PATRICK MURMANN

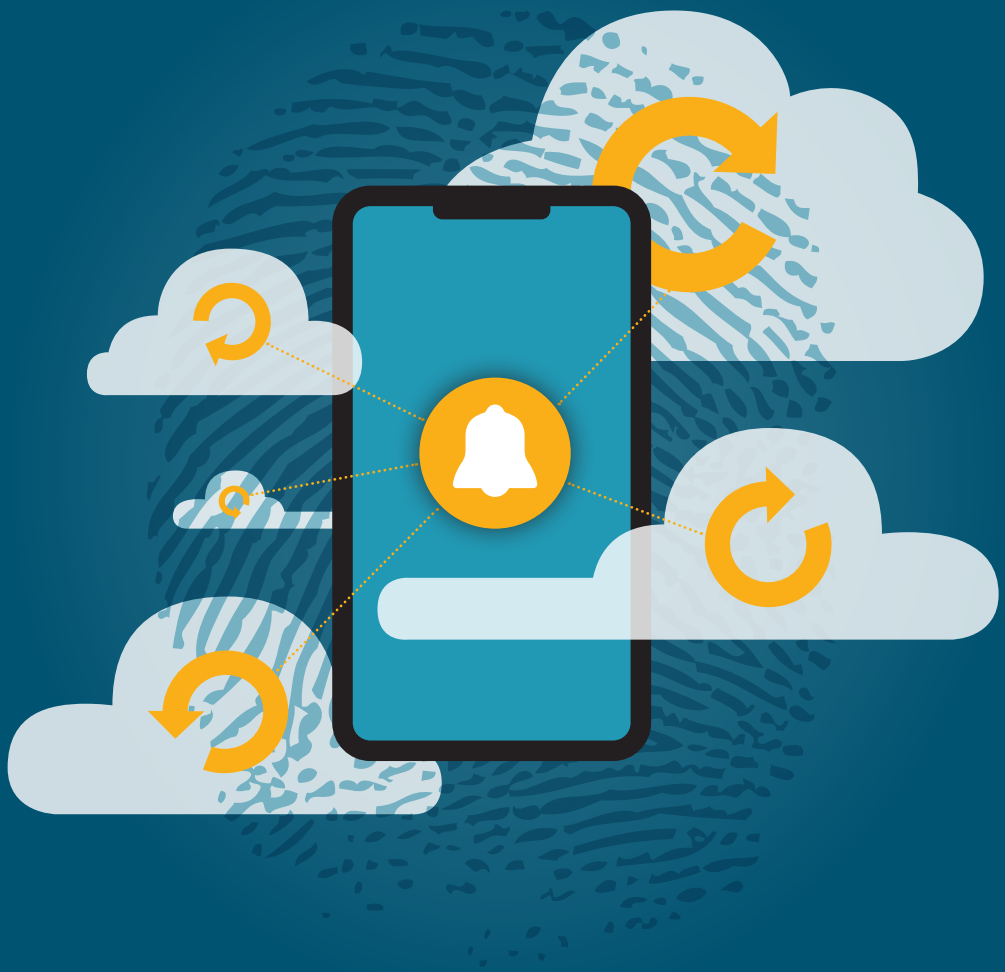

INFORMATION AT YOUR FINGERTIPS

Facilitating Usable Transparency via Privacy Notifications

Information at Your Fingertips

Facilitating Usable Transparency via Privacy Notifications

Patrick Murmann

Faculty of Health, Science and Technology

Computer Science

DOCTORAL THESIS | Karlstad University Studies | 2020:28

Information at Your Fingertips

Facilitating Usable Transparency via Privacy Notifications

Patrick Murmann

DOCTORAL THESIS | Karlstad University Studies | 2020:28

Information at Your Fingertips - Facilitating Usable Transparency via Privacy Notifications

Patrick Murmann

DOCTORAL THESIS

Karlstad University Studies | 2020:28

urn:nbn:se:kau:diva-80075

ISSN 1403-8099

ISBN 978-91-7867-144-1 (print)

ISBN 978-91-7867-148-9 (pdf)

Distribution:
Karlstad University
Faculty of Health, Science and Technology
Department of Mathematics and Computer Science
SE-651 88 Karlstad, Sweden
+46 54 700 10 00

Print: Universitetstryckeriet, Karlstad 2020

WWW.KAU.SE

Information at Your Fingertips: Facilitating Usable Transparency via Privacy Notifications

Patrick Murmann
Department of Mathematics and Computer Science

Abstract

The General Data Protection Regulation stipulates legal rights of transparency and intervenability. Transparency provides data subjects with insight into how their personal data have been processed, clarifying what consequences will or may arise due to the processing of their data, whereas intervenability enables them to intervene in the process. Technological artefacts, transparency-enhancing tools (TETs) serve the purpose of conveying respective information precisely and intelligibily. However, despite being a prerequisite for transparency, many TETs available today lack usability in that they do not stringently reflect the needs of their users, which raises the question as to whether individual TETs fulfil their designated purpose.

The objective of this dissertation is to systematically apply principles pertaining to human-centred design to ascertain the qualities necessary to design TETs that facilitate transparency and advise means of intervenability with regard to the needs of their target audience. We classify the state of the art of usable TETs published in the literature and discuss the gaps therein. Contextualising our research in the domain of personal health tracking, we investigate to what extent customisation can help accommodate the needs of users of TETs. We introduce privacy notifications as a conceptual means to inform data subjects about facts worthy of their attention, and examine the immanent properties required to accomplish actual usability. We categorise the characteristics of privacy notifications in terms of what insight they convey, and how respective facts need to be presented to facilitate informed decision-making on the recipient's part. Based on findings obtained via quantitative and qualitative user studies, we elicit concomitant factors related to the parameterisation of privacy notifications. We present the prototypical implementation of TETs whose iterative evaluation provides us with a catalogue of design requirements that demonstrably reflect the needs of their users.

Keywords: General Data Protection Regulation (GDPR), Human-centred design, Human-computer interaction (HCI), Information privacy, Intervenability, Mobile health (mhealth), Personal health tracking, Privacy notification, Transparency, Transparency-enhancing tool (TET), Usability.

Acknowledgements

I thank my supervisor, Prof. Dr. Simone Fischer-Hübner, whose profound insight and benevolent patience have guided me up to the point of defending my dissertation. I thank my external co-supervisor, Prof. Dr. Delphine Reinhardt, whose hospitality I have been enjoying during my secondment at the University of Göttingen.

I would like to thank my co-authors, co-workers and colleagues at and beyond Karlstad University for many insightful discussions.

I acknowledge the valuable feedback I received in my role as an early stage researcher of the Privacy&Us ITN, in the course of which my work was funded by the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant 675730.

Karlstad, August 2020 Patrick Murmann

List of appended papers

I. Patrick Murmann, Simone Fischer-Hübner. Tools for Achieving Usable Ex Post Transparency: A Survey. IEEE Access, 5:22965–22991, 2017.

II. Patrick Murmann. Usable Transparency for Enhancing Privacy in Mobile Health Apps. In Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct, MobileHCI 2018, pages 440–442, ACM, 2018.

III. Patrick Murmann, Delphine Reinhardt, Simone Fischer-Hübner. To Be, or Not to Be Notified—Eliciting Privacy Notification Preferences for Online mHealth Services. G. Dhillon et al. (editors), ICT Systems Security and Privacy Protection, pages 209–222. International Federation for Information Processing, IFIP SEC, Springer, 2019.

IV. Patrick Murmann, Matthias Beckerle, Simone Fischer-Hübner, Delphine Reinhardt. Reconciling the What, When and How of Privacy Notifications in mHealth. (Under submission).

V. Patrick Murmann. Eliciting Design Guidelines for Privacy Notifications in mHealth Environments. International Journal of Mobile Human Computer Interaction, 11(4):66–83, IGI Global, 2019.

VI. Patrick Murmann, Farzaneh Karegar. From Design Requirements to Effective Privacy Notifications: Empowering mHealth Users to Make Informed Decisions. (Under submission).

VII. Eva Schlehahn, Patrick Murmann, Farzaneh Karegar, Simone Fischer-Hübner. Opportunities and Challenges of Dynamic Consent in Commercial Big Data Analytics. M. Friedewald et al. (editors), Privacy and Identity Management. Data for Better Living: AI and Privacy, Proceedings of the IFIP Summer School, pages 29–44. Springer, 2020.

Comments on my participation

Paper I. I am the main author of the paper. I conducted the literature research, reviewed the papers, conceptualised the taxonomy for ex post TETs, and wrote up the manuscript. Simone Fischer-Hübner advised me on conducting the literature research, peer-screened the results obtained from the retrieval process, contributed the legal parts of the Background section, advised me on how to structure the gaps in the literature, and helped in writing the the paper.

Paper II. I am the sole author of the paper. Simone Fischer-Hübner provided ideas and reviewed the paper.

Paper III. I am the main author of the paper. I designed and conducted the online study, conducted the statistical analysis and wrote up the results. Delphine Reinhardt and Simone Fischer-Hübner advised me in conceptualising the study, helped in structuring the paper, and helped in writing the final manuscript. Simone Fischer-Hübner suggested to investigate privacy personas by segmenting users based on their preferences and predisposition.

Paper IV. I was responsible for planning and conducting the user study. I conducted the statistical analyses, collected and visualised the findings, and wrote major parts of the paper. Matthias Beckerle advised me on the statistical analyses, and contributed the machine learning techniques by means of which he confirmed many of the results obtained via the statistical analysis. He provided extensive advice on how to structure and refine the paper. Simone Fischer-Hübner helped compose the notification scenarios, peer-performed the translation of the survey from English into German, and provided supplementary input for multiple sections. She helped structure the manuscript and reviewed the article. Delphine Reinhardt provided guidance on writing the article and reviewed it.

Paper V. I am the sole author of the paper. Simone Fischer-Hübner provided valuable advice of how to structure the contents of this paper, discussed the model contained therein, and provided me with review comments.

Paper VI. I was responsible for eliciting the original set of requirements from the literature, and distilled the subsequent three iterations upon consultation. I implemented the prototype and applied the subsequent changes after each iteration of the evaluation. I moderated the user study. I wrote the introduction, background & related work, methodology (except for the section on study design), requirements, discussion and conclusion sections. Farzaneh Karegar provided feedback on each generation of the design requirements. She conceptualised the study design and served as a keeper of minutes during the user sessions. She provided the references on intervenability and wrote the study design subsection and results section of the paper. Simone Fischer-Hübner provided review comments during multiple iterations of writing the paper.

Paper VII. Eva Schlehahn, Rigo Wenning and Harald Zwingelberg are the original architects of the concept called dynamic consent. Eva Schlehahn provided advice on all legal matters throughout the entire project and provided feedback for each generation of the prototype. She moderated one of the focus groups and wrote the legal parts of the article. I evaluated the first three generations of mockups created for the project and designed the subsequent generations including the one used during the focus groups. I moderated one of the focus groups, analysed the results obtained in both groups, and structured the findings. I wrote major parts of the Methodology, Results and Discussion sections. Farzaneh Karegar helped create the two latest versions of

the prototype and served as a keeper of the minutes in one of the focus groups. She provided the section on related work. Simone Fischer-Hübner provided legal and organisational advice during the conceptualisation of the project, and reviewed multiple iterations of the prototype. She served as a minute taker in one of the focus groups, wrote the results of that group, and helped structure the overall results of both groups. She wrote the introduction and part of the discussion section.

Other publications

- Patrick Murmann, Simone Fischer-Hübner. Usable Transparency Enhancing Tools: A Literature Review. Technical report, Karlstad University, Department of Mathematics and Computer Science, 2017.

  Contribution: I am the main author of this paper. I conducted the retrieval of the literature and documented the results. Simone Fischer-Hübner advised me about conducting the literature research and peer-screened the publications retrieved in its course.

- Michael Bechinie (editor). User Interface Requirements V1.0. Privacy&Us deliverable D4.1., 2017.

  Contribution: I provided Section 2.1 of this paper.

- Ben Wagner (editor). Risk Assessment V1.0. Privacy&Us deliverable D5.2., 2017.

  Contribution: I participated in the workshop at the 3rd Privacy&Us training event in Tel Aviv, in the course of which the material for this paper was collected.

- Simone Fischer-Hübner, Leonardo A. Martucci (editors). User Interface Designs and Prototypes V1.0. Privacy&Us deliverable D4.2., 2018.

  Contribution: I provided Section 2 of this paper.

- Patrick Murmann. Towards Usable Transparency via Individualisation. Licentiate thesis, ISBN 978-91-7867-003-1 (print), ISBN 978-91-7867-008-6 (PDF), Karlstad University, 2019.

  Contribution: I am the sole author of this thesis.

- Simone Fischer-Hübner, Matthias Beckerle, Tobias Pulls, John Sören Pettersson, Patrick Murmann, Jonathan Magnusson. Project PAPAYA, Deliverable D3.4 – Transparent Privacy-Preserving Data Analytics, 2020.

  Contribution: I provided the major part of Sections 3.1–3.2 and the illustrations contained therein. I helped copy-editing the manuscript.

Contents

Glossary

Breaches. (Capitalised) One of the three categories of privacy notifications.

Consent. An informed and unambiguous affirmative statement, given freely and specifically, by which a data subject indicates her agreement to have her personal data processed (GDPR, Art. 4 (11)).

Consequences. Circumstances that arise for a data subject due to the processing of her personal data. Also (Capitalised): One of the three categories of privacy notifications.

Data concerning health. Personal data related to the physical or mental health of a natural person (GDPR, Art. 4 (15)).

Data controller. Entity that determines the purpose of personal data processing (GDPR, Art. 4 (7)).

Data processor. Entity that processes personal data on behalf of the data controller (GDPR, Art. 4 (8)).

Data subject. Role of an "Identified or identifiable natural person" with respect to the processing of her personal data (GDPR, Art. 4 (1)).

EHR. Electronic Health Record

Electronic health record. Such data are typically managed by the organisation that were responsible for their collection, such as clinics and hospitals. See also: *personal* health record.

Ex ante transparency. Transparency about personal data processing *before* personal data are disclosed.

Ex post transparency. Ditto *after* personal data have been disclosed.

GDPR. EU General Data Protection Regulation (Regulation (EU) 2016/679)

HCD. Human-Centred Design

HCI. Human-Computer Interaction

Human-centred design. Iterative and integrative design process that focuses on the needs of the final users of a system.

I(C)T. Information (and Communication) Technology

Information privacy. The ability to control the disclosure, processing and dissemination of one's personal data.

Intervenability. The legal right to intervene in the processing of one's personal data (GDPR, Art. 12 et seq.).

mHealth. Mobile health. We consider the non-clinical usage context of self-quantification and personal health tracking.

Mobile app. An application, native or otherwise, running on a → *mobile phone*.

Mobile phone. A customary smart phone capable of establishing data connections to the Internet and processing → *privacy notifications*.

Notification preferences. We consider three determinants to customise the delivery of → *privacy notifications*: (1) *whether* to notify, (2) *when* to notify (timing), and (3) *how* to notify (modality).

Persona. Refers to two independent concepts of personae: (1) The digital persona as described by Clarke (p. 7), and (2) the privacy persona used to describe the predisposition of a group of subjects in terms of their privacy (p. 8).

Personal data breach. A breach of security that leads to the accidental or deliberate destruction, loss, alteration, unauthorised disclosure of, or access to, personal data (GDPR, Art. 4 (12)/34).

Personal data. → *PII*

Personal health record. Such data are managed by the subject who collected them for the purpose of personal informatics. See also: *electronic* health record.

Personal health tracking. Measuring and processing one's behavioural, physiological or anecdotal personal data.

Personally identifiable information. Information that can be used to identify a data subject (GDPR, Art. 4 (1)).

PET. Privacy-Enhancing Technology

PHR. Personal Health Record

PII. Personally Identifiable Information

Privacy indicator. Any indicator pertaining to improving the privacy of a data subject, such as in the form of iconography or a pop-up message.

Privacy notice. Sometimes used as a synonym for either privacy *policy* or for privacy *indicator*.

Privacy notification. Customised notification sent by a TET to a data subject in response to a privacy incidence, or as an opportunity to improve the recipient's privacy.

Privacy policy. The legal manifest that specifies the actions, stakeholders, purposes and duration of how, by whom, why, and for how long personal data will be processed. Specifies also the contact details of the data controller.

Privacy preferences. Preferences expressed in terms of having one's personal data processed for a specific purpose (ISO/IEC 29100:2011 2.17).

Privacy-enhancing technology. ICT used to improve the privacy of an entity. A PET may or may not facilitate transparency (→ *TET*).

Privacy. → *Information privacy*

Processing (of personal data). Operational measures performed on personal data, including storage and transmission (GDPR, Art. 4 (2)).

Profiling. Automated processing of personal data, particularly to predict the respective data subject's behaviour, performance or status (GDPR, Art. 4 (4)).

Self-quantification. → *Personal health tracking*

Sensitive data. → *Special categories of data*

Special categories of data. Includes, among other types of data, 'data concerning health'. Processing special categories of data warrants extra care and protection (GDPR, Art. 9).

TET. Transparency-Enhancing Tool

Third party. Entity or stakeholder other than the data subject, the data controller or the data processor (GDPR, Art. 4 (10)).

