how a data subject's personal data have been processed.

The paper is related to the work by Feth et al. (Feth et al., 2017) in that both papers build upon conceptual models motivated by the realisation that human-centred design as specified in ISO 9241-210 (International Organization for Standardization, 2010) is indispensable for creating privacy-enhancing technologies with sustainable efficacy. However, the model presented by Feth et al. is generic in that it deals with the processes and stakeholder relationships related to information security for the Internet of Things at a highly conceptual level in which data transparency represents but a concomitant factor rather than a specifically intended design goal.

Similarly, the model presented in this paper is related to the model presented by Bravo-Lillo et al. (Bravo-Lillo et al., 2011) in that both models reflect the behaviour expressed by human actors who are confronted with security warnings. Both models cover the phases traversed while interpreting, deciding and acting upon a warning, and draw upon individualisation, and thus customisation, to satisfy individual users' needs. However, Bravo-Lillo et al.'s model focuses on the concepts of notice, choice, consent and control in the context of ex ante scenarios. These concepts essentially correspond to a user's decision of whether or not to use an app or service based on the terms and conditions presented to her. Hence, this model is only partially applicable to ex post transparency in that the transactional context in which previous decisions have been made may not be be immediately accessible to a user who receives a privacy notification.

Conversely, the paper at hand specifically positions itself in the usage context of fitness tracking and mhealth. Against this backdrop, it conceptualises a model that reflects the interaction between a user and an ex post TET under the premise of facilitating ex post transparency by means of privacy notifications.

2.2. Gaps of TETs

The authors conducted a survey on usable ex post TETs presented in the literature at that time (Murmann and Fischer-Hübner, 2017a), the sources of which were based on the result of a systematic literature review conducted previously (Murmann and Fischer-Hübner, 2017b). The contribution of the paper is two-fold: Firstly, their work provides a taxonomy of usable ex post TETs discussed in the literature. Researchers on usable privacy and data transparency can rely on a classification system to classify TETs according to particular characteristics and compare them with each other. Secondly, the researchers conducted an analysis of various gaps with respect to the usability of the TETs discussed in the literature. The analysis is based on the usability heuristics established by Nielsen (Nielsen, 1994) and the design principles for the ergonomics of human-system interaction specified by ISO 9241-110 (International Organization for Standardization, 2006). These references were complemented with the principles proposed by (Patrick and Kenny, 2003) for the design of interaction systems in the domain of information privacy and the statutes of the GDPR.

The following gaps have been discussed in the survey:

Self-descriptiveness. Many designers of TETs fail to visualise how, why, and by whom personal data are processed through forms that are easily accessible for users. Hence, many TETs fail in descriptively presenting scenarios that exceed a certain level of complexity, or that require technical or domain knowledge on the part of the user.

Standardized icons. Iconography has not been standardised to a degree that icons related to data privacy would allow for immediate recognition by lay persons. So far, many complex phenomena related to this topic have escaped standardised forms of depiction that are both self-descriptive and universal.

Transparency and control of judgemental statements. Judgemental statements are statements made by a TET to emphatically signal states of operation that reflect context-specific levels of importance and criticality, such as by means of colour coding, iconography or suitable graphics. In some cases, such statements are employed to stimulate awareness or even action on the part of a user. However, the cues and threshold values underlying particular levels of criticality are often not transparent. Moreover, these mechanics are often not adaptable to the needs of individual users, which questions the utility of judgemental statements across different users and usage contexts.

Suitability for individualisation. Individualisation supports the mental models of different groups of users. Hence, individualisation is crucial to satisfy the needs of individual users across different socio-cultural environments, usage contexts and levels of proficiency.

Support intervenability and accountability. Few TETs support their users in leveraging their legal right of intervenability, which limits their practical use in terms of acting upon the insight gained, and thus holding data controllers accountable for the processing of the user's personal data. Intervenability covers, e. g., the right of data subjects to withdraw their consent for the processing of their personal data (GDPR Art. 7 (3)), rectify their personal data (Art. 16), have their personal data erased (GDPR Art. 17), have their data transferred to another service provider (Art. 20), or object to automated decision-making and profiling (Art. 21, 22).

Error prevention. Preventing users from misinterpreting results and making decisions that do not reflect their best interest is important, yet not always implemented by TETs. Personal data are considered sensitive and valuable assets. However, gauging the way such data are processed and deciding upon how to react is not always straight forward. Hence, user errors may lead to consequences that are both unexpected and unpleasant. Ideally, TETs would not only detect incorrect settings, but would also guide users in configuring the settings based on their individual needs.

Personal data breach notification. Art. 34 of the GDPR mandates breach notifications being issued under certain circumstances (Data Protection Working Party, 2018). However, most TETs do not facilitate receiving breach notifications from data controllers, or support users in putting respective messages into context. Hence, users may not receive breach notifications at all, and if they do, they might be unable to contextualise the message, or fail to act upon it accordingly.

Considering the aspects of individualisation and intervenability from the perspective of human-centred design raises the question as to what extent the concepts above are reflected by the needs of the intended target audience, i. e. by users of online mhealth services. Moreover, former research showed that the participants of a study on ex post TETs expressed doubt with respect to the efficacy of data transparency unless they were actually able to "do something practical about it" (Fischer-Hübner et al., 2014). This motivated a study that investigated the preferences of users of mhealth services with respect to receiving customised privacy notifications. The study examined the aspects of personal predisposition and preferences, as well as the effect of intervenability on the users' choice to receive privacy notifications.

2.3. Preferences for Privacy Notifications

In the context of fitness tracking, the user's primary goal is to track her health, either momentarily or in the form of longitudinally quantifying her performance. Conversely, goals related to data security or information privacy rarely qualify as primary tasks (Whitten and Tygar, 1999). A task that depends

on proactive measures on the part of a user to enhance her privacy is unlikely to be conscientiously employed on a regular basis (Pfleeger et al., 2014), and is therefore unlikely to yield sustainable efficacy. Given the fact that information privacy is orthogonal to health tracking, mhealth environments necessitate techniques that enable data transparency without proactive scrutiny on the part of a user.

