Usable Transparency Enhancing Tools

A Literature Review

Patrick Murmann, Simone Fischer-Hübner

Faculty of Health, Science and Technology

Usable Transparency Enhancing Tools

A Literature Review

Patrick Murmann, Simone Fischer-Hübner

WORKING PAPER | July 2017

Usable Transparency Enhancing Tools - A Literature Review

Patrick Murmann, Simone Fischer-Hübner

WORKING PAPER | July 2017

urn:nbn:se:kau:diva-56757

ISBN 978-91-7063-804-6 (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 2017

WWW.KAU.SE

Usable Transparency Enhancing Tools: A Literature Review

Patrick Murmann, Simone Fischer-Hübner

Department of Mathematics and Computer Science,
Karlstad University

Abstract. This technical report documents the procedure of a literature review conducted on usable ex post-transparency enhancing tools (TETs). The review of scientific literature serves the purpose of providing insight into the characteristics of existing implementations of usable TETs. By providing a concise summary of existing implementations, the report aims to facilitate future research on the subject matter.

Keywords. literature review, privacy, transparency, transparency enhancing tools, usability, visualisation

Acknowledgement

The project leading to this report has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 675730.

1 Introduction

Over the last decade, more and more privacy-sensitive transactions have been realised online, not least thanks to the propagation and continuous success of cloud services. In many cases, data subjects, that is, users and usees of such services, have only little knowledge about the underlying mechanisms or the stakeholders involved in processing and storing of their personal data. Even though in Europe data controllers are legally obliged to provide privacy policies that state how and under what exact circumstances the users' personal data are being processed [7], such policies are rarely written in such a way that laypersons could clearly infer all potential consequences that will or might arise due to them using a respective service. The goal of transparency enhancing tools (TETs) is to make the underlying processes more transparent, and to enable data subjects to better understand the implications that arise due to their decision to disclose personal data, or that have arisen due to choice made in the past.

However, the usability of a TET is a necessary condition for the transparency it is supposed to provide. TETs that are deployed in real-world scenarios must provide self-explanatory user interfaces and visualise personal data and other circumstantial information in a clear, comprehensible way specifically designed for the experience and expectations of the intended target audience. In order to catch an audience that is as large and diverse as the Internet community data visualisation is often simplified. Abstraction and simplicity, however, come at the price of information loss or reduced interactional capabilities, entailing a trade-off between the functionality of TETs and their usability.

On a conceptional level, TETs can be subdivided into 'ex ante' and 'ex post'-TETs [13]. Ex ante-TETs guide the user's decision making process *before* she makes her choice pertaining to disclosing any personal data to a data controller. Guidance on part of the TET occurs by making transparent the exact circumstances of the situation. The objective is to guide the user's decision in an as understandable as possible way, informing her about the implications that will or might arise due to her decision. Conversely, ex post-TETs visualise disclosed personal data in such a way as to make transparent the processes that have taken place *since* the user has disclosed her data. In addition to normative or regulatory stipulations, such as the legally binding privacy policy provided by a data controller, the actual value, that is, the actually disclosed personal data that has been processed, stored, and potentially redistributed by the data controller, must be taken into account by such TETs. During the initial phase of this literature review, the reviewers discovered that most research had been conducted in the area of ex ante-TETs, and decided to narrow down the scope to usable implementations of ex post-TETs.

In order to get a better overview of what kind of ex post-TETs exist, which specific approaches they take, which functionality they provide in which usage context, and which respective trade-offs they entail, the authors performed a systematic literature review that covers the subject matter of usable ex post-TETs. The details of conducting this review is documented throughout the remainder of this report.

This report is structured as follows: Section 2 briefly refers to former literature surveys conducted on TETs and other privacy enhancing technology. Section 3 lays down the exact contextual scope of the review, and Section 5 specifies the criteria used to select individual publications from the body of scientific literature. Section 6 provides an overview over the methodology used to conduct the literature review. Section 7 goes into detail about the databases and literature search engines used to retrieve the publications. Section 8 elaborates on the characteristics of the search terms and search queries used to query said databases. Section 9 documents the screening process used to select or disregard individual publications for the review by applying the criteria specified in Section 5. Section 10 explains the snowballing process used to retrieve secondary papers contextually related to the initial set of screened publications. Section 11 discusses how the publications that were selected during the various stages of the retrieval process were aggregated into one final result set. Section 12 presents a qualitative taxonomy that was used to classify all reviewed papers. Finally, Section 13 concludes the review.