Tips. (Capitalised) One of the three categories of privacy notifications.

Transparency-enhancing tool. ICT that enhances the transparency of how a data subject's personal data will be (ex ante), or have been (ex post) processed. A TET may provide additional functionality to help improve the user's privacy, e. g. provide recommendations or guide users in exercising their right of intervenability.

Transparency. Clarity conveyed about personal data processing, in a way that is accessible, intelligible and relevant for the addressee.

UCD. User-Centred Design

UI. User Interface

Usability. A quality or measure of how well a particular user can use a particular tool in a particular context to accomplish a particular task.

User-centred design. → *Human-centred design*

Terminology

In this thesis, we draw on the terminology established by normative sources as much as possible. In the case of conflicting or overlapping sources, we prioritise legislative sources over technical standards, which in turn overrule terms coined in scientific publications. E. g., we use the term 'data subject' as specified in the GDPR rather than 'PII principal' as specified by ISO/IEC 29100. Conversely, the term 'ex post transparency' is specified in neither of these sources, nor is it proposed in advisory frameworks, such as the guidelines released by the Article 29 Working Party, ENISA or OECD. Requiring a term that is both concise and well established, the term has therefore been adopted from the literature.

Introductory Summary

Preface

Information at your fingertips[1] served as a commercial tagline that is typically attributed to William Henry Gates III. During his notable keynote for the COMDEX in fall 1990 [41], the co-founder, ex chairman and ex CEO of Microsoft Corporation coined the term in the context of what he described as his vision for the next decade of ICT. Conceptually speaking, said vision built on the notion of integration 'in a seamless fashion' [41]. He explained that by having various types of information at their fingertips users of graphical, fully integrated information systems would accomplish an unprecedented level of productivity in domains such as business administration, personal information management, and education. Gates described the key technology for this achievement as media-rich, content-driven tools, the use of which would enable users to work more effectively, efficiently, and satisfyingly.

The past 30 years since 1990 have introduced the rise and fall of various technologies, architectures and paradigms. As contemporary witnesses we have seen companies giving rise to formidable accomplishments in terms of speed, capacity, and complexity. Far beyond what Gates called 'company-wide networks', we saw the recurring ebb and flow of self-contained online environments that promise pleasant user experience. These flourishing ecosystems have been described as walled-in dooryards manicured by service providers for economical purposes [92, 128], and seem to compete with the comparatively unkempt, libertarian Internet originally rooted in academia [92]. Gates' vision epitomises the omnipresent algorithmic intelligence deployed worldwide to optimise business processes and revenues, which has led to users of online environments being consumers and creators of information alike. We have learned that *personal data*,[2] once disclosed, can be the source of inference, turning users into *usees*.[3] We have also learned that both users and usees are often unaware of the underlying processes that affect them, nor that they typically know how to intervene in the process.

The reason for relying on the phrase 'information at your fingertips' as the title of this thesis is three-fold:

First, the original meaning of the idiom reflects the primary topic covered in this thesis: *Transparency*. At its core, transparency pertains to the availability and visibility of information. Hence, the idiom adequately paraphrases the core content of the thesis.

Second, Gates' keynote suggests that the type of information he was referring to related to suitable information in the sense that such information ought to *facilitate informed decision-making* for the one having access to it. The notion of suitability is tightly connected to *usability and customisation*, both of which reflect major aspects of this thesis. Suitability is also related to the aspect of *clarity and intelligibility* with respect to transparency, both of

[1] At one's fingertips: "(especially of information) readily available; accessible" [94, 112]

[2] Also: 'personally identifying information': "any information relating to an identified or identifiable natural person" (Art. 4 (1) GDPR [33])

[3] "A usee is a person about whom the IT system produces and processes data, and who has usually no control over this process." [39]

which transcend availability and transform mere data into information. Furthermore, having information at their fingertips enables users to work towards an overarching objective: To make informed decisions of how to improve their privacy. Hence, despite the thematic disparity between the objective of Gates' dogma and transparency-enhancing tools, this thesis attributes a similar function to the role of customised information as a key enabler for informed decision-making.

And third, the mobile ecosystem, which serves as the contextual basis for our investigation of *privacy notifications*, is inherently tied to the tactile interaction between user and machine. Tactile feel is not only required to actuate operations on touchscreen displays, but also allows users to receive, to feel information, such as via touch feedback and vibration signals. Tactile information can thus be literally at one's fingertips. Despite the fact that tactile perception is not the primary informatory sense investigated in this thesis, it underscores the contextual relationship between a user's highly sensitive tactile organs and a handheld device. The latter serves as the designated intermediary between user and information system, and therefore constitutes the *medium* employed for informed decision-making on the part of the user.

1 Introduction

Many users of online data services find themselves in a precarious situation when it comes to deciding what is best for their privacy based on the information they have at their disposal. On the one hand, the EU General Data Protection Regulation (GDPR) [33] mandates that data controllers shall provide data subjects with sufficient information about their practices regarding the processing of personal data to enable them to make informed decisions as how to manage their personal data. On the other hand, data controllers seek legal compliance by covering in their privacy policies a plethora of meticulous details, which often fill dozens of pages of printed text [72]. Readers of such policies may struggle not only because of the sheer extent of the policy [42], but also because of a lack of readability, comprehensibility and transparency of the underlying information [35, 73]. Privacy policies are written by lawyers and are primarily motivated by the business interests of the firms employing them. Their purpose is to achieve completeness with respect to legal compliance, and thus seek maximum safety for the data controller while providing sufficient leeway for innovative future scenarios related to the processing of personal data. Hence, privacy policies are often rich in legal or technological jargon, or contain a large amount of seemingly unstructured details [15]. The resulting policies seldom reflect the mental models of their readers, the majority of which are laypersons with average linguistic skills and technical understanding.

This dichotomy introduces a phenomenon that Nissenbaum [87] refers to as a 'transparency paradox', a dilemma that questions the efficacy of the principle of notice and choice as it might render the very purpose of data transparency null and void. The futility of data transparency potentially entails an imbalance of power in that a considerable amount of data subjects do not have the means necessary to satisfactorily scrutinise the terms and conditions with respect to deciding whether to use the service in question [108], much less use them as the basis for informed decision-making. As a result, such individuals might face consequences and implications they may have been unable to anticipate [15, 43]. Alternatively, data subjects may choose not to give their consent, which often constitutes the only alternative to accepting the terms. This effectively establishes a dichotomy in the form of 'take it or leave it' [87], the blatant lack of options excerting additional pressure on the decision-makers.

Nissenbaum's findings and the large majority of research on usable data transparency pertain to *ex ante* transparency. They relate to scenarios in which a potential future user of a data service tries to decide objectively a priori whether to consent to the terms and conditions of a data service, be it an online service or an app she is considering to install on her mobile device. However, the conceptual dilemma of non-feasibility applies equally to the context of *ex post* transparency in that a data subject's attempt to obtain retrospective transparency about how her personal data have been processed once they were disclosed may be thwarted equally likely. Information on how personal data have been processed may not be readily available, may not be presented in an

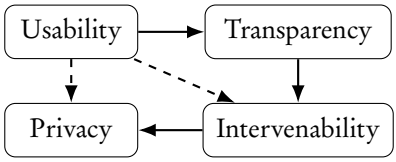

Figure 1: Usability enables transparency and intervenability, and thus privacy.

intelligible form, or may fail to point out clearly what options data subjects have in terms of intervening in the processing of their data. Hence, data subjects are effectively left without a means of helping them make informed decisions as regards managing the personal data they have disclosed.

Technological artefacts, transparency-enhancing tools (TETs) provide data subjects with the means necessary to access and sufficiently understand the circumstances and consequences related to the processing of their personal data [50]. Ideally, the TET in question would convey information such that content and form of the presentation reflect the mental model of the viewer, and match and complement her predisposition, previous knowledge and social context. In both ex ante and ex post transparency, the missing link to accomplishing meaningful transparency is often a lack of usability on the part of the TET, such as insufficient intelligibility or actionability of a privacy indicator. Despite legal requirements for intelligible transparency and a wealth of well-documented principles of usability engineering, only few actual guidelines and technical standards exist with respect to the design of usable TETs, especially ex post TETs. Whereas a considerable amount of research has been conducted on usable security and ex ante transparency, the number of ex post TETs discussed in the scientific literature is comparatively small [82]. Relatively few ex post TETs are available at present, and the number of usable TETs that have been designed specifically with the needs of their final users in mind is even lower (Paper I).

We argue that the principles underlying the usability of interactive systems equally apply to TETs, and that usability represents an indispensable prerequisite for an artefact that facilitates ex post transparency. Transparency, as a data protection goal stipulated in the GDPR (Art. 5 (a)), represents a cornerstone of information privacy. Hence, measures of usability, applied strigently and systematically, can help enable privacy (Figure 1).

Objective. The objective of this dissertation is therefore to systematically apply methods established for human-centred design with the goal of yielding a prototypical ex post TET that is demonstrably usable by its designated target audience. We seek to elicit the requirements necessary for the prototypical implementation of the tool, and validate them by means of user studies. We base the elicitation process on preliminary research of the status quo regarding usable TETs discussed in the literature. We classify existing TETs and point out gaps that indicate directions for our research. We specify personal health tracking and self-quantification as part of the domain of mobile health (mhealth) as the usage context for our design process, against the backdrop

of which we specify ex post transparency (Section 2.2) and intervenability (Section 2.3) as the goals we intend to achieve. We devise *privacy notifications* as a conceptual means to facilitate ex post transparency, and investigate the concept of customisation to capture the variety of preferences inherent in our target audience. As an exemplary precursor to future real-world applications, we present a prototypical TET that operates on the basis of privacy notifications. Its iterative evaluation motivates the applicability of a collection of design requirements on which its implementation is based.

The remainder of the introductory summary of this collection thesis is structured as follows: Section 2 provides an overview of the background of the subject matter, including a contextual arrangement of the concepts privacy, transparency, intervenability and usability, as well as contextual factors related to the usage context of mhealth. Section 3 derives four concrete research questions from our research objective. Section 4 elaborates on why human-centred design and various supplementary methods have been chosen for conducting our research. Section 5 discusses the contributions we make to the body of knowledge by addressing our research questions. Section 6 gives a summary of Papers I–VII. Section 7 provides an overview of work related to our contributions. Section 8 briefly points out limitations in our work, and Section 9 concludes the introductory summary with an outlook on possible future work.

2 Background

This section illustrates the backdrop of our research. We briefly discuss the interplay between the concepts of privacy, transparency, intervenability and usability (Sections 2.1–2.4). We anchor these concepts in the usage context of personal health tracking (Section 2.5), which defines the basis for addressing the aforementioned objective of this thesis.

2.1 Information privacy

As regards our contextualisation of the term 'privacy', we rely on the concept advocated by Warren and Brandeis in 1890 [125], and subsequently by Westin in 1967 [126]. Some of the time-honoured principles established by these authors carry over to the information age of the 21st century in that they conceptualise individuals in their roles as self-determined decision-makers whose personal choices not only affect their immediate surroundings, i. e. their intimate sphere of privacy, but also deal with reciprocal factors such as what intelligence other entities do or do not have about the individual in question [114]. They conceptualise privacy as a collective value, which is possessed and exercised by the individual in the form of a personal right [126]. Consequently, individuals have the freedom and power to leverage their right of privacy as a means of control to establish what kind of social interaction they prefer to cultivate. It is up to the individual to decide what kind of personal information she is

willing or not willing to share with the members of a particular social context, and by doing so, exercises control over the potential risks and harms to her privacy [17, 113].

In recent times, Nissenbaum conceptualises information privacy as the "right to control information about oneself" [87]. Conversely, Belotti [13] describes the facilitation of privacy as the operational approach to "enable people to determine [...] how to control [...]", referring to the individual's capability rather than to normative values. Irrespective of the underlying conceptualisation, exercising control implies the ability to define for oneself an online identity that manifests as the sum of activities conducted in a social context [49]. According to Clarke [23], the digital identity crystallises in the form of a digital persona, a reflection of the human personality or *alter ego*. It is comprised of the traces of personal data that have been disclosed deliberately and accidentally to the social context in question, such as an online social network. Ideally, data services should respect the individual's right of self-determination [18, 20], including the context-specific choice of individual self-portrayal.

Analogous to social contexts in the physical world, the conclusions onlookers are able to draw about the individual may or may not be congruent with her own subjective conception of her digital self [23]. The perception of the holistic picture of the individual's digital identity will be based on information available at different levels of concreteness, such as explicit and implicit data, and on data derived from and predicted about the data subject (see Paper I, Section 5.4.1, p. 81). The picture of the 'formal digital persona' will become more comprehensive once the observing party has access to personal data that identify the individual across the boundaries of multiple social contexts [23]. On a case-by-case basis, comprehensive perception of the data subject via traces left online in multiple, seemingly independent contexts may be in the individual's best interest. It may, however, also indicate that her attempt to control her digital identity has failed in terms of spilling over from one social context to another.

Revisiting the contemporary legal context, the current legislation likewise stipulates the concept of privacy as a legal right for informational contexts. Art. 12 of the Universal Declaration of Human Rights [30] protects human subjects against the interference with one's correspondence, honour or reputation, many of which are administered online today. For Europe and its residents, the Charter of Fundamental Rights of the European Union [34] mandates personal freedoms in the form of respect for, inter alia, private communication (Art. 7), and protection of personal data (Art. 8). Art. 8 establishes that the processing of one's personal data must be carried out on the basis of consent (or an other legitimate legal basis), and that "Everyone has the right of access to data which has been collected concerning him or her" (Charter, Art. 8 (2)).

More specifically, the protection of natural persons with regard to the processing of their personal data[4] falls under the legislative act of the EU

[4] "'processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, [...]" (GDPR Art. 4 (2))

General Data Protection Regulation 2016/679 [33], which has replaced the former Directive 95/46/EC since 25 May 2018. Hence, the legal principles drawn upon throughout this dissertation refer to the GDPR as a legal reference. These principles are stipulated in Art. 5 of the Regulation, and are partially congruent with the data protection goals specified in the Standard Data Protection Model [120]. The latter provides a method to consolidate the technological and organisational measures implemented by a data processor with the legal requirements pertaining to the processing of personal data. The data protection goals specified therein comprises the classic triad related to data security, availability, integrity and confidentiality, and complements them with goals related to information privacy: Unlinkability, transparency and intervenability all deal with protecting the data subjects themselves by applying legally compliant measures of processing personal data. The revised version of the Standard Data Protection Model [121] adds data minimisation to this list, which the initial generation treats as a principle rather than a protection goal.

In this dissertation, we focus primarily on the interplay between usability and transparency. However, since obtaining insight about how one's personal data are processed will not affect the processing in itself, we complement these two cornerstones with intervenability as a means to exercise one's legal right to bring about the change necessary to actually improve one's privacy.

2.2 Transparency

We conceptualise information transparency in the sense of *clarity of information* as described by Turilli and Floridi [119]. Turilli and Floridi refer to transparency as proven knowledge about data that have been disclosed previously. The property of *visibility* implies accessibility of the information relevant for the context in question, and is complemented by the properties of completeness and certainty. The purport of the term transparency differs from an alternative meaning that is sometimes used in the context of software engineering. In software architectures, the transparency of an interface or functional point of transfer is implementation-agnostic in that architects typically seek decoupled properties, i. e. independence from its underlying components. In this respect, transparency implies opaqueness or *in*visibility of the actual implementation.

This dissertation adopts the former view of transparency in that its designated goal is to aid data subjects in obtaining "an adequate level of clarity of the processes in privacy-relevant data processing" [45]. Hence, transparency poses an instrument that favours the interests of data subjects, and that imposes on data controllers obligations in terms of legal compliance. However, transparency may also serve as a mitigating factor against concern or even as a facilitator of trust (see e. g. [29, 68, 97]), and may therefore help create opportunities for data services to distinguish themselves from their competitors. The OECD Privacy Framework [91] therefore comprehends transparency, along with accountability, security, purpose limitation and accessibility, as an instrument of trust that helps establish confidence in an entity or process, e. g. with respect to that entity's reliability, rather than trust established on the

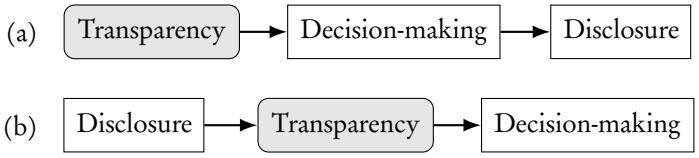

Figure 2: Conceptual phases of (a) ex ante, versus (b) ex post transparency.

basis of 'rules of thumb' [91].

The GDPR stipulates lawfulness, fairness and transparency as the three fundamental principles pertaining to personal data processing (Art. 5 (1a)). By submitting her declaration of consent, a data subject formally agrees to the processing of her personal data under the terms and measures specified in the data controller's privacy policy. The processing must be carried out accurately and adequately, and is limited to the retention period and purpose specified in the policy that served as the basis for giving consent (Art. 5 (1b–1e)). The data controller bears the full weight of demonstrating legal compliance of the measures carried out, and can be held accountable for deviations from the target specification (Art. 5 (2)). The term 'data processing' itself is a broad umbrella term that comprises a plethora of operational and organisational measures, including storing and transferring personal data to affiliated data processors, such as data centres or subcontractors whose collaboration enables the adequacy of the specified purpose. It also includes classification, profiling and automated decision-making (Art. 21–22).