Triggered by external events related to personal data processing, *privacy notifications* provide users with asynchronous information about how their personal data are processed by online data services. To yield sustainable efficacy, such notifications must reflect the recipient's individual needs in terms of what she wants to be notified about, and what, if any, additional guidance she requires to understand the informational content of the message delivered.

A study recently conducted by Murmann et al. (Murmann et al., 2019) investigated the aspect of individualisation and intervenability with respect to the notification preferences of users of online mhealth services. It assessed how scenarios related to personal data processing can be grouped, to what extent these groupings were reflected in the participants' choice to receive notifications about respective scenarios, how notification preferences related to privacy personas established in the literature (Morton and Sasse, 2014, Morton, 2015), and whether intervenability affected the participants' choice to be notified. The findings of the study are as follows:

Categorisation. The researchers considered three categories of scenarios related to online data processing, Personal data breaches, Consequences, and Practical tips, for which the participants indicated different notification preferences. *Breaches* refer to the deliberate misappropriation, accidental loss or alteration of personal data (GDPR, Art. 4 (12)) (Data Protection Working Party, 2018). *Consequences* describe conditions of how personal data are actually processed or could be processed based on information currently at the processor's disposal. *Practical tips* seek to improve a user's privacy by providing her with personal advice and customised soft nudges based on her situational context. Categorising privacy notifications may allow designers of TETs to provide users with presets of notification settings based on thematic groupings of the underlying circumstantial cues.

Personas. Personas aid designers of interaction systems in assessing the needs of their target audience (Cooper, 1998). In this regard, the researchers tried to find a correlation between the notification preferences indicated by the respondents of the study and *privacy personas* established in the literature for the context of user trust in data services by Morton et al. (Morton and Sasse, 2014, Morton, 2015). The researchers were unable to reproduce the privacy personas, but provided first results of an alternative segmentation based on the respondents' notification preferences. The two preliminary segments reflected clusters of notifications received in response to events that (1) had actually taken place, and, conversely, (2) comprised indicators of hypothetical risks of personal data processing that could be conducted in the future. Once refined, these segmentations may provide users of TETs with selectable presets for notification settings.

Impact of intervenability. The researchers showed that intervenability had an impact on the participants' choice to be notified about the processing of their personal data. At the same time, the participants expressed uncertainty as regards the meaning of intervenability and as to how respective legal rights can be exercised. Designers of TETs will have to take both factors into consideration when they offer respective options to leverage intervenability.

Along with the gaps discussed in Section 2.2, these findings motivate the asynchronous model proposed in Section 3.2 and the design guidelines for privacy notifications discussed in Section 4.

3. MODELLING

This section presents the two evolutionary phases of the model that describes the functionality of TETs based on privacy notifications. The second, asynchronous model, complemented with the findings from the literature (Section 2) reflects the basis for the guidelines of human-centred design of notification-based TETs discussed in Section 4.

3.1. Synchronous Model

The first model of how TETs operate (Figure 1) was described in the author's survey on TETs (Section 2.2) and covers three distinctive operation steps: (1) Disclosing personal data to a data service, (2) reviewing how the personal data have been processed, and (3) acting upon insight gained in the previous step by intervening with the processing of one's personal data. Each of the steps is accompanied and supported by the TET.

Abstract and highly conceptual, this model essentially implies proactive scrutiny on the part of a user in terms of reviewing and, if necessary, intervening with the processing of her data. It is therefore only partially suited to accommodate the asynchronous nature of the privacy notifications described in Section 2.3, and to motivate design guidelines that specifically address this nature. Consequently, this article proposes an alternative model that is derived from the previous one, and that specifically accounts for the particularities of TETs as technological artefacts running on mobile devices. The asynchronous model presented hereafter does not invalidate or replace the initial model. Its purpose is instead to address the characteristics of a TET that facilitates ex post transparency by means of privacy notifications.

3.2. Asynchronous Model

The findings from the study discussed in Section 2.3 show that most participants wanted to receive privacy notifications, which motivates a model that leverages notifications as a means of ex post transparency. Building on the categories of notifications elicited from the study, the role of the TET envisioned for this model is that of a technological artefact that employs privacy notifications as a means to notify data subjects about personal data breaches, i. e. about non-compliance with respect to how her personal data have been processed, about the consequences of disclosing her personal data, and about practical tips about how to improve her privacy. The asynchronous model comprises phases of user interaction between a user and a TET. Unlike the first model, it not does consider a Disclose phase during which a user would interact with the TET. The model is instead based on the premise that personal data have been disclosed in the past, and that the process of reviewing and intervening with the processing of these data is decoupled from the previous disclosure.

Figure 1. Model of the phases related to actions conducted by data subjects who disclose personal data to an online service according to (Murmann and Fischer-Hübner, 2017a). The lower half of the illustration signifies the role of a TET used to enable the actions depicted in the upper half.