The appendix of this document briefly reflects each selected publication in terms of its contents, as well as in terms of the distinctive characteristics of its implementation. Moreover, each publication is classified according the taxonomy introduced in Section 12.

2 Related Work

Hedbom's [12] survey on implementations of privacy enhancing tools (PETs) briefly discusses the conceptual, legal, and technical aspects of PETs. The author specifies an extensive taxonomy used to classify each reviewed implementation. The scope of the survey comprehends usable implementations, such as stand-alone products and browser plug-ins, and explicitly disregards enabler technologies and protocols.

Janic et al.'s [17] survey on TETs specifically focuses on the aspect of trust that data subjects put in a data service. The authors try to answer the questions of how personal concerns with regard to privacy are related to their trust in a respective service, whether and to what extend transparency and trust are mutually related, and what value TETs can provide users with when it comes to evaluating trust relationships between users and service providers.

Both Hedbom [12] and Janic et al. [17] briefly touch upon the usability of the reviewed tools. Hedbom [12] specifies 'Ease of access' as one aspect of the security requirements of a PET. Janic et al. [17] consider the aspect of comprehensibility of TETs, and argue that as a result of being comprehensible a respective tool will better support the user's awareness of the information it visualises.

Conversely, the literature review at hand was specifically conducted with the usability of TETs in mind. In the context of this review, the usability of a TET is not only considered an optional asset or secondary byproduct, but a hard requirement to the respective implementation. The querying and retrieval process of the review explicitly disregards publications that are not related to

usable prototypes of ex post-TETs. In this regard, the contextual scope of this review is narrower than the two aforementioned surveys. Chronologically, this review expands the range of reviewed literature since Janic et al.'s [17] survey from 2013, and takes into account a body of literature published until 2016.

3 Subject matter

The research topic considered during this literature review focuses on the overlap of three major sub topics:

Privacy. The central focus of the research is related to *privacy*. The term 'privacy' is not distinctively tied to a single academic discipline, but covers multiple aspects, such as legal, societal, psychological and technical facets, many of which are intertwined and interdependent. Consequently, privacy has applications in many disciplines simultaneously, such as law, social sciences, psychology, and computer science. Within the scope of this review, privacy specifically relates to informational privacy in the context of information technology [14].

Transparency. In the context of this review, *transparency* relates to the legal right of a data subject to gain insight into the processing and storing of her personal data by a data controller. This right is stipulated in Chap. III, Art. 12 et seq. of the General Data Protection Regulation (GDPR) of the European Union [8]. In the case of TETs employed in ex post-scenarios, transparency relates to visualising the individual data items that have deliberately or accidentally been disclosed to a particular data controller or downstream processor. Only by possessing comprehensive knowledge about the underlying processes, the stakeholders involved, as well as any data flows between them is a data subject able to obtain insight about the overall scope of her disclosed data, and is able to make a well-informed decision as to whether or not to exercise legal rights as regards, for example, the rectification or deletion of her data (GDPR, Chap. III, Art. 16 et seq.)

Usability. The usability of software, as it is comprehended in this review, refers to the respective software being effectively, efficiently, and satisfyingly usable by the target audience it was designed for, as specified in ISO 9241-11 [16]. Useable software must take into consideration the particular experience and previous knowledge (or lack thereof), as well as expectations of the actual target users. In order to be comprehensible by the respective audience, information must be displayed in such a way as to pose as little a cognitive load as possible, preferably in an as simple as possible way [25,32].

These three major topics form the overarching context of the literature review. Publications that are considered relevant for the review are expected to be found in the intersection of these topics, representing the contextual overlap of their various aspects.

4 Scope

This Section specifies the formal scope of the review. It builds upon the *Taxonomy of Literature Reviews* discussed by Randolph [28], which refers to the original taxonomy proposed by Cooper [6]. This taxonomy defines six *characteristics* whose interplay allow for an overall classification of literature reviews. Each of these dimensions has one or more possible properties called *levels* that specify how the review is positioned in the respective context. More than one level may apply for each characteristic.

The following compilation represents the formal classification of this review according to said taxonomy. The six characteristics are printed in bold face, whereas terms printed in italics stand for the canonical levels discussed in [28].

Focus. The review at hand focuses on the *outcomes* as well as the *practices or applications* of the reviewed publications. Although the underlying methods and theories are briefly being reflected, it is the final implementations and their applicability that is of primary interest for this review. The criteria of the screening process, as specified in Section 5, explicitly rely on the availability of an actual implementation or prototype.