The Article 29 Working Party [8] elaborates on the properties pertaining to information transparency, all of which are based on prescriptive facets specified in GDPR Art. 12. Information aimed to satisfy the demands of data subjects must be "concise, transparent, intelligible and easily accessible", require that "clear and plain language must be used", and "shall be provided in writing, or by other means, including where appropriate, by electronic means". The Working Party provides tangible examples of poor and good practice regarding the use of language, (in)appropriate wording, and recommendations regarding the measures by means of which respective information ought to be conveyed. The document [8] adopts a user-centric view, complementing and extending the few examples provided in GDPR Art. 12 et seq.

In terms of conceptualising decision-making as part of a data subject's cognitive process, we differentiate between ex ante transparency and ex post transparency [50], both of which are reflected in the legislation of the GDPR.

Ex ante transparency facilitates informed decision-making *before* personal data {will, may} be disclosed and processed.

Ex post transparency facilitates informed decision-making *after* personal data have been disclosed and processed.

Ex ante transparency refers to the notion of providing data subjects with the information necessary to make informed decisions about whether or not to use

a product or data service in the first place (GDPR, Art. 14). As such, ex ante transparency addresses decision-makers in their role as future data subjects, and is implemented in the form of a privacy policy issued by the data controller. The privacy policy specifies the details about how a subject's personal data will be processed in the future and serves as the legal reference for questions of purpose and scope.

Ideally, privacy policies present any details such that the information contained therein enable a data subject to unequivocally understand the measures pertaining to the processing of her personal data, and help her make an informed decision as to whether she gives her consent. Hence, privacy polices that aim to be transparent must be understandable and relevant for the data subject [108], enabling her to anticipate the consequences that will arise due to her personal data being processed. In this context, ex ante transparency relates to providing data subjects with transparency about how their personal data *will* or *may be processed in the future*, should they decide to use the product or service in question (Figure 2a).[5] Consequently, ex ante transparency carries the full weight in terms of providing an adequate clarification that sufficiently supports a data subject's decision of whether she gives her consent, and, by doing so, accepts the consequences of her choice.

Transparency goes hand in hand with choice. However more often than not, data subjects have little choice as to what personal data they choose to disclose for what purpose, but are instead obliged to decide whether to consent to disclosing a conglomerate of various personal data. This predicament hints at various conceptual and structural shortcomings of the notion of consent [19, 114]. Nissenbaum [87] describes the dichotomous nature of choice as a paradigm that is rooted in the micro-economical calculus of perceived costs. Data services are designed such that one's refusal to give consent, i. e. to refrain from participating in the social context in question, will result in personal costs that are, for the most part, perceived as unacceptably high. The OECD [91] builds on this notion of limited choice and encourages disclosure schemes that are more granular, and that allow data subjects to disclose only the bare minimum of data required to conduct the data service in question.

Conversely, *ex post transparency* describes the concept of providing data subjects with transparency about how their personal data *have been processed* (GDPR, Art. 15, 34). It applies to scenarios in which a data subject has decided to use a product or sign up for a data service, and thereupon has knowingly or unknowingly disclosed personal data. Ex post transparency enables users to subsequently obtain transparency about the processing of their data, and to hold data controllers accountable for any transgression in terms of deviating from the measures stipulated in the privacy policy. The insight thus obtained may be valuable in itself, but may be even more expedient in that respective knowledge may serve as a precursor for informed follow-up decisions, such as to intervene in the processing of one's personal data (Figure 2b). More generally speaking, suitable follow-up actions might also present themselves in

[5] See Figures 1 and 3 on pages 63 and 64, respectively, on how TETs serve as enabler technologies in scenarios of ex ante and ex post transparency.

the form of relatively simple operational measures, such as to change individual privacy settings related to the data service in question.

Ex post transparency mirrors ex ante transparency in that information provided to inform users about past data processing must be equally accessible and intelligible. However, whereas users of a data service rely on ex ante transparency exactly once, i. e. at the time of giving their consent, they may potentially benefit several times from the insight gained via ex post transparency in the course of using a data service. Hence, usable ex post transparency will have to be designed for cases of repeated use, and accommodate the needs of users whenever they choose to audit the processing of their personal data.

Throughout this dissertation, we draw on selected findings established for ex ante transparency. We also argue that in various scenarios, such as in the case of privacy notifications, ex ante and ex post transparency complement each other seamlessly because informed decision-making may equally depend on intelligence about the past as it depends on prospects of choices made in the present. At large, however, our research primarily deals with principles related to ex post transparency. We consider scenarios in which the user of a data service has decided to consent to the terms and conditions of a data controller, has effectively become a data subject, and is now experiencing the consequences of the processing of her personal data. Said user finds herself in a situation in which she relies on insight about the past to make decisions that will improve her privacy in the future.

Looking ahead towards matters pertaining to usability (Section 2.4), we conceptualise transparency as customised information provided by TETs such that the information conveyed is relevant and meaningful for the recipient. We conceptualise the presentation of such information as a means to enable data subjects to make informed decisions. Moreover, we complement ex post transparency with ex ante transparency to advise users about feasible means of how to intervene in the processing of their personal data by complementing retrospective insight with context-specific, actionable choices of how to transform the user's decision into actions.

2.3 Intervenability

Intervenability pertains to interfering "with the ongoing or planned data processing." [45], the objective being "the application of corrective measures and counterbalances where necessary" [45]. The GDPR mandates that data subjects can hold data controllers accountable for the processing of their personal data, and grants them data subject rights to

- Withdraw consent for the processing of their personal data (Art. 7 (3))
- Access to information as to what personal data are processed for what purpose (Art. 15)
- Have their personal data rectified (Art. 16)
- Have their personal data erased (Art. 17)

- Restrict the processing of their personal data (Art. 18)
- Have their data exported and/or transferred to another service provider (Art. 20)
- Object to the processing and profiling of personal data (Art. 21)
- Object to automated decision-making (Art. 22).

Ex post transparency, i. e. proven knowledge about past data processing, constitutes a causal prerequisite for exercising a data subject's right of intervenability [74]. Actions invoked in the course of intervenability are contextual in that many actions are applicable only in cases that warrant respective measures, such as exercising one's rights as stipulated in Art. 22 being applicable only if automated decision-making actually takes place.

In many contemporary scenarios, implementing measures to deal with intervenability is not straight forward. The complexity of the technological infrastructure deployed throughout information systems may require control mechanisms that are similarly sophisticated and costly as the ICT they are supposed to supervise [46]. The facilities involved in processing a data subject's personal data may be co-located or outsourced to subcontractors, implying that some of the data controllers and data processors may span across multiple legal entities, countries, or geographic economic areas. The means by which data subjects interact with data controllers are neither standardised nor automated. As regards contact points, GDPR Art. 13/14 mandate the specification of information about "the identity and the contact details of the data controller" (1a), and, where applicable, "the contact details of the data protection officer" (1b). Consequently, the organisational measures deemed necessary by data controllers to deal with data subject requests frequently present themselves in the form of 'clerk-operated help desk[s]' [46] rather than via actionable options that are seamlessly embedded into the user interface (UI) of the data service in question.

Exercising intervenability requires an explicit declaration of intent on the part of a data subject, i. e. the legal, administrative, and operational measures resulting from requests to intervene in the processing of her personal data do not come into effect without her prior action. TETs have practical limits regarding the extent to which they can actively support data subjects in exercising their data subject rights. This applies in particular to TETs that operate legally and technically independently from the data controller providing the data service.[6] Lacking means to relay data subject requests, third party tools will not be able to provide users with actionable options that offer immediate functionality to exercise a legal right. TETs may, however, provide their users with customised guidance on a case-by-case basis, such as by prioritising suitable options, or by clarifying subsequent legal or administrative measures required by data subjects in the course of taking an action.

[6] See Section 5.2.4 for conceptual and technical constraints related to privacy notifications.

2.4 Usability

In its most basic form, usability refers to "the quality or state of being usable" [75]. In a quantitative context, it pertains to "the degree to which something is able or fit to be used" [94]. Both definitions create an association between an implement or tool and a user, and assign a measure of appropriateness and correspondence to their relationship. The measure relies on a clear specification of the characteristics of both user and tool, and of the context in which the user relies on the object for a particular purpose. The ISO Guidance on ergonomics of human-system interaction (ISO 9241-11) [53] subdivides usability into the three distinctive measures of effectiveness, efficiency and satisfaction. All three measures build on the notion of a specific kind of user using a specific set of equipment to accomplish a specific task to achieve a specific goal or outcome. The interplay of the aforementioned factors is conducted in a specific context of use that forms the backdrop against which the measures apply. Hence, the notion of usability appears as a combination of multiple concomitant factors and is incomplete unless each of them is specified sufficiently.

Usability can be conceptualised as a measure of ergonomics with respect to human-computer interaction. ISO 9241-110 [55] stipulates seven dialogue principles[7] for the design of UIs that facilitate the interaction between a human subject and an interactive system. These principles are akin to and partially congruent with the golden design rules established by Shneiderman [111] and the usability heuristics established by Nielsen [85], the latter of which provide practitioners and evaluators with heuristics for inspecting the usability of interactive systems. All three sources build on the aforementioned notion of usability depending on the specification of a specific task conducted by a specific group of users. ISO 9241-110 complements these requirements with the concept of user needs, i. e. specific demands shared by the representatives of the intended target audience. It stipulates that said users have specific needs that must be respected at the time the system is designed such that the resulting product helps them accomplish their goal effectively, efficiently and satisfyingly.

The GDPR mandates usability mainly with regard to the concept of transparency, i. e. in the form of comprehensible information pertaining to the processing of personal data provided to data subjects by data controllers. Respective requirements apply to privacy policies and information on how a data subject's personal data are processed (Art. 12 et seq.), but likewise apply to personal data breach notifications (Art. 34) issued by a data controller. Furthermore, Art. 12 (7) suggests that information provided as part of a privacy policy or an audit of personal data processing "may be provided in combination with standardised icons in order to give in an easily visible, intelligible and clearly legible manner a meaningful overview of the intended processing. Where the icons are presented electronically they shall be machine-readable." To this day, however, most research conducted on privacy icons focused on ex ante transparency [51, 98]. The attempt to establish standardised iconography that

[7] Ttranscript provided in Section 2.2 on p. 65.

allows for codifying a sufficiently large variety of real-world scenarios related to personal data processing is still ongoing [105, 106].

The guidelines on transparency provided by the Art. 29 Working Party [8] interpret the high-level requirements stipulated by the GDPR, providing tangible examples in terms of how to implement usable measures of presenting information. Along similar lines, Patrick and Kenny [96] propose four design principles for UIs of information systems: Comprehension, Consciousness, Control, and Consent. These principles focus on the data subject in her role as a user of a data service that deploys an UI designed for the specific purpose of guiding decision-making related to personal data processing. The principles mainly pertain to ex ante transparency in that Comprehension and Consciousness essentially reflect aspects of accessibility and intelligibility, which eventually lead to Consent. However, the principles apply equally to ex post transparency in that Control constitutes corrective measures, which could potentially enable data subjects to intervene in the processing of their personal data. In this respect, Patrick and Kenny's principles may serve as a basis for establishing more refined design principles for TETs. We revisit the four principles in Section 5.2.5 when we discuss established models that motivated the conceptualisation of privacy notifications.

Revisiting the specification provided by ISO 9141-11 [53], we ultimately consider usable transparency against the backdrop of contextualisation. The fact that the ergonomic characteristics of an artefact rely on the specification of a particular group of users who seek to accomplish a particular goal under particular boundary conditions necessitates a distinctive context of use [67]. Contextualisation is required not only for measuring and evaluating the usability of a product, but also for engineering the requirements of an implementation that aims to be usable.[8] The context considered in this dissertation is personal health tracking as a subdomain of mobile health (mhealth) [63]. We consider users of online mhealth services whose primary task is to track and monitor their health, and whose secondary task is to improve their privacy in the course of tracking their health.

2.5 Transparency in personal health tracking

We consider the scope of personal health tracking to comprise personal informatics and self-quantification, and consider scenarios in which users of mhealth devices disclose their personal data to online mhealth services. We therefore limit the scope to personal data processed in the form of personal health records (PHR), i. e. data managed by the individual who collected them [117]. Conceptually, such data differ from electronic health records (EHR), which are typically managed by the institution that collected them, such as a hospital or clinic [9]. Pursuant to Art. 4 (1) of the GDPR, both types of data qualify as personal data. Furthermore, personal data related to the health and well-being of a human being qualify as special categories of data (Art. 9). Processing such

[8] See Section 4.1 on human-centred design for the sequential steps necessary to facilitate a design process that can be considered to be sustainably usable.

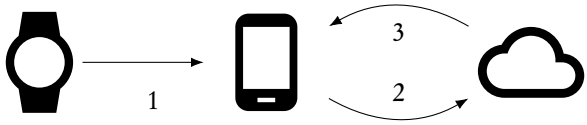

Figure 3: Data flows of user data in personal health tracking: (1) Sensor data, (2) storage, and (3) retrieval/review.

data is generally prohibited, and is subject to restriction and special protection.

The business sector of mhealth has been growing steadily over the last couple of years and is predicted to continue doing so in the foreseeable future [115, 116]. Millions of users of mhealth devices monitor their health and disclose their data to online mhealth services. In personal health tracking, the individual reasons for self-quantification are as diverse as the plethora of sensor data collected and processed by contemporary tracking devices [130]. The motivation to collect and process physiological, physical and anecdotal data about oneself ranges from tracking one's activities [76], tracking one's health and sleeping behaviour [64], comparison and competition [3, 44], motivation to exercise [11], or stimulating behavioural change [48, 104] to lifelogging [84, 93] or serve as the basis for research [52, 103]. The data thus collected are highly contextual in that they comprise meta data, such as time stamps, and in that reconciling such data with location data allows for pinpointing the user's activity and condition to geographical locations. The latter feature allows for obtaining insight about the cooperation and interaction of multiple individuals in terms of when they were spending time in the same place, and enables observers to draw conclusions about the activities they have been conducting [65, 66].

Technology is readily available as a myriad of ubiquitous mhealth devices such as fitness bracelets, smart watches, breast belts or headbands [4]. To date, broadband technologies such as 5G have not provided comprehensive coverage of a wireless infrastructure for stationary and mobile communication units [90]. We therefore consider scenarios in which mhealth wearables do not communicate directly with online mhealth services, but rely on a relaying intermediary, typically in the form of a mobile phone with cellular mobile telephony capabilities (Figure 3, step 1). Optionally, the onboard sensors of the device used as an intermediary may provide additional sensing data, such as location data [59], if the wearable does not collect such data itself. Moreover, users can curate their mhealth data with situational annotations to contextualise the traces of their physiological well-being. Ultimately, the mobile phone serves as a central information hub in that it not only relays the user data to an online service (Figure 3, step 2), but is also used to monitor longitudinal statistics of the data collected previously (Figure 3, step 3).

Smartphones have been serving for the purpose of personal information management (PIM) ever since they replaced personal digital assistants (PDA), which were typically without cellular connection and more limited in terms of their sensory capabilities [21]. We focus on scenarios in which users employ smartphones as data relays and disregard cases in which this role is facilitated

via laptops or desktop computers. Complementary to serving as a means for reviewing and visualising user data, the smartphone also serves as the target platform for the TET conceptualised for the usage context at hand. Due to the fact that all personal data pertaining to personal health tracking that are disclosed to the online mhealth service pass through the phone, and given the omnipresence of smartphones in today's socio-culture, that same device serves as an adequate host for a TET whose purpose is to notify users about noteworthy events at a moment's notice. Hence, all functional and non-functional requirements related to the design of this TET reflect boundary conditions imposed by a context related to mobile computing rather than a stationary application context, such as a home or office environment.

The primary goal users seek to accomplish when they operate in this scenario is to set up and maintain the devices necessary to facilitate accurate and seamless self-monitoring. We presume that obtaining transparency constitutes a secondary goal for these users, but that the prospect of improving their privacy is at least worthy of their consideration [2]. The task associated with this goal consists of retrieving and reviewing information about how their personal data have been processed. We further assume that users prefer to spend as little time as possible on privacy or security-related tasks [99], and to experience as little friction and disruption as possible when switching between their primary and secondary task [28].

We reason that users of mhealth services will not, on a regular basis, engage in proactive measures to seek out and review information about how their mhealth data have been processed [123]. We further reason that, if this group of users was to be aided in improving their privacy, respective information will instead have to be delivered to them automatically, and that the modalities leveraged in the course of this process will have to blend in smoothly with the mobile ecosystem they already employ. Based on this rationale, we therefore argue that privacy notifications can serve as a suitable means to fulfil the purpose of providing users with customised insight about matters related to their information privacy.

3 Research objective

The objective of this thesis is to contribute to the design of usable TETs that demonstrably accommodate the needs of their designated target audience by stringently applying principles of human-centred design. We subdivide this overarching goal into four distinctive research questions:

RQ1: Classification of usable TETs. What are the characteristics and gaps of usable ex post TETs published in the literature?