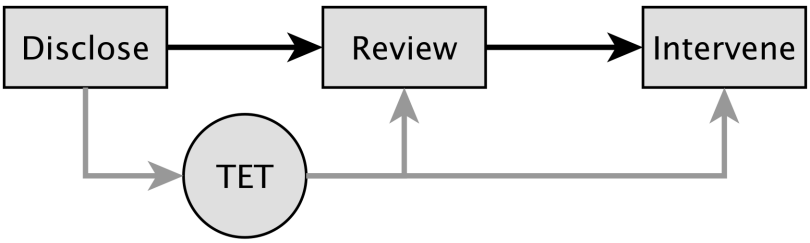

Figure 2. Phases of user interaction facilitated by a notification-based TET

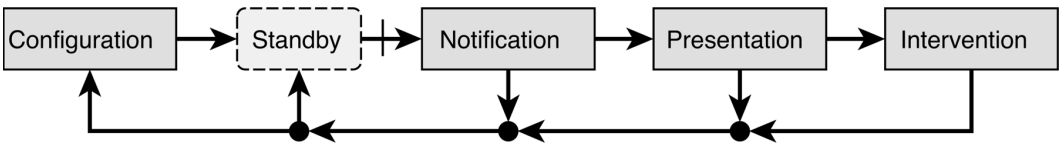

The model is based on the concept of delivering notifications that are either triggered by external events, or that are based on heuristics applied to intelligence gathered previously by querying an online service for updates about how the user's personal data have been processed. In the former case, privacy notifications take the form of *push notifications*, such as a personal data breach sent by the online service or by a regulatory authority (GDPR, Art. 34). Push notifications qualify as asynchronous events in that they originate from a remote entity and in that they are the result of an event whose occurrence and outcome are unexpected by the recipient of the notification. However, since neither the sender nor the recipient of a push notification are in full control of the communication channel between the two entities, it is unlikely that data services will rely on such means to transmit messages of critical significance.

It is more likely that TETs, which ultimately serve a data subject's, i. e. the user's interests, either run on the user's device, or are supplied by a trusted third party on whom both service provider and user rely. In this case, the TET actively queries updates from the data service in periodic intervals, screens the data for patterns of interest, and triggers respective notifications on the user's device. In this regard, such notifications rely on *pull technology*, which differs from push notifications in that even though the results of a query may be unexpected, the occurrence of a notification, i. e. its timing, is deterministic in that it is issued by the TET.

The phases of the asynchronous model are depicted in Figure 2 and reflect stages of interaction between the user and the TET. Arrows indicate a sequential order of the phases, and denote mutual dependence of one phase on its predecessor. The phases serve the following purposes:

Configuration. When the TET is first put into operation, it starts in the Configuration phase. During Configuration, the TET ascertains the user's preferences in terms of being notified and the amount of guidance provided during use.

All interactions observed during the subsequent phases can reflect back on the settings of the TET in that choices made by the user are taken into account to determine future behaviour. Hence, configuring the run-time behaviour of the TET can be both explicit and implicit. However, explicit configuration always takes precedence over recommended settings, and users can return to the configuration phase at any time.

Standby. Standby is a phase of non-interaction, and has been included in the model to represent the missing link between Configuration and Notification. During Standby, the TET lies dormant until it either receives a push notification from a remote entity, or until it queries a remote service for updates on how the user's data have been processed. In either case, the TET scans the information obtained, ascertains its context and severity, and compares the results with the user's notification preferences. The crossed out arrow signifies the absence of transactional synchronicity between the Notification phase and its predecessor.

Notification. Notification relates to choosing an opportune moment to notify the user about a circumstance worthy of her attention. Time and modality chosen to deliver a notification are based on the nature and severity of the incident and on the user's notification preferences.

Presentation. Presentation denotes the process of presenting intelligible facts about an incident, and to provide users with contextually meaningful choices and recommendations of how to proceed. Presentation can comprise multiple layers of contents, and represents the interaction conducted in each of them. It includes, e. g., the abstracted synopsis of privacy notifications used to provide recipients with a quick overview of the phenomenon in question, and optionally follows up this view with secondary facts of contextual details. Form and extent of the data thus displayed depend on the user's preferences.

Intervention. Intervening with the processing of the user's personal data is treated independently from Presentation in that this step relies heavily on the options provided by the online data service in question. In essence, TETs guide users in terms of exercising their right of intervenability based on the information obtained. However, they can only suggest certain choices, and will reach their practical

limits once those cues are acted upon, in which case users will have to interact with the infrastructure of the data service. TETs can support users in terms of exercising their right of intervenability by referring them directly to the appropriate contact point or official in charge to file their request.

4. DESIGN GUIDELINES

This section proposes guidelines for the design of usable TETs that operate on the basis of privacy notifications. The guidelines are derived from the gaps detected in the literature (Section 2.2) and the study on privacy notification preferences (Section 2.3), and are complemented with findings from the literature established on usable privacy and user perception of notifications received on mobile devices. The guidelines cover all phases of the asynchronous model introduced in Section 3.2 and therefore provide designers with holistic principles for designing TETs.

As regards the literature cited in the guidelines, the sources draw upon two separate fields. Firstly, many of the references relate to research conducted in usable privacy and security. However, most of this research deals with usage contexts rooted in ex ante scenarios. Research in this field focuses primarily on ex ante *security warnings* or *privacy notices* rather than ex post *privacy notifications* as discussed in this article. Security warnings deal, e. g., with phishing alerts issued by an anti-phishing tool as soon as users run the risk of performing an action that jeopardises their security, whereas privacy notices provide insight about how the user's personal data will be processed once she signs up to a service or starts using a particular mobile app.

Many of these findings rely on the concept of a transactional context, i. e. a contextually coherent sequence of operation steps carried out by a user, such as receiving a security warning *while* interacting with a website or being confronted with a privacy notice *in the process of* installing an app. Such reflexive coherence has no suitable equivalent in the context of asynchronous privacy notifications, which reach their recipients unexpectedly and often out of context, i. e. without an immediately recognisable causal relationship. Hence, establishing contextual comprehension *ad hoc* along with presenting privacy notifications plays a pivotal role in achieving usability of TETs that employ notifications as a means to facilitate ex post transparency (see guideline *Provide contextual cues* below).

This is why the guidelines draw upon a second research field that deals with the asynchronous nature of notifications, and that specifically considers the effect interruptions of a user's primary task may have on their recipients. Collectively, the superset of both research fields exemplarily captures the salience and efficacy of privacy indicators and the asynchronous nature of privacy notifications. The literature reviewed for both fields represent the result of a non-exhaustive query of the databases maintained by ACM, IEEE and Springer, and considered the first 50 search results of each publisher. These sources were complemented with selected papers retrieved during the literature review discussed in Section 2.2.