Goal. In terms of the intended goal, the review aims to provide an *integration* of the TETs published in scientific literature. Although the findings are briefly discussed and classified, the main purpose is to analyse and present the status quo of ex post-TETs that are available today. As such, this review may serve as a basis for further critical reflection of individual implementations.

Perspective. The review's perspective is intended to be as *neutral and unbiased* as possible, and the same classification scheme (as specified in Section 12) applies to all reviewed material in exactly the same manner.

Coverage. The review is designed to be *exhaustive with selective citations.* It is exhaustive in that it considers all publications available from the selected sources so long as they fulfill the screening criteria specified in Section 5. It is highly selective in that the thematic scope is specifically limited to publications whose contents are congruent with the topics specified in Section 3. Moreover, the choice of the databases and search terms used for the information retrieval process, and the deliberate exclusion of other sources of information may certainly be considered selective. For that matter, the review aims to cover a *representative* selection of publications in the specific context of computer science and engineering.

Organisation. The review is structured as a combination of *conceptual* as well as *methodological* elements. It is conceptual in that the criteria specified in Section 5 and the formal taxonomy laid down in Section 12 both constitute identifiable properties in TETs that, once ascertained, can be used to classify the respective technology. Abiding by a rigorous methodology as regards the information retrieval and screening process, the review itself is *methodological* in nature.

Audience. The review is directed towards *general* as well as *specialised* scholars and experts of respective research fields. For the former, the review condensates an overview of ready-to-use or prototypical implementations that are

available today. For the latter, it provides a taxonomy that served as a basis to scrutinise the examined technologies, and that allows for categorising individual TETs based on the characteristics of the taxonomy.

5 Selection criteria

This section describes the criteria applied for screening the retrieved papers during the initial search, as well as during the screening of the references and citations derived from the initial aggregation.

Peer-reviewed. In order to be considered for inclusion, a publication must be peer-reviewed, ensuring the level of quality and rigour expected from scientific publications. This criterion, among other aspects, influenced the choice of databases used for the retrieval process (Section 7).

Implementation. The publication must either have developed or extend a concrete solution that actually implements the TET. Papers that describe a concept, model, or theory are not considered here, because respective TETs would lack interactive and visualisation capabilities, and thus actual usability. In cases where prototypes were discussed by the authors, but the maturity of the TET could not be determined,[1] the respective publication was included and the lack of information discussed in the summary of the paper in the appendix.

Availability of full-text. The availability of a full-text version of the publication to actually evaluate its contents was considered a prerequisite. The rationale behind this decision is further discussed in Section 7.

User study. It was considered beneficial if user studies had been conducted to evaluate the usability of the respective TET. Correspondingly, the type of the conducted user study served as respective dimensions for the taxonomy of this review, as specified in Section 12. However, lacking a user study was not reason for excluding the paper.

Publication type. Neither the publication type, nor the reputation of the publisher, journal, or conference were taken into account while screening the papers. An implicit preselection occurred insofar as two out of four of the selected databases were actually publishers (see Section 7).

Article type. Only primary sources were considered for the review, that is, publications that described original research of TETs. Secondary sources, that is, publications that referred to other papers without actually extending them were only considered if they ultimately presented actual implementations of their own.[2] Tertiary sources, such as reviews and collections of articles, were not considered.

Domain. The disciplines of the researchers and their affiliations were not considered a criterion during the screening process. Publications from all fields

[1] Some papers lacked technical details about the implementation, as well as textual descriptions of the actual status of the TET, respectively.

[2] See the discussion on the removal of duplicates in Section 6.

and sources were considered, and no subject-based filtering was conducted on databases that offered a respective filtering based on subject or field.

Language. Only publications in English were considered.

Extent. Publications were not screened based on the amount of pages or their word count, so long as they fulfilled the other selection criteria. The rationale was that a brief and concise technical report could provide as valuable an insight into a TET-prototype as a verbose article.

6 Methodology

The phases of the information retrieval and review process are schematically depicted in Figure 1. While the initial search of the review was based on the guidelines by Kitchenham et al. [19], the snowballing phase is based on the recommendations of Webster et al. [37]. This strategy allows for a initial search based on well-defined search terms, and complements it with a simplified but systematic trace of references during the snowballing phase.

This slightly modified approach as compared to following the recommendations of a single authority was chosen for practical reasons: On the one hand, it allowed for relying on search terms and publication databases instead of a preselected set of relevant publications in the beginning of the review process. On the other hand, it allowed for a simplified, completely automated snowballing process instead of an exhaustive recursive descent of the relevant articles, as the latter approach would not have been feasible due to temporal constraints.