RQ2: Modelling of notification preferences. How can notification preferences be structured based on the underlying scenarios?

RQ3: Design of privacy notifications. How can privacy notifications be designed to facilitate transparency and informed decision-making?

Table 1: Mapping Papers I–VII to research questions RQ1–RQ4. Key: '•': Paper addresses RQ fully. '○': Paper addresses RQ partially.

Research question	Paper I	II	III	IV	V	VI	VII
RQ1	•	—	—	—	○	—	—
RQ2	—	○	•	•	—	○	—
RQ3	—	—	—	—	•	•	○
RQ4	—	—	○	•	○	•	○

RQ4: Efficacy of privacy notifications. To what extent can privacy notifications complement proactive measures on the part of data subjects to investigate how their personal data have been or will be processed?

RQ1 seeks to ascertain the state of the art of usable ex post TETs published in the literature. We aim to classify TETs based on conceptual, functional and technological characteristics, as well as aspects related to their usability. Based on gaps detected in the literature regarding their compliance with legal requirements and established principles of usability, we choose customisation as a means to accommodate the needs of users as the principal direction for our work.

Pursuing the notion of customisation, RQ2 drives our effort to discern *what kind of facts users prefer to be notified about*, which calls for structuring privacy notifications based on their content. We introduce a three-fold classification that subdivides scenarios into Breaches, Consequences and Tips. Orthogonally to the informational structure, we conceptualise notification preferences as a measure of *how users prefer to be notified*, which introduces timing and modality. We investigate to what extent user preferences can be predicted by relying on determinants such as a user's predisposition or usage behaviour to provide her with suitable settings for privacy notifications.

In the course of the user studies conducted to address RQ2, we learn that the right to intervene in the processing of one's personal data has an effect on our respondents' request for notification. We therefore introduce user control, corrective measures, and actionable choices as designated requirements for designing privacy notifications.

To address RQ3, we investigate *how privacy notifications ought to be designed* to facilitate transparency. RQ3 seeks answers as to how facts conveyed as part of privacy notifications need to be presented to accommodate the needs of recipients. We scrutinise the particularities and constraints imposed by the ecosystem of mobile phones, and present our findings in the form of concrete design guidelines. We investigate the scope and organisational structure of the design space to infer tangible requirements of how intelligible facts need to be curated. Moreover, we evaluate how privacy notifications can enable informed decision-making and advise taking action in response to the insight thus obtained.

RQ4 seeks answers to the question of how privacy notifications, irrespective of the technical factors addressed in RQ2 and RQ3, cater to the needs of the target audience in terms of usefulness and utility. More specifically, we seek to ascertain to what extent privacy notifications can complement the traditional way of actively enquiring information related to the processing of one's personal data on one's own.

Table 1 shows to what extent Papers I–VII address each of the research questions. Full coverage (•) denotes that it is the expressed purpose of the paper to address a research question, whereas partial coverage (◦) means that the paper touches on individual facets of a research question. Due to the considerable scope of all four research questions, more than one paper addresses each of them from a slightly different angle.

4 Methodology

This section discusses methods and techniques we utilised in the course of conducting our research. It covers human-centred design as a strategic approach to organise and structure our endeavour long-term (Section 4.1), and complements it with various auxiliary research methods chosen to accomplish specific sub goals (Section 4.2). Moreover, the human-centric nature of our research made ethical approval a prerequisite rather than an option, which we briefly touch upon in Section 4.3.

4.1 Human-centred design

Drawing on the notion of usability as the quality of a tool or system used by a particular group of users to accomplish a particular task, the research presented throughout this dissertation has been conducted against the backdrop of human-centred design. Advocated by Norman in 1986 [88], *user*-centred design (UCD) pertains to an iterative design process that seeks to accommodate the needs and requirements of the final users as a central driving force.[9] UCD is integrative in that the final users of a product are not only considered normative stakeholders, but in that individual representatives of the target audience actively contribute to its design, e. g. when design requirements are specified or when prototypes are evaluated. UCD considers the characteristics inherent in the interplay of the situational and environmental factors that affect the use of the final product, such as working conditions and particularities of the task that users try to accomplish.

UCD has motivated the creation of standardised procedures related to the design of interactive systems. In this thesis, we conduct the methodology laid down in ISO 9241-210 [56], a successor to the discontinued ISO 13407 standard [54]. ISO 9241-210 specifies the life cycle for *human*-centred design for interactive systems (HCD) as an iterative, stepwise process that facilitates

[9] User-centred design is sometimes referred to as user-centred development to emphasise the fact that the development of a product is driven by the needs, tasks and goals of a user.

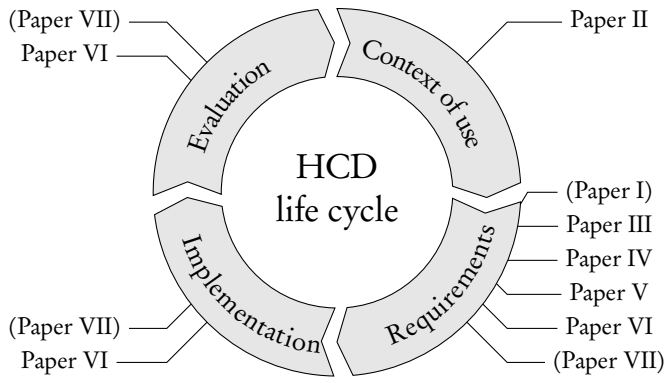

Figure 4: Development life cycle for human-centred design for interactive systems according to ISO 9241-210 [56]. The labels on both sides of the cycle indicate what paper contributes to each of the four phases.

the design of a usable interactive system (Figure 4). The life cycle models the principles of UCD for the context of interactive systems, integrating the methodological approach into the ISO 9241-series for the ergonomics of human-system interaction. Generally speaking, HCD is employed as part of a design process when a design team realises that demonstrable usability of the intended product depends on input and feedback from the actual users of a product rather than to solely rely on the opinions of independent domain experts [89]. The reason for integrating representatives of the target audience into the design process may e. g. be to improve the functional or operational qualities of the product, to make it easier to understand or operate by its users, to accommodate a larger variety of capabilities or disabilities exhibited by its users, or to improve the user experience by mitigating concomitant circumstances, such as discomfort or stress [56].

The life cycle of HCD commences at the 12 o'clock position with the specification of a definitive usage context, and traverses the subsequent phases in clockwise order. Requirements serve as target values for the subsequent implementation and evaluation, and are specified such that they reflect the needs of the target audience in terms of accomplishing their respective goals in the specified context. The purpose of evaluating the implementation is to ascertain whether the implementation suffers from any deficiencies that prevent users from accomplishing their goals effectively, efficiently and satisfyingly. In the case of such shortcomings, the respective phase would be revisited, the faulty facet revised, and the subsequent phases reiterated [56].

In the context of TETs, UIs convey manifold forms of information that are both delicate and sophisticated. Designers and domain experts recruited from the field of information privacy can predict the needs and expectations of users only up to a certain extent. Hence, the target audience of the TET may provide complementary insight about their actual needs, such as to help designers decide what kind of facts are necessary to facilitate transparency and

what level of explanatory detail will be required to understand the information presented to them [1]. Our decision to rely on HCD surfaced at the stage when we reviewed the literature available on usable TETs [12]. At that point in time, we realised that many TETs were not strictly designed with their users' needs in mind (Paper I), which suggested a more stringent approach that focused on user needs, if actual usability was to be accomplished. Our participatory design harnessed online surveys (Section 4.2.2), a lab study (Section 4.2.3), and focus groups (Section 4.2.5) to draw on feedback and opinions of various groups of stakeholders, such as actual users and domain experts.

In this regard, HCD has turned out to be an effective strategy for pursuing our research. It enabled us to draw on a well-established framework to plan our strategy long-term. In terms of coverage, the contents related to specifying the design requirements occupy by far the largest part of our work. The reason for this seemingly disproportional allocation is that the combination of context and approach constitutes a niche of original research that required considerable time and effort in terms of specifying the foundation on which the subsequent implementation and evaluation were to be built. Due to the scientific rigour required for carrying out each phase, we managed to complete only a single cycle of the process. Ideally, further iterations would have followed to revise the requirements and design based on the findings obtained during the evaluation. The cycle would have been re-traversed until we reach a state in which the evaluation satisfactorily attests that the design requirements have been met [56]. However, we implemented iterative design on a smaller scale, such as during the implementation of specific projects related to our research (Papers VI and VII), which is described in Sections 4.2.3 and 4.2.5, respectively.

4.2 Auxiliary methods

Under the umbrella of human-centred design, we employed various auxiliary methods to conduct specific parts of our research.

4.2.1 Literature research

Reviewing the literature provided us with a set of publications related to specific fields of research to address RQ1 and, to a minor degree, RQ2 and RQ3. By adapting the scope and screening criteria of the review process, we achieved a satisfactory trade-off in terms of breadth in terms of completeness, and the amount of papers that could realistically be covered during the review. The second factor that affected the number of papers we had to evaluate was the amount of scientific rigour applied in the selection and screening process, which in turn depended on the purpose of why we conducted the literature research in the first place.

The purpose of the literature research conducted in 2016/2017 was to ascertain the state of the art of usable ex post TETs [82]. It was crucial to implement a reproducible process, the results of which served as the baseline for the subsequent structural analysis conducted in Paper I. Methodologically, the process was systematic but non-exhaustive. It was systematic in that it adhered

to a stringent and documented retrieval process that yielded reproducible results. It was non-exhaustive in that it covered only a limited set of online databases and publishers, and in that it traced the bibliographic references only one generation forward and backward.

Conversely, the literature research conducted in December 2018 served the purpose of providing an overview of two orthogonal fields of the literature that we were investigating at the time. The review was carried out under strict temporal constraints and did therefore not exhibit the same level of rigour and depth as the one conducted in 2016/2017. However, the material retrieved in the process produced a sufficiently large set of publications that enabled us to discuss the interplay of two branches of the literature covered in Paper V.

4.2.2 Online survey

The purpose of our online surveys was to collect quantitative data from users of fitness trackers and mhealth devices to ascertain determinants related to RQ2 and RQ4. Compared to our lab study (Section 4.2.3), this method enabled us to collect a large amount of data in a short period of time. Using online crowd-sourcing platforms for recruiting our participants enabled us to specify screening criteria when we advertised our surveys, which we assigned in terms of age ($\geq$ 18 years), usage context (users of fitness trackers), and the country of residence (within Europe) of our participants in spe. These criteria provided us with test subjects who constituted the intended target audience of a prospective TET, i. e. users with hands-on experience with mhealth devices who were actually affected by the legislation of the GDPR.

Collecting the data in the form of quantitative choices meant that we did not need to transcribe, interpret and code qualitative results, such as verbal statements collected during an interview. The data thus collected could be used for the subsequent statistical analysis conducted for Papers III and IV with little effort in terms of screening and pre-processing. We explicitly refrained from collecting opinions via free-text fields for two reasons: First, textual data would have required interpretation. Since each of our online studies covered 300 participants, this would have significantly slowed down the analysis of the data. Second, free-text fields would have complicated the pseudonymisation and ethical approval (Section 4.3) because participants could have reported their well-being and ailments (GDPR Art. 9), or could have mentioned their name or the names of other people.

4.2.3 Lab study

Conversely, the purpose of our lab study was to collect qualitative in-depth feedback from a relatively small amount of participants to address RQ3. Employing methods related to participatory design and evaluative research, we sought our participants' opinions on individual facets of a prototypical implementation (Section 4.2.4). Investigating the efficacy of the prototype at a relatively early stage, the study was not conceptualised as a usability test,

i. e. we did not measure how effectively, efficiently and satisfactorily our test subjects accomplished a set of predefined tasks.

We did, however, draw on the fact that even a small number of evaluators help identify a considerable proportion of the issues inherent in a prototype [86]. This motivated us to expose three successive iterations of five (±1) participants to a rapidly evolving prototype, each iteration of which was refined according to the feedback obtained from the previous iteration. This objective called for a calm and controlled environment in which the test subjects could freely interact with the prototype, externalise their thoughts and opinions, and respond to our questions. By watching our participants interact with preselected facets of the prototype created for Paper VI, we obtained insight on what elements and features caught their attention or incurred their displeasure, and what concepts needed further clarification or refinement to achieve the intended purpose. We could not have achieved the same level of in-depth reflection of particular features with a quantitative survey, nor with an alternative qualitative creativity method such as a focus group (Section 4.2.5).

4.2.4 Prototypical implementation

The goal for implementing the prototype for our lab study conducted for Paper VI (Section 4.2.3) was to find an acceptable trade-off between making the artefact sufficiently usable to enable test subjects to engage in and relate to its functionality, while allowing us to make rapid changes once we had received feedback from all participants from one iteration of the evaluation. Stimulating users to become invested called for a fully functional prototype that felt and acted like a real-world application running on a mobile phone, whereas rapid prototyping required a simple, straight forward design.

We struck this balance by scripting the interaction with the prototype in that we conceptualised each interaction phase as a succinct, self-contained scene during which we probed our participants' experience. This enabled us to design each phase as responsive hi-fi UIs, while blanking out the transitions between individual phases. It also enabled us to quickly refine the design space of each phase depending on the feedback we received. We would not have been able to elicit the same amount of feedback either with a fully functional prototype, or with a non-functional lo-fi mock-up such as the one we used during the focus groups (Section 4.2.5).

4.2.5 Focus groups

Our focus groups served the purpose of having groups of domain experts, i. e. junior privacy researchers, engage in in-depth discussions about the concepts and prototypical implementations we confronted them with, thereby addressing RQ3 and, to some extent, RQ4. Similar to the lab study described in Section 4.2.3, this provided us with feedback on a lo-fi wireframe mock-up that served as a tangible reference for the thematic concept we were trying to explore for Paper VII. Consequently, the study took a hybrid form between conventional focus groups and a cognitive walkthrough.

Table 2: Mapping contributions to research questions RQ1–RQ4. Key: '•': Contribution addresses RQ fully. '○': Contribution addresses RQ partially.

Contribution	RQ1	RQ2	RQ3	RQ4
Conceptualisation of privacy notifications	—	•	—	•
Elicitation of design requirements for TETs	—	•	•	—
Literature survey and identification of gaps	•	○	○	—
Evaluated prototypical designs	—	—	•	•

Unlike the discourses we had with individual participants of our lab study, the focus groups benefitted from the expertise of multiple disciplines gathered around the table, which enabled us to examine our topic from technical, legal, and socio-economical angles. We experienced lively debates that were either triggered by concrete questions asked by the moderator, or in that participants picked up on the statements made by their peers. The hybrid form of the study enabled us to collect a plethora of feedback related to the topic and prototype with relatively little effort, resources and time. We could not have captured the same diversity of opinions within the same amount of time via alternative methods.

4.3 Ethical approval

We applied for ethical approval for all qualitative and quantitative user studies we conducted. Depending on the author's base of operations at the time, ethical approval was either obtained from the University of Göttingen (Paper III, reused for Paper IV), from Karlstad University (Paper VI), or from Unabhängiges Landeszentrum für Datenschutz Schleswig-Holstein (Paper VII). We were not involved in the process of applying for the latter. The formalities regarding the application process required for the two educational institutions differed greatly in terms of how the prospective study design, expected impact of the research, and the potential effect on the participants of the study had to be specified.

5 Contributions

We briefly summarise our contributions in Section 5.1. Due to the central role that privacy notifications play in this dissertation, we encapsulate the most important aspects related to the *conceptualisation of privacy notifications* in a dedicated subsection in Section 5.2.

5.1 Summary of contributions

This dissertation advances the body of knowledge by making the following main contribution to address our four research questions (Table 2):

- Conceptualisation of privacy notifications.

In addition to this main contribution, we make the following supplementary contributions:

- Elicitation of design requirements for TETs.
- Literature survey and identification of gaps.
- Evaluated prototypical designs.

Conceptualisation of privacy notifications

We present a stringent conceptualisation of the notion of 'privacy notifications'. Our work commences by deriving the concept from the literature (see *Literature survey and identification of gaps* below) in Papers II and V. In Paper V, we demarcate the unique features of privacy notifications from conventional concepts pertaining to privacy indicators, and contextualise privacy notifications in the application context of mhealth and personal fitness tracking. Moreover, we model TETs that harnesses privacy notifications based on factors that arise due to selecting mobile phones as our target platform.

We investigate two sets of orthogonal notification preferences in terms of customising privacy notifications (Papers III–VI) to address RQ2. The first type of preferences specifies what facts users prefer to be notified about (scenarios), whereas the other group of preferences determines when (timing) and how (modality) respective messages ought to be delivered. We further ascertain the efficacy of privacy notifications in terms of perceived usefulness of the concept, especially in contrast to enquiring respective information oneself, which addresses RQ4. We present preliminary results of how notification preferences can be predicted based on a user's demographics, usage characteristics, and predisposition to support customisation. Future work might investigate further determinants of notification preferences, and explore how adaptation may support customisation by adjusting the run-time behaviour of a TET depending on a user's choices.