The guidelines are as follows:

4.1. Use Interruption Conscientiously

(Sahami Shirazi et al., 2014) point out that receiving notifications is generally considered disruptive on the part of the recipients. If a notification is eventually dealt with, its response time is inversely proportional to its perceived importance (Sahami Shirazi et al., 2014). However, notifications are typically strongly related to personal messages sent by peers (Sahami Shirazi et al., 2014), whereas other types of notifications are often considered less urgent in terms of response time (Weber et al., 2018), or bear a negative connotation in terms of interruption (Pielot et al., 2014).

Weber et al. (Weber et al., 2016) report that interruptions via notifications received on smart TVs were mainly perceived negatively by test subjects, and suggest to carefully consider whether and when to display on-screen notifications. Despite their apparently different usage context, mobile devices are frequently used to play visual content or to engage in ongoing foreground activity, such

as web browsing or messaging. In cases of directed attention towards the device, mobile computing therefore exhibit similar perceptual characteristics in terms of posing as a continuous primary task.

In terms of urgency, (Wilson and Miller, 2006) specify '(hard) alerts' as "time-sensitive notifications, where failure to act on the indicator within a certain (predictable) time bound may be costly" (Wilson and Miller, 2006). (Egelman et al., 2008) argue that a user's primary task may be interrupted by security indicators if the primary and the interrupting task are contextually related. This is in line with the recommendation by (Schaub et al., 2017) who suggest that displaying privacy notices is acceptable as long as they pertain to the same transactional context, and thus correlate with the user's expectation. Pielot et al. (Pielot et al., 2017) report that notifications are highly context-dependent in that benevolently receiving and acting upon them depends on various factors, such as the current time of day, location, current and previous activities, and the device settings. However, Pielot and et al. show that by relying on previous user decisions machine-learning can be successfully leveraged to predict opportune moments with high accuracy.

As for privacy notifications, TETs will have to judge interruption based on the severity of the incident, i. e. the cost of failing to act upon the notification (see Prevent user errors). Whether receiving privacy notifications is perceived as costly, however, is likely to be individual. This suggests that TETs should rely on multiple notification categories that enable users to assign them different priorities or modalities, and thereby allow for different ways of handling incoming notifications. The findings of the study on privacy notification preferences in Section 2.3 show that users of mhealth services have different notification preferences for different categories of scenarios, such as data breaches being considered more worth receiving than other categories. (Patil et al., 2015) report that incidences with respect to information privacy that are not critical may be moderately delayed. Similarly, (Micallef et al., 2017) report that (soft) privacy nudges are considered low priority notices by most recipients. Hence, incidences should be weighed up as to whether they require the user's immediate response, whereas practical tips and notifications that reflect non-critical circumstances should be allowed more temporal leeway. The study on notification preferences (Section 2.3) shows that the majority of the participants wanted to be notified across all categories. This suggests that delivering privacy notifications of any type is a reasonable preset, but that users should be able to opt out of receiving further notifications based on the underlying category or content.

In essence, TETs should seek to interrupt users sparingly, and only if the user's privacy is immediately at stake. If in doubt as to whether a user wants to be notified about a circumstance related to her privacy, a TET should send the notification in question, but at the same time offer her to opt out of receiving further notifications. The final say as to the when and how of receiving privacy notifications should be up to the user, preferably by linking individual notification categories to individual modalities.

4.2. Leverage Multiple Modalities

(Schaub e tal.,2017) suggest to leverage multiple modalities available on today's mobile devices to notify users about privacy-related events, such as visuals (text and icons), sound, vibration and blinking LEDs. However, they warn that unfamiliarity with a particular modality might lead to uncertainty or even ignorance on the part of a user. In their study on how users of mobile devices prefer to be notified about various circumstances, (Micallef et al., 2017) report that users prefer salient modalities, such as audio or speech, for alarms but not for nudges. More importantly, the notification settings preferred for nudges was found to be almost opposite to the settings used for notifications received from conventional apps.

(Micallef et al., 2017) report that many recipients of salient messages feel uncomfortable in the presence of other people. This is in line with the findings of (Weber et al., 2018) who report that users of mobile devices prefer not to be disturbed during the forenoon, and, if provided the option, rather prefer to look after low priority notifications in the evening or night. Similarly, Weber et al. (Weber et al., 2016) suggest subtlety of on-screen notifications received on smart TVs in the presence

of onlookers. Highly personal information, such as privacy notifications received while residing in a public space, warrant particular diligence in terms of being displayed on a mobile device that might be accidentally observed or deliberately watched by an audience other than the intended recipient.

TETs should respect the user's daily routine. Since privacy notifications are asynchronous in nature, TETs have a certain leeway as regards the exact time of delivery, particularly if a notification results from an active query that can be scheduled at any time. The modality chosen to signal the arrival of a privacy notification may correspond to the category and content of the incident it reports, and should accommodate the recipient's choice of how she wants to be notified. Customisability regarding signalling privacy notifications can help distinguish privacy-related contents from messages received for other purposes, a large variety of which already vie for their recipients' attention.

4.3. Present Facts Intelligibly

(Bravo-Lillo et al., 2011) suggest that privacy notices should be brief. Privacy notices are often presented in the form of verbose legalese written by lawyers for the purpose of legal compliance but not for readability or transparency. The authors point out that the more verbose the text, the less probably such notices will be read. However, in both ex ante and ex post transparency, reading and understanding privacy indicators is crucial for acting upon them. (Bravo-Lillo et al., 2011) and (Bal et al., 2014) recommend that legalese and technical jargon should be avoided when addressing lay persons. (Schaub et al., 2017) recommend that privacy indicators should be specific in that they should be perceivable as part of the user's transactional context (see *Provide contextual cues*). GDPR Art. 12 (1) mandates that 'clear and plain language' be used to facilitate transparency. Moreover, Art. 12 (7) encourages the use of 'standardised icons' in the context of data transparency. However, respective iconography have not been established to date (Section 2.2), which is why their use for the purpose of visualising complex scenarios should be carefully considered such that the graphical depiction in question is clearly intelligible by the recipient. Singular icons and sets of icons should follow a clear pattern to support recognition rather than recall (Nielsen, 1994) and support learnability (International Organization for Standardization, 2006).