Elicitation of design requirements for TETs

Harmonising legal requirements with design principles, heuristics and good practices, and consolidating them with findings established in the literature (see *Literature survey and identification of gaps* below), we deduce a set of design requirements for implementing TETs that operate on the basis of privacy notifications. Commencing as a list coarse-grained principles used to identify gaps (Paper I), the requirements gradually evolve into more granular design guidelines (Paper V), and finally mature into a set of requirements (Paper VI).

A select subset of the requirements serves as the basis for implementing two prototypes (see *Evaluated prototypical designs* below) in Papers VI and VII to address RQ3. By evaluating one of the prototypes in terms of examining and confirming the properties constituted in the requirements, we verify the

efficacy of the underlying requirements themselves (Paper VI). This addresses RQ2 in that our design requirements are indicative of the general functioning and concepts of privacy notifications. Future work might pick up where we left off and continue evaluating the remaining set of requirements in appropriate usage contexts.

Literature survey and identification of gaps

We compare the existing body of knowledge with legal requirements and design principles, and infer from the incongruity gaps in the literature (Paper I). These gaps motivate the direction and type of our research.

Reviewing the literature is conducted in two major steps, both of which are discussed in Section 4.2.1. The first step of reviewing the literature yields a taxonomy of usable ex post TETs (Paper I). This classification system allows for categorising publications in terms of conceptual, technical and usability-related characteristics, which addresses RQ1. Complementarily, we categorise publications based on the usage context of the TET in question. As a result of analysing the gaps elicited in the process, we orient ourselves towards achieving usability of TETs by means of customisation.

Addressing RQ3, the purpose of the second review is to harmonise the findings published in two branches of the literature – privacy indicators and the delivery of notifications on smartphones (Paper V). From the outcome we deduce possible directions for our research to fill respective gaps. As a result, we introduce and refine the concept of privacy notifications as a means to facilitate transparency in mhealth environments (see *Conceptualisation of privacy notifications* above), which constitutes the basis of RQ2.

Evaluated prototypical designs

The prototypical implementations we produce in the course of conducting our research for Papers VI and VII represent distillations of principles and concrete design requirements (see *Elicitation of design requirements for TETs* above) into tangible artefacts. Evaluating our implementations confirms, in whole or in part, the efficacy of the prototypes in terms of supporting users in accomplishing specific goals.

Our two prototypical designs serve similar purposes: To notify data subjects about facts related to their privacy. Despite the similarity of both designs in terms of drawing on established principles, the design of each of them is highly context-specific, which calls for individual design and evaluation processes. Our research motivates and documents these processes and provides valuable indicators of how to address RQ3 and RQ4.

5.2 Conceptualisation of privacy notifications

An early version of the concept was introduced when we decided on mhealth as the application context to investigate usable transparency in Paper II, and has since evolved as the central theme of our research. This warrants a dedicated

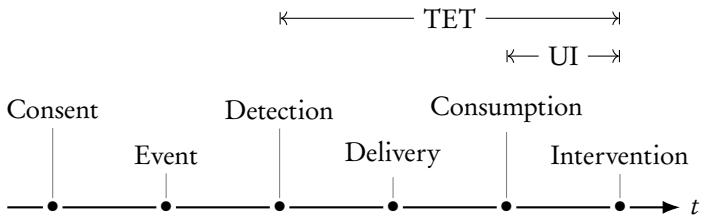

Figure 5: Steps pertaining to the life cycle of privacy notifications, including the scope covered by TETs and the UI they provide.

subsection on the major contribution of this thesis, which distils the findings obtained throughout Papers II–VII in a single place of the introductory summary. Starting with the basic principles of notification as a means to raise awareness (Section 5.2.1), we discuss various implications and constraints (Section 5.2.2–5.2.4). These considerations lend themselves to modelling a process that aims to harness privacy notifications to facilitate transparency (Section 5.2.5) and illustrates the unique character of the concept (Section 5.2.6).

5.2.1 Purpose and scope

In common parlance, notifications give (formal) notice of a fact, such as the occurrence of an event [75, 94]. The term 'privacy notification' has been used in multiple contexts (see e. g. [10, 57, 61, 77]). However to this day, no formal definition exists of the scope and exact definition of the term. A demarcation of how privacy notifications relate to, and differ from other forms of privacy indicators is provided in Section 5.2.6.

Receiving notifications constitutes an established paradigm of raising situational awareness of changes of circumstance or state on mobile devices [71]. The act of notification relies on passing personal messages without expecting anticipation on the part of the recipient regarding the arrival of the message. Envisioning a technological ecosystem in which transparency and intervenability go hand in hand with the use of a mobile device, we therefore draw on the notion of privacy notifications.

In the context of this thesis, we specify the purpose of privacy notifications as a means to raise awareness about noteworthy conditions related to a data subject's privacy by informing her about how her personal data have been processed (ex post), or about how the data may be processed in the future (ex ante). Moreover, privacy notifications provide recipients with customised recommendations of how to act to improve their privacy based on the information conveyed to them. To achieve this objective, the information provided as part of a notification must be conveyed such that it facilitates transparency about the underlying facts to enable informed decision-making on the part of the recipient.

We distinguish between multiple phases related to the operational life cycle

of privacy notifications. The phases shown in Figure 5 are not equidistant in that the time that passes between each of the steps can vary significantly or may even be non-deterministic. They constitute transitory phases that are temporally ordered and inter-dependent.

The designated scope of this thesis is limited to matters related to HCI. We therefore take a highly abstract stance regarding practical implementation details, as well as legal and administrative processes. This means that while strictly speaking the first two phases are necessary to support the rationale established here and to illustrate the functionality of privacy notifications, they are beyond the scope of this thesis. The first step, consent, is required to warrant the lawful processing of the user's personal data (GDPR, Art. 6), and lays the basis for designing retrospective means of ex post transparency. For the second step, we imply that at some point after a data subject has consented to the terms and conditions of a data service a noteworthy event takes place. Contextually, this event is related to the processing of the user's personal data, which could be any circumstance ranging from a personal data breach to an opportunity to improve the user's privacy. It will be ultimately up to each user to define for herself to what extent what kinds of events are worthy of her attention, and how urgently she wants to be notified about them.

TETs that rely on privacy notifications to facilitate transparency effectively come into play at the time when said event is detected. For investigating matters related to HCI, the actual mechanism of how the event is detected is beyond the scope of this thesis. The review of the literature conducted in Paper I shows that the ex post TETs available until 2017 relied at least on three different conceptual paradigms to detect an event: (1) Local processing, i. e. on the mobile phone that also runs the health tracking software. (2) Remotely by the service provider that operates and administers the online mhealth service. (3) Remotely by a third party, i. e. an entity on which both user and service provider rely to facilitate transparency. While we do not elaborate on the pros and cons of each of these paradigms, we discuss the factors that may affect a user's perception of a system using either method. One of these dimensions is the question to what extent the TET itself is trustworthy, which we discuss in Papers I, IV and VI.

We distinguish between the delivery of a privacy notification and the detection of the cause that triggered it because a user may prefer to defer the message in question, i. e. to have it delivered at an opportune moment rather than immediately at the time when the event is detected. Design decisions regarding notifications (Papers V and VI) deal e. g. with how the device should deliver a particular message to signify a specific level of urgency, and how the modalities used for its delivery will affect whether it is actually noticed.

It is only at the time when the notification is actually consumed that users start interacting with the TET, and that considerations about the interaction between user and device become relevant. Once a notification is consumed, TETs provide users with customised information. Depending on the exact nature of the event that triggered the notification, TETs may guide users towards making informed decisions, such as to exert influence on how their

personal data will be processed in the future. This may or may not entail that a user exercises her legal right of intervenability.

We consider the following basic characteristics of TETs that leverage privacy notifications as a means to facilitate transparency:

Customisation. The spectrum of users of mhealth services shows great diversity. TETs will have to deal with a large variety of user needs. We address these needs via customisation.

Asynchronicity. Privacy notifications are asynchronous in nature. The exact time of when a notification will be delivered is non-deterministic. Recipients will typically not anticipate its arrival, or may potentially perceive its delivery as disruptive.

Constraints. Privacy notifications are subject to conceptual, technical and legal limitations in terms of how well they can achieve their designated goal. Designers of TETs will have to take these constraints into consideration, and will have to balance them with the expectations and needs of users.

Each of these themes will be covered in the following subsections, as will be the consequences that arise for the design of usable TETs.

5.2.2 Customisation

Users of mhealth services exhibit different usage patterns, stem from different age groups and genders, and are comprised of various socio-cultural backgrounds. They own and use different kinds of health tracking devices, such as fitness bracelets and smart watches, as well as more specialised devices such breast belts and headbands. A considerable number of users willingly share their data with their family or with relatives, with friends and acquaintances, in online social network [3] or open data platforms [52, 103], or with medical professionals or coaches. In some instances, users put sharing their data on a level with sharing their experience, transforming mhealth devices into 'experience sharing instruments' [62]. This yields a considerable number of possible scenarios in terms of how personal data may be processed. As regards proficiency, users have different levels of domain knowledge in areas such as ICT and law, both of which are relevant for understanding the intricacies related to the processing of personal data. Consequently, TETs will have to accommodate a large variety of user needs. In terms of usability, 'suitability for individualisation' is an established dialogue principle aiming to equip interactive systems with the means to meet the needs of a particular target audience for the purpose of letting them accomplish their goal [55]. Hence, customisation works towards achieving two other principles established in ISO 9241-110, namely a system's 'suitability for a task', and its 'conformity with user expectations'.

To address the heterogeneity in terms of what facts pertaining to personal data processing users are and are not interested in, we propose a three-fold

structure of categories that classifies privacy notifications based on the type of content they deliver. The categorisation was first introduced in Paper II, and has since been used as a baseline for grouping individual scenarios based on high-level themes. For the sake of concretization, we introduce the term *scenario* to denote a particular incident, event or circumstance related to the processing of personal data.

The three categories[10] of privacy notifications are as follows:

Breaches. (Personal data breach) "… a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data transmitted, stored or otherwise processed" (GDPR, Art. 4 (12)). This citation implies that Breaches cover both accidental mishap and deliberate misappropriation related to the processing of personal data.

Consequences. Information about how personal have been, or might be processed based on information that is currently at a particular entity's disposal. Consequences differ from Breaches in that respective notifications are the result of lawful processing.

Tips. Customised recommendations sent to a user to help her improve her privacy.

This three-fold segmentation allows a more precise discussion of *what* types of scenarios users choose to be notified about. This dimension of customisation is part of research question 2, which is presented in Section 3. Orthogonally, the dimensions of timing and modality serve as measures of *how* users prefer to be notified about individual scenarios or categories of scenarios. Jointly, timing and modality constitute a customisable measure of urgency and severity attributed to the delivery of privacy notifications and the information they convey. The concept of timing and modality has been introduced in Paper V and was revisited in Paper IV.

Timing. Timing specifies *when* a privacy notification will be delivered. We distinguish between immediate and deferred delivery. Immediate delivery refers to a notification being delivered as soon as a TET has obtained respective insight and has reached the conclusion that the information in question is relevant for the user. Revisiting the phases of Figure 5, immediate delivery coincides with the time of detection. Conversely, deferred delivery postpones the actual delivery to an opportune moment in the future, where 'opportune' refers to any point in time that suits the needs of the recipient of the message [36]. In this case, a specific amount of time passes between detecting the cause and delivering the message.

Modality. The modality of *how* TETs signal the delivery of a privacy notification can take multiple forms. The user study conducted in Paper IV

[10]For the sake of clarity, the three categories are capitalised.

distinguished between various signalling mechanisms leveraged by contemporary mobile phones: Vibration, blinking of the LED, playback of an audio signal, a visual pop-up on the screen, system notifications, and emails. Any combination of these modalities are imaginable for individual scenarios or groups of scenarios. Depending on the configuration of a phone, relying on system notifications may involve additional secondary modalities, such as vibration and LED. Emails differ from the other modalities in that they leverage the archival and organisational mechanisms of the user's email client. Like system notifications, email clients may or may not rely on secondary signalling modalities to notify about the arrival of a new message, or might even invoke a system notification for doing so.

Implementing individualisation ultimately means to enable users to customise the settings of a TET and to put them in control of its run-time behaviour. One approach that aims to facilitate individualisation explored in this dissertation is via user profiles. Profiles comprise preselected sets of default settings that reflect the predisposition of a particular group of users. Relying on profiles allows novice users to choose their settings from a short list of presets instead of spending a considerable amount of time on configuring a large number of individual settings from scratch. TETs may even propose suitable profiles to users whose predisposition is known to adequately reflect the characteristics of the profile in question. Drawing on the homonymous term from design theory [26], we rely on the term *persona* to refer to groups of users who share the same mental model [25, 124]. The *privacy personas* underlying individual user profiles pertain to the representatives' attitude and behaviour towards information privacy, which may or may not correlate with ethnographic or socio-cultural factors. Individuals can e. g. be segmented according to their general predisposition in terms of privacy [127], by considering their proficiency with ICT and their willingness to put their knowledge to practical use [31], or based on the perceived reliability of, and trust in their online data service [79, 80]. In the context of TETs that operate on the basis of privacy notifications, persona-based segmentation relates to sets of settings known to reflect the needs of the representatives of the persona in question. Paper III explores the relationship between the privacy personas established by Morton et al. [79, 80] and exemplary scenarios in the context of personal health tracking. Paper IV disregards pre-established personas and investigates alternative clusters of notification settings based on user preferences.

As an alternative approach to accommodate the needs of users of TETs, Papers V and VI discuss the impact of adaptive feedback of a user's behaviour on the settings of a TET. This approach seeks to adapt the run-time behaviour of a TET based on users' previous choices [6], or by relying on contextual cues to ascertain opportune moments in which users should be engaged [83]. It motivated a model that allows for adapting the run-time behaviour of a TET at any stage of the interaction (Figure 7 on p. 35).

5.2.3 Asynchronicity

In terms of cognitive processing, responding to situational change requires switching from a foreground activity to a secondary task, which often requires conscious decisions before the person concerned can resume the original activity. This can pose a challenge in that rational reflection, compared to instinctive or intuitive reaction, requires a considerable amount of time and imposes considerable cognitive load [78, 99]. Moreover, informed decision-making relies on a clear presentation of all relevant facts to allow for ascertaining reciprocal relationships between informational constructs that would provide more limited insight if the facts in question were considered only in isolation [5]. In terms of visualising informational contents, the paradigmatic shift from static structures to dynamic contexts has been named as one of the main challenges for designers of information systems [22].

While being actively engaged in a primary task, any interruption by introducing a secondary task is commonly considered distractive, disruptive, or otherwise detrimental to the performance of the foreground activity [36]. Mentally processing the interruption and engaging in action implies that the foreground process needs to be suspended to allow for carrying out the secondary task, and needs to be resumed once the secondary task has been finished [78]. The context switching poses additional cognitive load in addition to processing each task sequentially [78, 99], which calls for methods and tools that help optimise the processing of information related to circumstantial change.

The fact that change, and thus interruption, is rarely anticipated by the affected individuals adds to the complexity of the problem space. Throughout the ecosystem of mobile phones, the issue of being interrupted in carrying out one's foreground activity is exemplified prominently by situational contexts related to receiving notifications about matters that are contextually unrelated to the recipient's primary activity. Despite the omnipresence of mobile phones today and despite the utility of receiving notifications as a means to obtain awareness [69], receiving notifications is largely perceived as disruptive [70], typically because the arrival of a message occurs at inopportune moments [100]. Hence, classifying the urgency of individual notifications and selecting opportune moments for their delivery is crucial for ensuring efficacy and convenience [101]. As regards content-related determinants, notifications received on mobile phones are typically consumed relatively quickly [107]. However, the actual response time depends primarily on the type of the message in question [37, 102]. Moreover, notifications inevitably compete with each other, and multiple apps and tools will have to differentiate themselves in terms of the modality leveraged to signal delivery [77].

For privacy notifications, this means that the immediacy of the delivery depends on the severity of the type of scenario in question, such as a Breach notification in contrast to a Tip. Our research proposes initial sets of heuristics that may help determine notification settings based on the underlying scenario (Papers III–VI). Ultimately, however, the notification settings of a TET that operates on the basis of privacy notifications are subject to personal preference, which necessitates customisation.

The likelihood of receiving privacy notifications and the frequency in which respective messages will be delivered depends on the recipient's choice of noteworthy scenarios. However, even once that choice is made, the actual probability and frequency of such messages will not be deterministic, as notifications pertaining to all three categories depend on how the user's personal data have actually been processed. Ideally, the number of notifications would be sufficiently low to avoid detrimental effects related to longitudinal factors, such as fatigue and habituation [16, 32, 122]. Conversely, however, infrequent exposure or the lack of occasional reflection of the subject matter might result in suboptimal reaction on the part of the recipient in the rare event of actual exposure, as such exotic messages would be received even more unexpectedly. A recipient might lack the wits necessary to interact with them effectively, or might experience shock or even panic in response to learning about the worrying facts related to her privacy [7].

A crucial factor of conveying pertinent information as part of privacy notification is therefore to provide recipients with contextual cues that help them successfully perform the context switch from their respective primary tasks. This implies not only that the nature and purpose of the message must be clarified, but also that recipients will require cues of how to align their mental models with the causal relationship pertaining to the scenario in question. Recipients ought to be informed that the notification represents the effect of a cause related to detecting noteworthy facts about their privacy, such as an anomaly detected with respect to the processing of personal data. This event, in turn, will depend on what data the user had disclosed to the mhealth service. Any of these details may provide valuable information in itself, but only jointly will they provide holistic insight about the realities underlying the processing that triggered the notification. In this regard, the findings obtained in Papers V–VI indicate that informed decision-making depends on constructively aligning a recipient's state of mind with the causal relationship between individual facets of information related to the scenario at hand.

5.2.4 Constraints

The notion of privacy notifications as dealt with in this dissertation takes into account the following constraints:

Conceptual. We presuppose that events worthy of a data subject's consideration can be detected unequivocally, and that intelligence of how her personal data have been processed in the past allows TETs to draw meaningful conclusions about the status quo. We realise that this assumption is highly abstract. It is, however, necessary to stipulate the conceptual basis of our work.

Technical. At the time of writing, no technical standards exist in terms of protocols or APIs to access, query, rectify or erase personal data that are processed by data controllers and data processors. Consequently, the autonomous functionality required to implement the concurrent

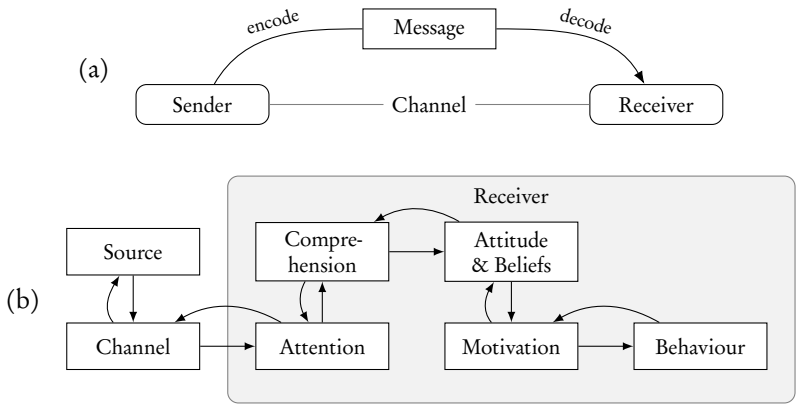

Figure 6: Models of communication according to (a) Shannon [110], Schramm [109] and Berlo [14], and (b) Wogalter et al. [24, 129].

run-time behaviour of privacy notifications is not feasible at this point. Following up on the previous item, we presuppose the availability of such information. In terms of scope, our research covers the phases starting at the time of delivering privacy notifications.

Legal. TETs have limited access to information pertaining to the actual processing of a user's personal data. Moreover, the operational measures that TETs are able to perform on behalf of data subjects are not only limited by practical constraints, but also by the fact that most scenarios will involve the interplay of multiple legal entities. Data subjects will have to rely on legal measures to exercise their legal rights. These measures may be proprietary, may not be particularly user-friendly, or may not necessarily be accessible via electronic means at all. Paper VI discusses implications that arise due to these constraints.

5.2.5 Modelling

We considered several models of communication for modelling privacy notifications according to the characteristics described in the previous sections. The models presented by Schramm [109] and Berlo [14] both build on and extend the information theoretical model by Shannon [110]. All three devise a sender who uses a transmission channel to transmit informational content in the form of a message to a receiver (Figure 6a). Whereas Shannon's model considers messages as syntactic constructs, Schramm's [109] and Berlo's [14] introduce human factors such as attitude and experience for both communication partners. In doing so, they elevate the notion of encoding and decoding information from a syntactic to a semantic level, which serves the purpose of indicating why the content is relevant and useful for the receiver. Consequently, the successful conveyance of informational content depends on factors related to the socio-cultural background and skill set of the sender. Conversely, the

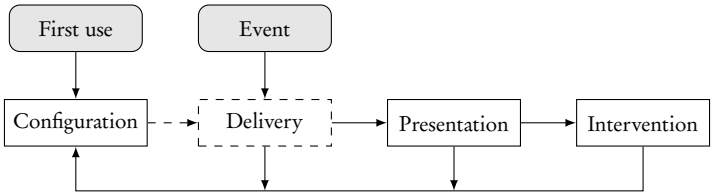

Figure 7: Design of interaction phases of privacy notifications.

receiver requires corresponding skills to decode the message such that the original meaning is restored. This duality creates a unique dependency between sender and receiver in that successfully transmitting high-level information relies on congruent skills and experience of both communication partners.

Neither Schramm's [109] nor Berlo's model [14] satisfactorily account for the asynchronous nature of privacy notifications. We therefore decided to abandon them in favour of a model that more adequately reflected the complexity of scenarios pertaining to personal data processing and the unexpected arrival of messages, the Communication-Human Information Processing model (C-HIP) proposed by Wogalter et al. [129]. Tailored for risk communication, C-HIP serves the purpose of raising awareness about the risk of imminent danger by issuing expedient warnings to instigate behavioural change. Starting from a source and a transmission channel, which reflect the transmitting entity and the modality of the message, respectively, this model comprises multiple phases traversed along the cognitive process from receiving a warning to acting in response (Figure 6b). C-HIP was conceptualised for working environments that pose a risk for a person's safety, yet the model applies similarly to risk communication in the context of privacy notifications. The Attention phase raises awareness about an incident, while the Comprehension phase advocates intelligible semantics. Attitudes & Beliefs is concerned with accommodating the recipient's predisposition, such as her familiarity with an environment. Motivation aims at relaying the actual costs of complying with the warning by taking action in contrast to accepting the consequences that will or may arise by refraining from acting. The concomitant factors of each phase can potentially interfere with the receiver's choice of taking action. Conversely, later phases may influence previous phases (curved backward arrows in Figure 6b), such as if familiarity with a working environment (Attitude & Belief) expedites understanding (Comprehension) or leads to habituation (Attention) [24].

The phases of C-HIP associated with the receiver's cognitive process exhibit similar capacities as the HCI design principles proposed by Patrick and Kenny [96]. The principles Comprehension, Consciousness, Consent and Control, all of which apply to contexts specifically related to information privacy, resemble the last four phases of C-HIP. Both constructs deal with scenarios that affect a person's quality of life and both deal with conceptual means to facilitate informed decision-making by communicating suitable facts. Furthermore, the Attention and Comprehension phases of C-HIP adequately account for the asynchronous nature of privacy notifications, and many of

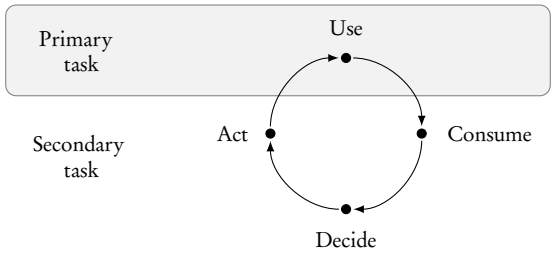

Figure 8: Cycle of user activities operationalised via privacy notifications: (1) Use service, (2) Consume notification, (3) Process and decide, (4) Act.

the recommendations published for C-HIP in terms of how warnings can be implemented effectively carry over to the design of privacy notifications [24]. Our conceptualisation of privacy notifications draws on the informational-cognitive properties of C-HIP and contextualises the model in the domain of personal health tracking. Introduced in Paper VI on p. 211, we propose a schema for modelling privacy notifications which reflects the asynchronous nature of a TET that harnesses privacy notifications. Instead of modelling individual states of a cognitive process, this model (Figure 7) traverses the principal phases during which users interact with the TET. Following the Delivery and consumption of a notification, the Presentation phase conveys customised, multilayered information about the incidence in question, whereas the Intervention phase advises users about how to intervene in the processing, should they choose to. The dedicated Configuration phase enables users not only to customise the run-time behaviour at the time when the TET is first put into operation, but also to perform subsequent adjustments during its entire service period. Customisation enables users to exercise control on individual phases of C-HIP by choosing prior to being notified what modalities will be used to signal the arrival of incoming notifications, when such messages will be delivered (Attention), and what kind of scenarios they prefer to be notified about (Attitude & Beliefs). By doing so, we seek to avoid cognitive barriers that might occur in the course of traversing respective phases due to discrepant behaviour of the TET or discordant attitudes on the part of users.

In essence, TETs that leverage privacy notifications constitute vigilant watchdogs on continuous duty. The process commences when a user starts using both the mhealth service and the TET monitoring the processing of her personal data (12 o'clock position in Figure 8). At this point, initial notification preferences need to be established to ascertain what kind of facts merit interrupting the user's primary task. Once the TET detects an anomaly or opportunity to improve the user's privacy, the user receives a notification about the facts pertaining to the scenario along with options of how to act in response. After consuming and processing the notification, the user decides whether and how to act based on the information at her disposal, and ideally continues using the service.

5.2.6 Demarcation

The main criterion by which privacy notifications differ from most work published in the literature is their longitudinal, potentially recurring modus operandi. Being inherently asynchronous in nature, they facilitate informed decision-making by combining concepts related to both ex post and ex ante transparency. Respective TETs aim to facilitate insight without prior impetus on the part of users other than their general decision to improve their privacy, which distinguishes them from most TETs that operate in contexts related to either ex post or ex ante transparency. In the case of ex post transparency, the more sophisticated tools presented in the literature typically rely on proactive measures on the part of a user, such as to request a transcript of the processing of one's personal data from a data controller and then use the TET to parse and visualise the contents [38]. In such contexts, users are responsible for going through all operation steps themselves. For ex ante transparency, the literature typically presupposes an ongoing transactional context in the course of which transparency-enhancing methods are leveraged to enable informed decision-making [108]. Such scenarios represent uninterrupted, self-contained processes that do not require context switching or contextual cuing on the part of a user. Both paradigms differ conceptually from the research presented in this dissertation in that existing solutions imply dedicated users who willingly take upon themselves the task of improving their privacy. Conversely, privacy notifications imply that user preferences are specified in advance, and that from that time onwards TETs involve users only in the event of circumstances that require their attention.

6 Summary of appended papers

Paper I. Tools for Achieving Usable Ex Post Transparency: A Survey

The paper builds on the result of a systematic literature review on usable ex post TETs. It analyses the TETs discussed in the literature with respect to their characteristics and overarching themes, and derives from them a taxonomy that classifies the final set of publications obtained via the literature research. Emphasising the aspect of usability principles and legal requirements, the paper elaborates on the particularities of each of the classification characteristics, and discusses to what extent they are reflected in each of the publications. The paper analyses them with respect to the principles established in the literature, and presents a list of gaps comprising properties that were not or only poorly implemented by TETs available at the time.

Paper II. Usable Transparency for Enhancing Privacy in Mobile Health Apps

Targeting the committee of the doctoral consortium held during the MobileHCI 2018 conference in Barcelona, this concept paper outlines future work pertaining to the principles of an ex post TET. It conceptualises a TET that

operates on the basis of privacy notifications as a means to inform users of online mhealth services about incidences related to the processing of their personal data, and proposes a three-fold classification to categorise respective messages. The paper was written during the preparatory stage of the study that led to Paper III.

Paper III. To Be, or Not to Be Notified

Privacy notifications represent an incident-based means to provide users of data services with ex post transparency. Drawing upon the usage context of personal health tracking, the paper reports on the results of an online study conducted to elicit preferences for privacy notifications received by users of online mhealth services. The paper shows that the causes pertaining to the notifications can be grouped thematically, that the user preferences do not correlate with privacy personas established in the literature, and to what extent the participants' right of intervenability affected their choice to be notified. Based on the findings, the paper infers principles for the design of usable TETs that operate on privacy notifications.

Paper IV. Reconciling the What, When and How of Privacy Notifications in mHealth

This paper refines the results obtained in Paper III in that it investigates multiple parameters in terms of serving as potential determinants for the settings of privacy notifications. Data are collected via two online surveys that draw on users of mhealth devices from two geographical regions in Europe. Analysing the quantitative responses, the paper infers from the results qualitative guidelines for designing TETs that harness privacy notifications to facilitate transparency. In doing so, we provide quantitative evidence for concepts that had been suggested in Paper V by drawing on multiple sources of the literature.

Paper V. Eliciting Design Guidelines for Privacy Notifications in mHealth Environments

This paper builds on the gaps and principles presented in Papers I and III, and conflates them with established findings from research on privacy indicators and push notifications received on mobile devices. It conceptualises a model reflecting the interaction phases of a TET that harnesses privacy notifications to facilitate transparency, which draws upon a user's preferences and past behaviour to provide her with customised settings. Superimposing the model with the findings, it derives guidelines for the design of TETs that employ privacy notifications to facilitate ex post transparency.

Paper VI. From Design Requirements to Effective Privacy Notifications: Empowering mHealth Users to Make Informed Decisions

This paper builds on the design guidelines for TETs obtained in Paper IV and augments them with additional prerequisites elicited from the literature, yielding a list of concrete design requirement. Their applicability is validated by conducting an iterative lab study that employs methods related to participatory design and exploratory research to obtain qualitative feedback on a prototypical implementation that is based on a selected subset of these requirements. The secondary purpose of the prototype is to serve as a blueprint for simulating interaction with a TET that relies on privacy notifications to facilitate transparency and informed decision-making. The paper presents a list of revised design requirements and a feasibility analysis for a prototypical design.

Paper VII. Opportunities and Challenges of Dynamic Consent in Commercial Big Data Analytics

The paper conceptualises 'dynamic consent' as a means to legitimise progressively evolving scenarios in which large companies process their customers' personal data to facilitate different online services. It covers the elicitation and specification of the term itself, a demarcation of dynamic consent from traditional forms of consent, and reports about the steps that led to the prototypical lo-fi design of a consent management tool. The prototype is exposed to the scrutiny of two focus groups consisting of domain experts from the field of information privacy. The paper reports and discusses the feedback received from the participants about the concept of dynamic consent as such, as well as about the strengths and weaknesses they noticed while discussing the interaction with the prototype.

7 Related work

This section briefly touches on literature that reflects areas of research related to the contributions listed in Section 5.

7.1 Conceptualisation of privacy notifications

Various groups of researchers investigated individual characteristics and particularities related to privacy notifications.

Patil et al. [95] investigate the interplay of timing and actionability on notifications related to the privacy of a user. They report that immediate feedback is most effective, but runs the risk of interrupting ongoing foreground tasks. In terms of finding a compromise between minimising interruption and raising awareness, they suggest that moderately delayed notification is sometimes deemed acceptable. We take from this finding that the timing of a privacy notification ought to be customisable to accommodate various types

of scenarios that call for different levels of urgency. We discuss this issue in Paper V, and ascertain and report actual user preferences in Paper IV.

Micallef et al. [77] investigate privacy notifications in the form of so-called 'privacy nudges'. They examine the perceived usefulness of soft-paternalistic nudges in terms of whether, when, and how the message in question ought to be delivered. The researchers conclude that users should be in control of whether and how they receive notifications to minimise the annoyance and intrusiveness of respective messages, and that the preferences of their test subjects depended on the context. They report salient messages as being generally undesirable, especially for nudges and receiving messages in the presence of other people. Salient signalling may, however, be appropriate for raising awareness or to ring alarm. As regards deferred delivery of notifications, their test subjects preferred no particular time of the day. The findings reported by Micallef et al. [77] largely mirror the results we obtained from conducting our online survey for Paper IV. According to our respondents' preferences, the context slightly affected whether, when and how they preferred to be notified, which we broke down into 12 distinctive scenarios.

Komninos et al. [60] investigated the design implications of using various modalities to receive notifications on mobile phones. They report that modalities differed in terms of salience and the implications they imposed on the recipients in terms of affecting their privacy when they receive such messages. The researchers rate audio signalling, vibration and LED in descending order both in terms of salience and privacy implications. Our results from Paper IV mirror these findings in that scenarios that seemed more urgent called for a higher total number of combinations of these and other modalities. Conversely, non-salient modalities, such as email, were selected more consistently across all scenarios.

In the context of raising awareness about app permissions on mobile devices, Jackson and Wang [57] rely on privacy notifications as an indicative means to emphasize discrepancies between a user's presumable attitude towards privacy and her actual choices in terms of trusting apps that require specific permissions. The purpose of their work is to raise situational awareness about potential deviations from a user's intended behaviour by harnessing preemptive just-in-time notifications to enable users to rethink their choices. They operationalise their model of users' predisposition by leveraging MUIPC [131],[11] a model that is primarily based on concern. We chose concern as one of the dimensions to model predisposition in Paper IV, but were unable to establish significant correlations between it and our test subjects' notification preferences.

7.2 Elicitation of design requirements for TETs

The Article 29 Working Party [8] discusses the legal basis for push and pull notices, and derives from them high-level guidelines for just-in-time notifications

[11] The Mobile Internet Users' Information Privacy Concerns scale (MUIPC) [131] constitutes a specialisation of Malhotra et al.'s [68] Internet Users Information Privacy Concerns (IUIPC) for the application context of mobile devices.

that facilitate ad hoc transparency. Along similar lines, Patrick and Kenny [96] base their design guidelines for click-through notifications on requirements stipulated by the law. They consider four phases pertaining to decision-making processes (comprehension, consciousness, control and consent), which relate to both ex ante and ex post transparency in that they model the deduction of actions – control or consent – from a preceding phase during which data subjects obtain sufficient insight about the matter in question. This particular view of conceptualising transparency encouraged us to introduce the asynchronous nature of privacy notifications in Paper V and motivated the model presented in Section 5.2.5.