Facts presented in privacy notifications should accommodate the mental state of their target audience. Recipients will most likely receive the message unexpectedly, may be overwhelmed by the details, and might struggle to put it into context (Angulo and Ortlieb, 2015). Hence, information presented by a TET should be clear and concise, leading to a cognitive state that Patrick and Kenny call 'comprehension' (Patrick and Kenny, 2003), i. e. the ability to make sense of the contextual cues displayed by an information system.

4.4. Use Multiple Levels of Detail

Complementing the findings on intelligibility, (Bal et al., 2014) suggest to first present meaningful summaries of the circumstantial cues of an event. (Anthonysamy et al., 2017) point out that principles that facilitate the design of usable privacy controls should not only aim for comprehensibility, but should also seek to minimise their users' cognitive load. Similarly, (Schaub et al., 2017) suggest to present optional, more detailed information on secondary screens. The Article 29 Working Party (Data Protection Working Party, 2016) seconds this in Sections 30 et seq. of their 'Guidelines on transparency', and advises policy makers to make use of multilayered privacy notices.

Although most of the aforementioned recommendations reflect concepts related to ex ante transparency, multi-layered information can help improve comprehensibility at large, and therefore help facilitate transparency. (Murmann and Fischer-Hübner, 2017a) point out that multimodal presentation and multiple layers of detail help scrutinise complex phenomena from multiple perspectives and satisfy the needs of multiple groups of users. Different preferences with respect to different levels of detail have also been reported in these authors' latest study on privacy notifications, in which some participants expressed interest in obtaining further information to decide how to act upon receiving certain types of privacy notifications. Awareness of the details pertaining to data processing correspondents to

what Patrick et al. refer to as 'consciousness' (Patrick and Kenny, 2003), a prerequisite for making informed decisions with respect to deciding how to intervene with the processing of one's personal data.

TETs that rely on privacy notifications should first provide a brief synopsis of the cause of the notification, and follow it up with further details upon request. It should be only at a secondary or tertiary stage that TETs present a meticulous review of the circumstances underlying an incident, such as detailed logs of individual data transactions.

4.5. Provide Contextual Cues

In scenarios related to ex ante privacy, users seek to finish a primary task, such as using an online data service, and are distracted by cognitive load imposed upon them by privacy-enhancing technology in the form of warnings and notices. (Schaub et al., 2017) suggest to render such information 'relevant' by contextualising them, i. e. by establishing contextual coherence between the primary task and a privacy indicator. (Bal et al., 2014) refer to 'exo information' as supplementary details that accompany the raw data of a fact to elevate it from a context-free state and put it into reproducible context. The recommendations of both groups of authors seek to enable users to better understand the circumstances responsible for triggering the indicator in question by clarifying the causal chain that links the displayed indicator to its cause.

More often than not, a transactional context will not exist for asynchronous notifications. Users will be preoccupied with tasks entirely unrelated to ex post transparency, or will not attend their devices at all. Hence, TETs will have to provide them with meaningful cues about the circumstances that lead to the notification being sent, and thus create a reproducible chain of effects. By contextualising cues, TETs illustrate prospective outcomes accessibly for the user experiencing them, which may help improve the efficacy of such notifications. Ideally, such cues should include information about how choices made in the past and events initiated by third parties are related to each other, and how both have led to the circumstances observed and reported in the present.

4.6. Communicate Risks and Consequences

Lay persons rarely foresee the full spectrum of possibilities that may result from disclosing their personal data to data services, such as their data being shared with third parties or them being profiled based on data obtained from multiple entities. (Bal et al., 2014) and (Bravo-Lillo et al., 2011) recommend to clarify the consequences as a result of a future action, and to provide instructions of how to avoid such risks. Moreover, (Bravo-Lillo et al., 2011) suggest to explain the consequences of performing or refraining from performing individual actions related to a risk. (Schaub et al., 2017) recommend to prioritise privacy indicators based on the level of risk or severity of the outcome, and to help users weigh individual options against each other.

These advices refer to ex ante privacy notices that seek to avoid predictable risk in the first place. However, the concept of supporting users in gauging personal risk for their privacy, and to help them better understand the consequences of past and present actions, equally applies to ex post transparency. Hence, TETs should indicate the consequences resulting from their previous choices, as well as the risk of choices made in the present. The consequences for a user's privacy reported as part of ex post transparency can be perceived differently depending on whether a notification deals with the outcomes of actual past events or hypothetical future events. This is reflected in the study by Murmann et al. (Murmann et al., 2019), in which the researchers reported a preliminary user segmentation based on different notification preferences for both types of notifications (Section 2.3).

4.7. Provide Actionable Choices

(Egelman et al., 2008) suggest that security warnings alone may lead to ignorance, whereas choices and recommendations offered alongside indicators can lead to contextual advices being heeded and risk being avoided. (Schaub et al., 2017) suggest that such options should not only be intelligible but

also be actionable in that the user's status quo, the options available at the time of decision-making, and prospective outcomes should be perceivable as being contextually linked (see *Provide contextual cues*).

Actionable choices, as they are discussed in the literature, primarily reflect the notion of control discussed by Patrick and Kenny (Patrick and Kenny, 2003) for ex ante scenarios. In many cases, this type of exertion of influence boils down to binary choices, such as using or not using the service or app in question. In the context of ex post transparency, the spectrum is wider in that intervenability potentially pertains to multiple applicable options that enable users of data services to hold data controllers accountable for the processing of their personal data.

Having options at their disposal that facilitate intervenability may have a positive impact on the users' willingness to review, and if necessary, act upon insight about how their personal data are processed. Murmann et al. (Murmann et al., 2019) report that intervenability impacted their participants' choice to receive privacy notifications, even though some respondents indicated that the effect and scope of the concept was not clear to them. TETs will therefore have to convey such background knowledge to users, and inform them about the consequences different forms of intervenability will have on their privacy.