The purpose of the design requirements presented by Fischer-Hübner at al. [40] is to facilitate transparency and accountability in cloud service scenarios. Some of these requirements apply to the usage context and methodology we have chosen for investigating privacy notifications in Paper VI.

The taxonomies presented by Bravo-Lillo et al. [15] and Schaub et al. [108] include guidelines and principles related to the design of, and interaction with systems that facilitate informed decision-making in the context of ex ante transparency. Many of these principles have been adapted into our design guidelines presented in Paper V and the requirements presented in Paper VI. We identified the contextualisation of information and the amalgamation of cause and effect, which we call 'contextual cues', as the missing link to make us of these authors' principles for the purpose of presenting informational contents in the context of privacy notifications.

The requirements presented by Cruzes and Jaatun [27] and Thomas et al. [118] for the context of designing privacy-preserving information systems are both based on the analysis of qualitative data collected in the course of empirical studies. They are based on interviews that reflect the views of laypersons and domain experts [27], and on the severity of perceived threats and harms in the context of online data services [118]. The dimensions covered in these publications motivated us to investigate proficiency and concern as possible determinants of notification preferences in Paper IV.

7.3 Literature survey and identification of gaps

As has been shown in our literature research [82], relatively few publications are available on ex post TETs that have systematically been designed with usability in mind. Likewise, few synoptic surveys and classification systems that touch on this topic exist.

Hedbom [47] provides a taxonomy of PETs that were available in 2009, which he uses to classify these PETs according to conceptual, socio-cultural and technological aspects. Published in 2013, Janic et al. [58] present a catalogue of TETs that contextualises the interplay of the factors trust, privacy concern, and transparency. Both groups of authors cover both ex ante and ex post TETs, and both publications deal only briefly with aspects of usability. They discuss PETs and TETs that were available at the time, and thus predate the legislation constituted by the GDPR.

Conversely, Paper I is based on a systematic literature research that specifically addresses the aspect of usability of ex post TETs published until November 2017. The paper provides a taxonomy of the TETs, points out gaps in the literature and suggests future directions regarding the design of ex post TETs that suffice applicable legal requirements and usability principles.

7.4 Evaluated prototypical designs

Of the TETs reviewed in our literature review [82] and classified in Paper I, individual implementations had actually undergone evaluation (column 'User study' in Table 2 on p. 71). However, only a select few of these artefacts were the result of a stringent design approach that included eliciting design requirements based on user needs, preferences and predispositions.

This gap in terms of demonstrable usability (see *Literature survey and identification of gaps* above) motivated us to open up the design processes employed for Papers VI and VII to accommodate input from actual users and domain experts, respectively. Hence, the evaluation of the prototypical designs provided us with cues as to what extent the design reflected the needs and expectations of our target audience (Paper VI), or gained traction in terms of accomplishing its designated goal (Paper VII). The interplay of context (mhealth and dynamic consent) and medium (privacy notifications) is not covered by existing TETs, which indicates the novelty of our work.

8 Limitations

As part of the research presented in this dissertation, we investigated selected dimensions related to the life cycle of privacy notifications (Section 5.2.1). This included e. g. matters related to customising the delivery of messages, to methods of presenting circumstantial information to facilitate transparency and informed decision-making, and to advise data subjects how to intervene. These factors are as much complementary as they are orthogonal. Findings obtained from previous research, our own as well other researchers', have carried over to research we conducted afterwards, effectively implementing an inductive approach of reflexive validation. However, the conclusions drawn about individual facets were, for the most part, elicited in isolated contexts. At no point did we verify the interplay of our findings holistically to confirm their actual applicability or efficacy. Consequently, many concepts discussed in relation to privacy notifications are based on hypotheses on our part, and constitute fragments of deductive reasoning that have not yet been validated. The reasons for considering individual facets of our research independently rather than holistically are as follows:

Constraints. The conceptual and technical constraints discussed in Section 5.2.4 complicate a holistic analysis of the subject matter. Consequently, individual research questions target respective issues in largely isolated, loosely coupled contexts.

Legal. Legal measures constitute orthogonal barriers to technical and operational processes, and add to the overall complexity of technological artefacts that facilitate such measures. The legal constraints mentioned in Section 5.2.4 entail that our work only covers HCI-related aspects of privacy notifications.

Feasibility. Practical constraints of what could realistically be covered as part of our studies prevented us from conducting, e. g., longitudinal studies or in situ field studies that would have allowed us to analyse the use of privacy notifications in real-world scenarios. The technological sophistication required to tap into the processing of actual mhealth data, if only locally on a user's mobile phone, would have been immense, and would have raised questions in terms of legal and ethical compliance.

Strategical planning. Despite the fact that the methodological human-centred approach employed throughout our research was decided on at an early stage, the specification of the usage context and the research conducted in this area commenced at a later point in time. This delay was to be expected, but made it challenging to deal exhaustively with matters related to the conceptualisation, implementation and evaluation of a prototype.

9 Conclusion and future work

Products and applications related to mhealth show an ongoing potential for growth. At the same time, usable means of obtaining transparency about the actual practices of data services do not seem to have found wide-spread usage. The European legislation mandates that intelligible information about how personal data are processed must be readily available to allow for informed decision-making on the part of data subjects. However, our review of the literature shows that the accessibility and intelligibility of respective information are not always sufficient to meet the needs of users. Moreover, respective users may not be able or willing to spend much time on auditing extensive amounts of sophisticated information. We therefore propose privacy notifications as a conceptual means to facilitate transparency and informed decision-making, the qualities of which we explore in the context of personal health tracking.

Privacy notifications build on the ecosystem of mobile phones by leveraging push notifications to raise awareness about circumstances worthy of the recipient's attention. Following up on this paradigm, we argue that TETs that harness privacy notifications may provide data subjects with situational consciousness of past events and promote opportunities to improve their privacy in the future. Information conveyed as part of a privacy notification is not only relevant for the recipient in that it is customised to meet the individual's expectations, but is also presented such that the scope and form of the message in question is tailored to suit the needs of the one receiving it. Despite the non-deterministic occurrence of events that warrant sending a privacy notification

to a user, our findings indicate that privacy notifications may be timed to arrive at opportune moments in that the nature of the event triggering the notification can serve as a determinant for ascertaining when and how respective messages should be delivered. Facts pertaining to the processing of personal data may be conveyed such that they facilitate transparency and provide valuable advice on how to take action in response to the insight obtained. The majority of the feedback received in the course of our research indicates that customised notification is considered useful and convenient, especially compared to actively enquiring for respective information. Hence, privacy notifications may pose a stepping stone not only towards legal compliance in terms of facilitating transparency, but also to satisfy the needs and expectations of users.

A noticeable limitation of the research presented throughout this dissertation is that individual factors were mainly considered in isolated contexts, and that no holistic evaluation was conducted to validate the interplay of respective findings. Future work could follow up on our work by conducting longitudinal studies that investigate relational aspects, such as how users perceive the delivery of notifications depending on the situations they find themselves in at that time. Respective field studies could explore the presentational properties related to providing contextual cues that help recipients relate the insight obtained in the present to events that occurred in the past.

Acknowledgement of contents

Prior work

PhD dissertations in computer science at Karlstad University consist of two phases. The second phase is based on work published in the form of a licentiate thesis submitted at the end of the first phase. This dissertation is based on and acknowledges work previously published as part of the author's licentiate thesis [81].

Work of others

Material Design icons. The three graphical symbols shown in Figure 3 on p. 16 belong to the collection of Material Design icons,[12] which are published under the Apache License V2.0.

Pictorial citations. The copyright for the figures cited in Paper I lies with the original proprietors.

[12] `https://material.io/`

References

[1] Chadia Abras, Diane Maloney-Krichmar, and Jenny Preece. User-centered design. *Bainbridge, W. Encyclopedia of Human-Computer Interaction. Thousand Oaks: Sage Publications*, 37(4):445–456, 2004.

[2] Mark S. Ackerman and Scott D. Mainwaring. Privacy Issues and Human-Computer Interaction. *Computer*, 27(5):19–26, 2005.

[3] Angeliki Aktypi, Jason RC Nurse, and Michael Goldsmith. Unwinding Ariadne's identity thread: Privacy risks with fitness trackers and online social networks. *Multimedia Privacy & Security Workshop at 24th ACM Conference on Computer on Computer & Communication Security (CCS)*, 2017.

[4] Manal Almalki, Kathleen Gray, and Fernando Martin Sanchez. The use of self-quantification systems for personal health information: big data management activities and prospects. *Health Information Science and Systems*, 3(1), 2015.

[5] Robert Amar and John Stasko. A Knowledge Task-Based Framework for Design and Evaluation of Information Visualizations. In *IEEE Symposium on Information Visualization*, pages 143–150. IEEE, 2004.

[6] Julio Angulo, Simone Fischer-Hübner, Erik Wästlund, and Tobias Pulls. Towards usable privacy policy display and management. *Information Management & Computer Security*, 20(1):4–17, 2012.

[7] Julio Angulo and Martin Ortlieb. "WTH..!?!" Experiences, Reactions, and Expectations Related to Online Privacy Panic Situations. In *Eleventh Symposium On Usable Privacy and Security (SOUPS)*, pages 19–38. USENIX Association, 2015.

[8] Article 29 Data Protection Working Party. Guidelines on transparency under Regulation 2016/679. 17/EN WP260 rev.01, 2018.

[9] Article 29 Working Party. Working Document on the processing of personal data relating to health in electronic health records (EHR). Technical report, 00323/07/EN WP 131. European Commission, 2007.

[10] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. The Impact of Timing on the Salience of Smartphone App Privacy Notices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, SPSM '15, pages 63–74. ACM, 2015.

[11] Christine Bauer and Simone Kriglstein. Analysis of Motivation Strategies in Running Tracking Applications. In *Proceedings of the 13th International Conference on Advances in Mobile Computing and Multimedia*, pages 73–79. ACM, 2015.

[12] Michael Bechinie (editor) et al. D4.1 User Interface Requirements. Technical Report D4.1/V1.0, Privacy&Us (deliverable), EU H2020 research and innovation programme, 2017.

[13] Victoria Bellotti. Design for Privacy in Multimedia Computing and Communications Environments. In *Technology and Privacy: The New Landscape*, pages 63–98. MIT Press, 1998.

[14] David Kenneth Berlo. *The Process of Communication: An Introduction to Theory and Practice*. Holt, Rinehart and Winston, 1960.

[15] Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security & Privacy*, 9(2):18–26, 2011.

[16] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W. Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. Your Attention Please: Designing Security-decision UIs to Make Genuine Risks Harder to Ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (SOUPS)*, pages 6:1–6:12. ACM, 2013.

[17] Ryan Calo. The Boundaries of Privacy Harm. *Ind. LJ*, 86:1131, 2011.

[18] Kim Cameron. The Laws of Identity. *Microsoft Corporation*, 2005.

[19] Fred H. Cate. The Limits of Notice and Choice. *IEEE Security & Privacy*, 8(2):59–62, 2010.

[20] Ann Cavoukian. Privacy by Design in Law, Policy and Practice: A White Paper for Regulators, Decision-makers and Policy-makers. Technical report, Information and Privacy Commissioner, Ontario, Canada, 2011.

[21] Yung Fu Chang, CS Chen, and Hao Zhou. Smart phone for mobile commerce. *Computer Standards & Interfaces*, 31(4):740–747, 2009.

[22] Chaomei Chen. Top 10 Unsolved Information Visualization Problems. *IEEE Computer Graphics and Applications*, 25(4):12–16, 2005.

[23] Roger Clarke. The digital persona and its application to data surveillance. *The information society*, 10(2):77–92, 1994.

[24] Vincent C. Conzola and Michael S. Wogalter. A Communication–Human Information Processing (C–HIP) approach to warning effectiveness in the workplace. *Journal of Risk Research*, 4(4):309–322, 2001.

[25] Kovila PL Coopamootoo and Thomas Groß. Mental Models: An Approach to Identify Privacy Concern and Behavior. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014.

[26] Alan Cooper. *The Inmates Are Running the Asylum: Why High-Tech Products Drive Us Crazy and How to Restore the Sanity*. Sams, 1998.

[27] Daniela Cruzes and Martin Gilje Jaatun. Cloud Provider Transparency-A View from Cloud Customers. In *CLOSER*, pages 30–39, 2015.

[28] Edward Cutrell, Eric Horvitz, and Mary Czerwinski. Notification, Disruption, and Memory: Effects of Messaging Interruptions on Memory and Performance. In *Human-Computer Interaction: INTERACT*, volume 1, page 263, 2001.

[29] Rachna Dhamija and Lisa Dusseault. The Seven Flaws of Identity Management: Usability and Security Challenges. *IEEE Security & Privacy*, 6(2):24–29, 2008.

[30] Drafting Committee of the Universal Declaration of Human Rights. *Universal Declaration of Human Rights*, 1948.

[31] Janna Lynn Dupree, Richard Devries, Daniel M Berry, and Edward Lank. Privacy Personas: Clustering Users via Attitudes and Behaviors toward Security Practices. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 5228–5239. ACM, 2016.

[32] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 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*, pages 1065–1074. ACM, 2008.

[33] European Parliament and the Council of the European Union. *Regulation (EU) 2016/679 of the European Parliament and of the Council*, 2016.

[34] European Parliament, the Council of Ministers and the European Commission. *Charter of Fundamental Rights of the European Union, 2000/C 364/01*, 2000.

[35] Pascal Faurie, Arghir-Nicolae Moldovan, and Irina Tal. Privacy Policy – "I agree"?! – Do alternatives to text-based policies increase the awareness of the users? In *International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, pages 1–6, 2020.

[36] Joel E. Fischer, Chris Greenhalgh, and Steve Benford. Investigating Episodes of Mobile Phone Activity as Indicators of Opportune Moments to Deliver Notifications. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, pages 181–190. ACM, 2011.

[37] Joel E. Fischer, Nick Yee, Victoria Bellotti, Nathan Good, Steve Benford, and Chris Greenhalgh. Effects of Content and Time of Delivery on Receptivity to Mobile Interruptions. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, pages 103–112. ACM, 2010.

[38] Simone Fischer-Hübner, Julio Angulo, Farzaneh Karegar, and Tobias Pulls. Transparency, Privacy and Trust – Technology for Tracking and Controlling My Data Disclosures: Does This Work? In *IFIP International Conference on Trust Management*, pages 3–14. Springer, 2016.

[39] Simone Fischer-Hübner and Helena Lindskog. Teaching Privacy-Enhancing Technologies. In *Proceedings of the IFIP WG 11.8 2nd World Conference on Information Security Education*, pages 1–17, 2001.

[40] Simone Fischer-Hübner, John Sören Pettersson, and Julio Angulo. *HCI Requirements for Transparency and Accountability Tools for Cloud Service Chains*, pages 81–113. Springer International Publishing, 2015.

[41] William H. Gates. Microsoft keynote for COMDEX 1990. `https://www.youtube.com/watch?v=uGA1Chm_8RE`, last visited June 1, 2020.

[42] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. How Short Is Too Short? Implications of Length and Framing on the Effectiveness of Privacy Notices. In *Twelfth Symposium on Usable Privacy and Security (SOUPS)*, pages 321–340. USENIX Association, 2016.

[43] Margaret Hagen. User-Centered Privacy Communication Design. In *Twelfth Symposium on Usable Privacy and Security (SOUPS)*. USENIX Association, 2016.

[44] Juho Hamari and Jonna Koivisto. "Working out for likes": An empirical study on social influence in exercise gamification. *Computers in Human Behavior*, 50:333–347, 2015.

[45] Marit Hansen. Top 10 Mistakes in System Design from a Privacy Perspective and Privacy Protection Goals. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 14–31. Springer, 2012.

[46] Marit Hansen, Meiko Jensen, and Martin Rost. Protection Goals for Privacy Engineering. In *IEEE Security and Privacy Workshops*, pages 159–166. IEEE, 2015.

[47] Hans Hedbom. A Survey on Transparency Tools for Enhancing Privacy. In Vashek Matyáš, Simone Fischer-Hübner, Daniel Cvrček, and Petr Švenda, editors, *The Future of Identity in the Information Society: 4th IFIP WG 9.2, 9.6/11.6, 11.7/FIDIS International Summer School*, pages 67–82, 2009.

[48] John Hicks, Nithya Ramanathan, Donnie Kim, Mohamad Monibi, Joshua Selsky, Mark Hansen, and Deborah Estrin. AndWellness: An Open Mobile System for Activity and Experience Sampling. In *Wireless Health 2010*, pages 34–43. ACM, 2010.

[49] Mireille Hildebrandt. Privacy and Identity. *Privacy and the criminal law*, 43, 2006.

[50] Mireille Hildebrandt. Behavioural biometric profiling and transparency enhancing tools. Deliverable D7.12, FIDIS, 2009.

[51] Leif-Erik Holtz, Katharina Nocun, and Marit Hansen. Towards Displaying Privacy Information with Icons. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 338–348. Springer, 2010.

[52] Open Humans. Explore, analyze, and donate your data – doing research together. `https://www.openhumans.org/`, last visited February 4, 2019.