Respective options will have to be contextual in that applicable choices should rely on the situation individual users act upon. If, e. g., the retention period of data disclosed to an mhealth service was exceeded, it would be pointless to encourage a user to rectify these data, whereas advising her to withdraw her consent for the processing of her personal data might be an expedient option.

4.8. Provide Support and Guidance

(Cranor, 2008) and (Krol et al., 2012) point out education and training as two factors that aid users in better understanding security and privacy-related tasks. (Bravo-Lillo et al., 2011) suggest that UIs should distinguish between advanced users and novice users, and should provide the latter with additional knowledge to help them understand the situation they find themselves in. (Brustoloni and Villamarín-Salomón, 2007) point out that context-sensitive guidance affords a "better balance between security and usability" than would alternative approaches based on security dialogues or automated decisions made by software (Brustoloni and Villamarín-Salomón, 2007).

As for ex post transparency, users will receive privacy notifications unexpectedly and extemporaneously, which might cause feelings of anxiety and helplessness on their part (Angulo and Ortlieb, 2015). Guidance can potentially mitigate the level of agitation and should be motivated by two objectives. Firstly, TETs should aid users in comprehending and contextualising notifications. Help pages should provide secondary information about the nomenclature and semantics of the phenomena notified about (see *Present facts intelligibly*). Secondly, a priori guidance enables users to better understand future decisions, such as the ones offered to them in the form of contextual choices (see *Provide actionable choices*). For many users, such options will not be self-explanatory, which is why they need additional support to contextualise their previous choices and the options currently available, as well as to weigh up multiple options against each other (see *Communicate risks and consequences*). By default, TETs should provide additional support in the form of help texts, but should respect the user's choice in terms of not being bothered later on. In any case, TETs should provide supplementary guidance upon request.

4.9. Prevent User Errors

(Schaub et al.,2017) recommend that privacy notices should highlight unexpected practices that deviate from established norms or practices. (Egelman et al., 2008) and (Brustoloni and Villamarín-Salomón, 2007) suggest that failing to read and act upon a privacy notice should be costly in that the inconvenience felt by circumventing privacy indicators should be noticeable for a user. Both findings are rooted in ex ante transparency, but the necessity to mitigate user error with respect to not acting upon privacy notifications equally applies to ex post transparency.

(Weber et al., 2018) suggest that incoming notifications should be available in an event history for later reference. If mobile devices handle incoming privacy notifications identical to regular system notifications, handling them might be prone to user errors, such as them being accidentally dismissed instead of being reviewed and acted upon (Pielot et al., 2018). Errors like these might be the result of misinterpretation or habituation, the latter of which has been shown to be opposable by means of polymorphic messages (Anderson et al., 2014).

TETs should highlight conspicuous results to prevent users from missing and failing to act upon them. They may provide options that, if enabled, require additional acknowledgement on the part of a user to dismiss privacy notifications from a message queue. Moreover, TETs should implement a secondary queue that stores dismissed notifications and reviewed incidences for archival purposes. This two-stage approach mirrors the paradigm of moving files or messages to a waste bin known from graphical desktop environments and email clients. Erasing singular notifications from the archive would still be possible, but doing so would effectively be more costly, and hence less error prone.

The same principle of benevolent paternalism should apply to actions chosen by users in response to receiving a privacy notification (see *Provide actionable choices*). In many cases, certain options may be more favourable than others, and TETs should saliently emphasise the ones that seem most appropriate. Similarly, TETs should scrutinise potentially unwise choices and demand additional confirmation in cases where a user's choice seems overly questionable or outright dangerous for her privacy.

4.10. Respect User Needs

In summary, TETs should provide meaningful sets of presets, such as to enable all categories of privacy notifications per default, and to allow users to opt out of receiving notifications about a certain category or type of scenario. Preferably, the default settings should be based on the user's actual predisposition, such as her privacy attitude. Ideally, a user's privacy attitude will link to a specific preset of notification settings that meets her demands in terms of striking a balance between privacy needs and additional cognitive load imposed by incoming notifications. The preliminary findings reported in the study about notification preferences (Section 2.3) indicate that respective segmentations can potentially be made, such as notifications about proven facts in contrast to potential privacy risks. It will, however, require further research to establish personas that reliably reflect requirements for multiple groups of users.

Moreover, the user's proficiency with the TET can serve as an indicator for the extent of active guidance or recommendation she is offered along with actionable choices. The user's previous knowledge would primarily reflect the extent of information offered for each operation step. However, proficiency should not affect the availability of supplementary support per se, which should always be available upon request.

Settings should be reflexive in that choices made in the past should be reflected in options offered in the present, which, in turn, should carry over to options offered for future decisions. Correspondingly, Angulo et al. (Angulo et al., 2012) propose a combination of predefined privacy settings with 'on the fly' privacy management. In their prototype, choices made by a user can carry over to the defaults of future options. Applied to the model for privacy notifications, choices made during the Notification, Presentation and Intervention phases may feed back to the Configuration phase. Regardless of the presets of a set of options, however, it should be up to the user to decide whether, how and when she wants to be notified, as well as which facts she wants to be notified about.

5. DISCUSSION

Most of the design guidelines discussed in Section 4 are overarching in that they apply to all phases of the asynchronous model. Despite the fact that some of the guidelines map exclusively to individual phases, such as conscientious interruption being exclusively associated with the Notification

phase, most phases of the model and the guidelines as a whole are closely interrelated and highly interdependent. This applies particularly to the Configuration phase, which not only controls the subsequent phases with respect to how they operate, but in that it also potentially receives feedback from them in terms of how users decide along the way, such as not to receive further notifications of a particular kind, or to receive more extensive explanation about a particular term or circumstance.

As *regards research question 1*, no model currently discussed in the literature reflects the particular nature of privacy notifications. Existing models from the domain of usable security are unsuitable in that they pertain to data security in general and do not cover privacy notifications in particular. Alternatively, some of the existing models originate from the domain of ex ante transparency. The candidates in question are tailored to the needs of coherent transactional contexts, and therefore do not satisfactorily reflect the asynchronous nature inherent to privacy notifications. Hence, this paper proposes a model that specifically deals with privacy notifications in the context of ex post transparency.

As for *research question 2*, the design guidelines proposed to accommodate the aforementioned model draw upon multiple fields of the literature. The sources reviewed for usable security and privacy have been complemented with literature on the interruptive nature of notifications and the effect such interruptions have on their recipients. Jointly, both research fields are shown to serve a basis for inferring suitable design guidelines for a TET that operates on privacy notifications.

As *for research question 3*, the guidelines proposed in this paper satisfy the particular nature of the asynchronous model introduced for implementing privacy notifications as a means to facilitate ex post transparency in mhealth contexts. The guidelines specifically consider the particular modalities and forms of interaction inherent to mobile devices as the designated target platform for the TET.

The design guidelines are based on established findings in the literature but remain preliminary in that the conclusions drawn by combining independent sources from multiple research fields will have to be consolidated by future research. Future research will, e. g., have to establish concrete sets of presets for notification settings, such as the ones discussed for guideline Respect user needs, as will be mappings between individual notification scenarios and actionable choices of how to intervene.

It will require mockups and prototypical designs of TETs to investigate whether the guidelines proposed above are mature enough to serve as viable requirements for the design of usable TETs, and whether they reliably reflect the functionality conceptualised in the model. Moreover, it will require further validation as to whether the model itself proves to be comprehensive enough to serve as a conceptual template for implementing TETs that exhibit the envisioned characteristics inherent to privacy notifications.

At a later time, the prototypes will have to be tested by the intended target audience, i. e. users of online mhealth services. The evaluation will show whether laypersons acknowledge the envisioned benefit in terms of ex post transparency of their personal data, and whether the concept of intervenability can be reliably conveyed by means of the prototypical implementation. This is when the circle of human-centred design will close and yield validated guidelines that can demonstrably serve as design principles for TETs that operate on the basis of privacy notifications. Alternatively, the second iteration of the design lifecycle will introduce a refined model and will aim to establish a revised set of guidelines.

6. CONCLUSION

The steady increase of mobile health (mhealth) devices are in stark contrast to the lack of data transparency experienced by many users of online mhealth services. The article draws upon multiple research fields represented in the body of knowledge and infers from them guidelines suitable for the human-centred design of ex post transparency-enhancing tools (TETs). The guidelines are discussed against the backdrop of a model that reflects the particular nature of TETs operating on the basis of privacy notifications, which serve as a means to provide users with transparency in terms of how

their personal data have been processed by online data services. The discussion of the model and guidelines focuses on the aspects of individualisation to accommodate the individual needs of users of mhealth services, and on intervenability as a form of exercising a data subject's legal right in response to how her personal data have been processed. Accounting for the particular characteristics of mobile devices, TETs that facilitate ex post transparency by means of privacy notifications aim to provide users of mhealth services with customised advice about auspicious future choices based on insight obtained retrospectively.

ACKNOWLEDGMENT

This work was supported by the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant 675730. The author would like to thank Simone Fischer-Hübner for her advice on structuring and reviewing the article.

REFERENCES

Anderson, B., Vance, T., Kirwan, B., Eargle, D., and Howard, S. (2014). *Users Aren't (Necessarily) Lazy: Using NeuroIS to Explain Habituation to Security Warnings*. Academic Press.

Angulo, J., Fischer-Hübner, S., Wästlund, E., & Pulls, T. (2012). Towards usable privacy policy display and management. *Information Management & Computer Security*, *20*(1), 4–17. doi:10.1108/09685221211219155

Anthonysamy, P., Rashid, A., & Chitchyan, R. (2017). Privacy Requirements: Present & Future. In *Software Engineering: Software Engineering in Society Track (ICSE-SEIS), 2017 IEEE/ACM 39th International Conference on*, (pp. 13–22). IEEE.

Art. 29 Working Party. (2011). *Advice paper on special categories of data ("sensitive data").* Technical Report Ares(2011)444105–20/04/2011, Article 29 Data Protection Working Party.

Bal, G., Rannenberg, K., & Hong, J. (2014). Styx: Design and Evaluation of a New Privacy Risk Communication Method for Smartphones. In *IFIP International Information Security Conference*, (pp. 113–126). Springer. doi:10.1007/978-3-642-55415-5_10

Bravo-Lillo, C., Cranor, L. F., Downs, J., & Komanduri, S. (2011). Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security and Privacy*, *9*(2), 18–26. doi:10.1109/MSP.2010.198

Brustoloni, J. C., & Villamarín-Salomón, R. (2007). Improving Security Decisions with Polymorphic and Audited Dialogs. In *Proceedings of the 3rd symposium on Usable privacy and security*, (pp. 76–85). ACM. doi:10.1145/1280680.1280691

Cooper, A. (1998). *The Inmates Are Running the Asylum: Why High-Tech Products Drive Us Crazy and How to Restore the Sanity*. Sams.

Cranor, L. F. (2008). *A Framework for Reasoning About the Human in the Loop*. Advanced Computing Systems Professional and Technical Association.

Data Protection Working Party. (2016). *Guidelines on transparency under Regulation 2016/679*. Technical report, 17/EN WP260.

Data Protection Working Party. (2018). *Guidelines on Personal data breach notifications under Regulation 2016/679*. Technical report, WP250rev.01.

Egelman, S., Cranor, L. F., & Hong, J. (2008). You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (pp. 1065–1074). ACM. doi:10.1145/1357054.1357219

European Parliament and the Council of the European Union (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council*. Author.

Feth, D., Maier, A., & Polst, S. (2017). A User-Centered Model for Usable Security and Privacy. In T. Tryfonas (Ed.), *Human Aspects of Information Security, Privacy and Trust* (pp. 74–89). Springer International Publishing. doi:10.1007/978-3-319-58460-7_6

Fischer-Hübner, S., Pettersson, J. S., Angulo, J., Edbom, J., Toresson, M., & Andersson, H. (2014). *D:C-7.3 Report on end-user perceptions of privacy-enhancing transparency and accountability*. Technical Report D37.3, Karlstad University.