[53] International Organization for Standardization. Guidance on usability. Technical Report ISO 9241-11:1998(E), ISO, 1998.

[54] International Organization for Standardization. Human-centred design processes for interactive systems. Technical Report ISO/TR 13407:1999(E), ISO, 1999.

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

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

[57] Corey Brian Jackson and Yang Wang. Addressing The Privacy Paradox Through Personalized Privacy Notifications. *Proceedings of the ACM Conference on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):68:1–68:25, 2018.

[58] M. Janic, J. P. Wijbenga, and T. Veugen. Transparency Enhancing Tools (TETs): An Overview. In *Third Workshop on Socio-Technical Aspects in Security and Trust*, pages 18–25, 2013.

[59] Kanitthika Kaewkannate and Soochan Kim. A comparison of wearable fitness devices. *BMC Public Health*, 16(1):433, 2016.

[60] Andreas Komninos, Jeries Besharat, Vassilios Stefanis, and John Garofalakis. Perceptibility of Mobile Notification Modalities during Multitasking in Smart Environments. In *14th International Conference on Intelligent Environments (IE)*, pages 17–24. IEEE, 2018.

[61] Braden Kowitz and Lorrie Cranor. Peripheral Privacy Notifications for Wireless Networks. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*, WPES '05, pages 90–96. ACM, 2005.

[62] Ioannis Krontiris, Marc Langheinrich, and Katie Shilton. Trust and Privacy in Mobile Experience Sharing: Future Challenges and Avenues for Research. *IEEE Communications Magazine*, 52(8):50–55, 2014.

[63] Ian Li, Anind Dey, and Jodi Forlizzi. A Stage-based Model of Personal Informatics Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 557–566. ACM, 2010.

[64] Zilu Liang, Bernd Ploderer, Wanyu Liu, Yukiko Nagata, James Bailey, Lars Kulik, and Yuxuan Li. SleepExplorer: a visualization tool to make sense of correlations between personal sleep data and contextual factors. *Personal and Ubiquitous Computing*, 20(6):985–1000, 2016.

[65] Deborah Lupton. M-health and health promotion: The digital cyborg and surveillance society. *Social Theory & Health*, 10(3):229–244, 2012.

[66] Deborah Lupton. Self-tracking Cultures: Towards a Sociology of Personal Informatics. In *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: The Future of Design*, pages 77–86. ACM, 2014.

[67] Martin Maguire. Context of Use within usability activities. *International Journal of Human-Computer Studies*, 55(4):453–483, 2001.

[68] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 15(4):336–355, 2004.

[69] D. Scott McCrickard, C. M. Chewar, Jacob P. Somervell, and Ali Ndiwalana. A Model for Notification Systems Evaluation – Assessing User Goals for Multitasking Activity. *ACM Transactions on Computer-Human Interaction*, 10(4):312–338, 2003.

[70] D. Scott McCrickard and Christa M. Chewar. Attuning Notification Design to User Goals and Attention Costs. *Communications of the ACM*, 46(3):67–72, 2003.

[71] D. Scott McCrickard, Mary Czerwinski, and Lyn Bartram. Introduction: Design and evaluation of notification user interfaces. *International Journal of Human-Computer Studies*, 58(5):509–514, 2003.

[72] Aleecia M. McDonald and Lorrie Faith Cranor. The Cost of Reading Privacy Policies. *ISJLP*, 4:543, 2008.

[73] Aleecia M. McDonald, Robert W. Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor. A comparative study of online privacy policies and formats. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 37–55. Springer, 2009.

[74] Rene Meis and Maritta Heisel. Understanding the Privacy Goal Intervenability. In *International Conference on Trust and Privacy in Digital Business*, pages 79–94. Springer, 2016.

[75] Merriam Webster Inc. Merriam Webster Online Dictionary. `https://www.merriam-webster.com/`, last visited February 4, 2019.

[76] Jochen Meyer, Anastasia Kazakova, Merlin Büsing, and Susanne Boll. Visualization of Complex Health Data on Mobile Devices. In *Proceedings of the Workshop on Multimedia for Personal Health and Health Care*, pages 31–34. ACM, 2016.

[77] Nicholas Micallef, Mike Just, Lynne Baillie, and Maher Alharby. Stop Annoying Me!: An Empirical Investigation of the Usability of App Privacy Notifications. In *Proceedings of the 29th Australian Conference on Computer-Human Interaction (OZCHI)*, pages 371–375. ACM, 2017.

[78] Yoshiro Miyata and Donald A. Norman. Psychological Issues in Support of Multiple Activities. *User centered system design: New perspectives on Human-Computer Interaction*, pages 265–284, 1986.

[79] Anthony Morton. *Individual Privacy Concern and Organisational Privacy Practice – Bridging the Gap*. PhD thesis, University College London, 2015.

[80] Anthony Morton and M. Angela Sasse. Desperately Seeking Assurances: Segmenting Users by their Information-Seeking Preferences. In *Privacy, Security and Trust (PST), 2014 Twelfth Annual International Conference on*, pages 102–111. IEEE, 2014.

[81] Patrick Murmann. Towards Usable Transparency via Individualisation. Licentiate thesis, ISBN 978-91-7867-003-1 (print), ISBN 978-91-7867-008-6 (PDF), Karlstad University, 2019.

[82] Patrick Murmann and Simone Fischer-Hübner. Usable Transparency Enhancing Tools: A Literature Review. Technical report, Karlstad University, Department of Mathematics and Computer Science, 2017.

[83] Inbal Nahum-Shani, Shawna N. Smith, Bonnie J. Spring, Linda M. Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A. Murphy. Just-in-Time Adaptive Interventions (JITAIs) in Mobile Health: Key Components and Design Principles for Ongoing Health Behavior Support. *Annals of Behavioral Medicine*, 52(6):446–462, 2017.

[84] ENISA (European Network and Information Security Agency). To log or not to log? – Risks and benefits of emerging life-logging applications. `https://www.enisa.europa.eu/`, November 2011.

[85] Jakob Nielsen. Heuristic Evaluation. *Usability Inspection Methods*, 17(1):25–62, 1994.

[86] Jakob Nielsen and Rolf Molich. Heuristic Evaluation of User Interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 249–256. ACM, 1990.

[87] Helen Nissenbaum. A Contextual Approach to Privacy Online. *Daedalus*, 140(4):32–48, 2011.

[88] Donald A. Norman. *User-Centered System Design: New Perspectives on Human-Computer Interaction*. CRC Press, 1986.

[89] Donald A. Norman. *The Design of Everyday Things*. Basic Books Inc., 2013.

[90] nPerf SAS Company. Cellular data networks in Sweden. `https://www.nperf.com/en/map/SE/`, last visited May 12, 2020.

[91] OECD. The OECD Privacy Framework. Technical report, OECD Publishing, 2017.

[92] Kieron O'Hara and Wendy Hall. Four Internets. *Communications of the ACM*, 63(3):28–30, 2020.

[93] Kieron O'Hara, Mischa M Tuffield, and Nigel Shadbolt. Lifelogging: Privacy and empowerment with memories for life. *Identity in the Information Society*, 1(1):155–172, 2008.

[94] Oxford Dictionaries. Oxford Online Dictionary. `https://en.oxforddictionaries.com/`, last visited February 4, 2019.

[95] Sameer Patil, Roberto Hoyle, Roman Schlegel, Apu Kapadia, and Adam J. Lee. 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, pages 1415–1418. ACM, 2015.

[96] Andrew S Patrick and Steve Kenny. From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions. In *International Workshop on Privacy Enhancing Technologies*, pages 107–124. Springer, 2003.

[97] Siani Pearson. Privacy, Security and Trust in Cloud Computing. In *Privacy and Security for Cloud Computing*, pages 3–42. Springer, 2013.

[98] John Sören Pettersson. A Brief Evaluation of Icons in the First Reading of the European Parliament on COM (2012) 0011. In *IFIP International Summer School on Privacy and Identity Management*, pages 125–135. Springer, 2014.

[99] Shari Lawrence Pfleeger, M Angela Sasse, and Adrian Furnham. From Weakest Link to Security Hero: Transforming Staff Security Behavior. *Journal of Homeland Security and Emergency Management*, 11(4):489–510, 2014.

[100] Martin Pielot, Bruno Cardoso, Kleomenis Katevas, Joan Serrà, Aleksandar Matic, and Nuria Oliver. 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, 2017.

[101] Martin Pielot, Karen Church, and Rodrigo de Oliveira. An in-situ study of mobile phone notifications. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI)*, pages 233–242. ACM, 2014.

[102] Martin Pielot, Amalia Vradi, and Souneil Park. 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)*, pages 3:1–3:11. ACM, 2018.

[103] Quantified Self. Self Knowledge Through Numbers. `https://www.quantifiedself.com/`, last visited July 2, 2020.

[104] Mashfiqui Rabbi, Min Hane Aung, Mi Zhang, and Tanzeem Choudhury. MyBehavior: Automatic Personalized Health Feedback from User Behaviors and Preferences using Smartphones. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 707–718. ACM, 2015.

[105] A. Rossi and M. Palmirani. A Visualization Approach for Adaptive Consent in the European Data Protection Framework. In *Conference for E-Democracy and Open Government (CeDEM)*, pages 159–170, 2017.

[106] Arianna Rossi and Monica Palmirani. Can Visual Design Provide Legal Transparency? The Challenges for Successful Implementation of Icons for Data Protection. *Design Issues*, 36(3):82–96, 2020.

[107] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. Large-scale Assessment of Mobile Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3055–3064. ACM, 2014.

[108] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing*, 21(3):70–77, 2017.

[109] Wilbur Schramm. How Communication Works. *The Process and Effects of Mass Communication*, 3, 1954.

[110] Claude Elwood Shannon and Warren Weaver. *The Mathematical Theory of Communication (and Recent Contributions to the Mathematical Theory of Communication)*. University of Illinois Press, 1949.

[111] Ben Shneiderman. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Addison-Wesley Publishing Co., 1987.

[112] Judith Siefring, editor. *Oxford Dictionary of Idioms*. Oxford University Press, 2nd edition, 2004.

[113] Daniel J. Solove. A Taxonomy of Privacy. *University of Pennsylvania law review*, pages 477–564, 2006.

[114] Daniel J. Solove. Privacy Self-Management and the Consent Dilemma. *Scholarly Commons*, 2013.

[115] Statista. Number of connected wearable devices worldwide from 2016 to 2021. `https://www.statista.com/statistics/487291`, last visited June 28, 2018.

[116] Statista. Projected size of the global market for wearable devices in the healthcare sector from 2015 to 2021. `https://www.statista.com/statistics/607982`, last visited June 28, 2018.

[117] Paul C. Tang, Joan S. Ash, David W. Bates, J. Marc Overhage, and Daniel Z. Sands. Personal Health Records: Definitions, Benefits, and Strategies for Overcoming Barriers to Adoption. *Journal of the American Medical Informatics Association*, 13(2):121–126, 2006.

[118] Keerthi Thomas, Arosha K. Bandara, Blaine A. Price, and Bashar Nuseibeh. Distilling Privacy Requirements for Mobile Applications. In *Proceedings of the 36th International Conference on Software Engineering (ICSE)*, pages 871–882. ACM, 2014.

[119] Matteo Turilli and Luciano Floridi. The ethics of information transparency. *Ethics and Information Technology*, 11(2):105–112, 2009.

[120] UAG "Standard-Datenschutzmodell" des AK Technik der Konferenz der unabhängigen Datenschutzbehörden des Bundes und der Länder. The Standard Data Protection Model V1.0. `https://www.datenschutzzentrum.de/sdm/`, last visited 14 July 2020, 2016.

[121] UAG "Standard-Datenschutzmodell" des AK Technik der Konferenz der unabhängigen Datenschutzbehörden des Bundes und der Länder. The Standard Data Protection Model V2.0b. `https://www.datenschutzzentrum.de/sdm/`, last visited 14 July 2020, 2020.

[122] Anthony Vance, Jeffrey L Jenkins, Bonnie Brinton Anderson, Daniel K Bjornn, and C Brock Kirwan. Tuning Out Security Warnings: A Longitudinal Examination of Habituation Through fMRI, Eye Tracking, and Field Experiments. *MIS Quarterly*, 42(2):355–380/A01–A15, 2018.

[123] Asimina Vasalou, Anne-Marie Oostveen, Chris Bowers, and Russell Beale. Understanding Engagement with the Privacy Domain Through Design Research. *Journal of the Association for Information Science and Technology*, 66(6):1263–1273, 2015.

[124] Melanie Volkamer and Karen Renaud. Mental Models – General Introduction and Review of their Application to Human-Centred Security. In *Number Theory and Cryptography*, pages 255–280. Springer, 2013.

[125] Samuel D. Warren and Louis D. Brandeis. The right to privacy. *Harvard law review*, pages 193–220, 1890.

[126] Alan F. Westin. Privacy and freedom. *Washington and Lee Law Review*, 25(1), 1968.

[127] Alan F. Westin. Social and political dimensions of privacy. *Journal of social issues*, 59(2):431–453, 2003.

[128] Svea Windwehr and Christoph Schmon. Our EU Policy Principles: Interoperability. `https://www.eff.org/deeplinks/2020/06/our-eu-policy-principles-interoperability`, last visited June 22, 2020.

[129] Michael S. Wogalter, David M. DeJoy, and Kenneth R. Laughery. Organizing Theoretical Framework: A Consolidated Communication-Human Information Processing (C-HIP) Model. *Warnings and Risk Communication*, pages 15–23, 1999.

[130] Gary Wolf. Know Thyself: Tracking Every Facet of Life, from Sleep to Mood to Pain, 24/7/365. *Wired Magazine*, 365, 2009.

[131] Heng Xu, Sumeet Gupta, Mary Beth Rosson, and John Carroll. Measuring Mobile Users' Concerns for Information Privacy. In *International Conference on Information Systems (ICIS)*, pages 2278–2293, 2012.

Information at Your Fingertips

The General Data Protection Regulation stipulates legal rights of transparency and intervenability. Transparency provides data subjects with insight into how their personal data have been processed, clarifying what consequences will or may arise due to the processing of their data, whereas intervenability enables them to intervene in the process. Technological artefacts, transparency-enhancing tools (TETs) serve the purpose of conveying respective information precisely and intelligibily. However, despite being a prerequisite for transparency, many TETs available today lack usability in that they do not stringently reflect the needs of their users, which raises the question as to whether individual TETs fulfil their designated purpose.

The objective of this dissertation is to systematically apply principles pertaining to human-centred design to ascertain the qualities necessary to design TETs that facilitate transparency and advise means of intervenability with regard to the needs of their target audience. We classify the state of the art of usable TETs published in the literature and discuss the gaps therein. Contextualising our research in the domain of personal health tracking, we investigate to what extent customisation can help accommodate the needs of users of TETs. We introduce privacy notifications as a conceptual means to inform data subjects about facts worthy of their attention, and examine the immanent properties required to accomplish actual usability. We categorise the characteristics of privacy notifications in terms of what insight they convey, and how respective facts need to be presented to facilitate informed decision-making on the recipient's part. Based on findings obtained via quantitative and qualitative user studies, we elicit concomitant factors related to the parameterisation of privacy notifications. We present the prototypical implementation of TETs whose iterative evaluation provides us with a catalogue of design requirements that demonstrably reflect the needs of their users.

ISBN 978-91-7867-144-1 (print)

ISBN 978-91-7867-148-9 (pdf)

ISSN 1403-8099

DOCTORAL THESIS | Karlstad University Studies | 2020:28

INFORMATION AT YOUR FINGERTIPS

The General Data Protection Regulation stipulates legal rights of transparency and intervenability. Transparency provides data subjects with insight into how their personal data have been processed, clarifying what consequences will or may arise due to the processing of their data, whereas intervenability enables them to intervene in the process. Technological artefacts, transparency-enhancing tools (TETs) serve the purpose of conveying respective information precisely and intelligibily. However, despite being a prerequisite for transparency, many TETs available today lack usability in that they do not stringently reflect the needs of their users, which raises the question as to whether individual TETs fulfil their designated purpose.

The objective of this dissertation is to systematically apply principles pertaining to human-centred design to ascertain the qualities necessary to design TETs that facilitate transparency and advise means of intervenability with regard to the needs of their target audience. We classify the state of the art of usable TETs published in the literature and discuss the gaps therein. Contextualising our research in the domain of personal health tracking, we investigate to what extent customisation can help accommodate the needs of users of TETs. We introduce privacy notifications as a conceptual means to inform data subjects about facts worthy of their attention, and examine the immanent properties required to accomplish actual usability. We categorise the characteristics of privacy notifications in terms of what insight they convey, and how respective facts need to be presented to facilitate informed decision-making on the recipient's part. Based on findings obtained via quantitative and qualitative user studies, we elicit concomitant factors related to the parameterisation of privacy notifications. We present the prototypical implementation of TETs whose iterative evaluation provides us with a catalogue of design requirements that demonstrably reflect the needs of their users.

ISSN 1403-8099 | ISBN 978-91-7867-144-1 (print) | ISBN 978-91-7867-148-9 (pdf)

DOCTORAL THESIS | Karlstad University Studies | 2020:28