International Organization for Standardization. (2006). *Ergonomics of human-system interaction – Part 110: Dialogue principles. Technical Report ISO 9241-110:2006(E)*. ISO.

International Organization for Standardization. (2010). *Ergonomics of human-system interaction – Part 210: Human-centered design for interactive systems. Technical Report ISO 9241-210:2010(E)*. ISO.

Knowles, B., Smith-Renner, A., Poursabzi-Sangdeh, F., Lu, D., & Alabi, H. (2018). Uncertainty in Current and Future Health Wearables. *Communications of the ACM*, *61*(12), 62–67. doi:10.1145/3199201

Krol, K., Moroz, M., & Sasse, M. A. (2012). Don't work. Can't work? Why it's time to rethink security warnings. *7th International Conference on Risks and Security of Internet and Systems (CRiSIS)*, 1–8.

Lau, J., Zimmerman, B., & Schaub, F. (2018). Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of ACM Human Computer Interaction*, *2*, 102:1–102:31. doi:10.1145/3274371

Micallef, N., Just, M., Baillie, L., & Alharby, M. (2017). Stop Annoying Me!: An Empirical Investigation of the Usability of App Privacy Notifications. In *Proceedings of the 29th Australian Conference on Computer-Human Interaction*, (pp. 371–375). New York, NY: ACM. doi:10.1145/3152771.3156139

Morton, A. (2015). *Individual Privacy Concern and Organisational Privacy Practice – Bridging the Gap* (PhD thesis). University College London.

Morton, A., & Sasse, M. A. (2014). Desperately Seeking Assurances: Segmenting Users by their Information-Seeking Preferences. In *Privacy, Security and Trust (PST), 2014 Twelfth Annual International Conference on* (pp. 102–111). IEEE.

Murmann, P., & Fischer-Hübner, S. (2017a). Tools for Achieving Usable Ex Post Transparency: A Survey. *IEEE Access : Practical Innovations, Open Solutions*, *5*, 22965–22991. doi:10.1109/ACCESS.2017.2765539

Murmann, P., & Fischer-Hübner, S. (2017b). *Usable Transparency Enhancing Tools: A Literature Review. Technical report*. Karlstad University, Department of Mathematics and Computer Science.

Nielsen, J. (1994). *Usability Engineering*. Elsevier.

Patil, S., Hoyle, R., Schlegel, R., Kapadia, A., & Lee, A. J. (2015). Interrupt Now or Inform Later?: Comparing Immediate and Delayed Privacy Feedback. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, (pp. 1415–1418). New York, NY: ACM. doi:10.1145/2702123.2702165

Patrick, A. S., & Kenny, S. (2003). From privacy legislation to interface design: Implementing information privacy in human-computer interactions. In *International Workshop on Privacy Enhancing Technologies*, (pp. 107–124). Springer. doi:10.1007/978-3-540-40956-4_8

Pfleeger, S. L., Sasse, M. A., & Furnham, A. (2014). From Weakest Link to Security Hero: Transforming Staff Security Behavior. *Journal of Homeland Security and Emergency Management*, *11*(4), 489–510. doi:10.1515/jhsem-2014-0035

Pielot, M., Cardoso, B., Katevas, K., Serrà, J., Matic, A., & Oliver, N. (2017). Beyond Interruptibility: Predicting Opportune Moments to Engage Mobile Phone Users. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, *1*(3), 91:1–91:25. doi:10.1145/3130956

Pielot, M., Church, K., & de Oliveira, R. (2014). An In-situ Study of Mobile Phone Notifications. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services, MobileHCI '14*, (pp. 233–242). New York, NY: ACM. doi:10.1145/2628363.2628364

Pielot, M., Vradi, A., & Park, S. (2018). Dismissed!: A Detailed Exploration of How Mobile Phone Users Handle Push Notifications. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI '18*, (pp. 3:1–3:11). New York, NY: ACM. doi:10.1145/3229434.3229445

Reinhardt, D., & Fischer-Hübner, S. (2019). To Be, or Not to Be Notified: Eliciting Privacy Notification Preferences for Online mHealth Services. *IFIP International Conference on ICT Systems Security and Privacy Protection*. (forthcoming)

Sahami Shirazi, A., Henze, N., Dingler, T., Pielot, M., Weber, D., & Schmidt, A. (2014). Large-Scale Assessment of Mobile Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (pp. 3055–3064). ACM.

Schaub, F., Balebako, R., & Cranor, L. F. (2017). Designing Effective Privacy Notices and Controls. *IEEE Internet Computing*, *21*(3), 70–77. doi:10.1109/MIC.2017.75

Statista. (2018). *Number of connected wearable devices worldwide from 2016 to 2021*. Retrieved from https://www.statista.com/statistics/487291/

Weber, D., Mayer, S., Voit, A., Ventura Fierro, R., & Henze, N. (2016). Design Guidelines for Notifications on Smart TVs. In *Proceedings of the ACM International Conference on Interactive Experiences for TV and Online Video* (pp. 13–24). New York, NY: ACM. doi:10.1145/2932206.2932212

Weber, D., Voit, A., Auda, J., Schneegass, S., & Henze, N. (2018). Snooze! Investigating the User-defined Deferral of Mobile Notifications. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI '18*, (pp. 2:1–2:13). ACM.

Whitten, A., & Tygar, J. D. (1999). Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. Usenix Security.

Wilson, T., & Miller, R. (2006). *Reducing the cost of interruption using gradual awareness notifications*. Retrieved from http://up.csail.mit.edu/projects/slowgrowth/gradual-awareness.pdf

Patrick Murmann is a PhD student at the Department of Mathematics and Computer Science at Karlstad University in Sweden. His research interest is the usability of transparency-enhancing tools that facilitate ex post transparency of personal data disclosed to online data services. His academic and professional background are computer science and media engineering.