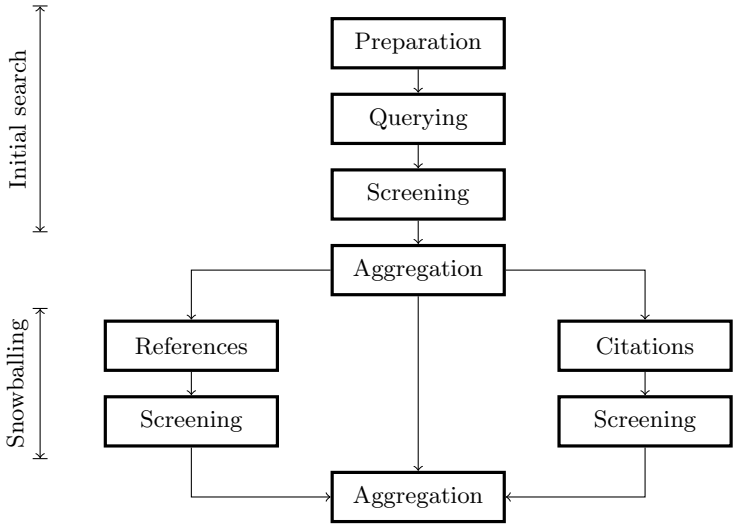

Fig. 1. Phases of the literature review

Based on a well-defined subject and scope as laid down in Section 3, publication databases (Section 8) and keywords (Section 8) were chosen during the preparation phase. Abiding by the recommendations of [19,28,36,37], all stages of the review, as well as all micro operations were rigorously documented.

Each selected database was queried independently using the specified keywords (Section 8), yielding a relatively large number of results. The publications so retrieved were sieved for relevant papers during the subsequent screening process (Section 9).

In order to retrieve additional relevant publications, the references and citations of the screened papers were being traced by conducting 'snowballing' both backwards and forwards (Section 10). The publications thus retrieved were screened by applying the exact same criteria as during the intial search. During the aggregation phase, duplicates and redundancies were removed from the result set yielded during the initial search, as well as the ones retrieved via snowballing (Section 11). Aggregation denoted the end of both the initial search and snowballing phase, each yielding a set of publications that sufficed the selection criteria specified in Section 5.

7 Selection of databases

The set of databases and their respective query settings used during the initial querying phase is listed in Table 1. The databases ACM[3], DBLP[4], IEEE[5], and Inspec[6] were selected for the following reasons:

The four literature databases represent the intersection of the recommendations from several senior researchers, colleagues, subject librarians, and bibliography experts. In the authors' home institution, these databases were the ones fellow researchers were most familiar with. They were recommended because of their reputation as regards the quality of the papers either originally published or referenced therein. More precisely, two of the databases, Inspec and DBLP, are independent literature databases that reference the publications of other publishers. Conversely, ACM and IEEE are publishers in their own right whose search engines comprehend only the publications released in their own journals, conference proceedings, and other publication channels. When queried for respective search terms (see Section 8), the four databases returned a high, yet manageable amount of results.

Springer was initially considered as an additional database, but was ultimately disregarded because it did not allow queries to be limited to the titles, abstracts and keywords of the papers it contained. When queried, Springer applied the search terms to the entire body text of the indexed papers, and

[3] ACM Digital Library (`https://dl.acm.org/`)
[4] DBLP Computer Science Bibliography (`http://dblp.dagstuhl.de/`)
[5] IEEE Xplore (`http://ieeexplore.ieee.org/Xplore/home.jsp`)
[6] Inspec via EBSCOhost (`https://www.ebscohost.com/`)
[7] In accordance with the selection criteria in Section 5, only English publications were considered during the screening phase (Section 9).

Database	Settings
ACM	Scope: the full ACM collection
IEEE	Search type: metadata only; content: all; year of publication: any
Inspec	Limiter: full-text; types: any; subject: any; publication: any; publisher: any; language: any[7]; classification: none
DBLP	Database: Dagstuhl

Table 1. List of the preselected databases and their query settings.

consequently yielded an unmanageably large amount of results. Likewise, other interdisciplinary publishing houses such as Scopus or Google Scholar, would not only have extended the breadth of the covered disciplines, but would also have returned an unfeasibly large amount of potential results.

With only few exceptions, Karlstad University had access to full-text versions of all publications found on the selected databases. Respective shortcomings might have been overcome by means of interlibrary loans. However, due to the relatively tight time frame of the project, such loans would have complicated the information retrieval process considerably.

IEEE and Inspec are primarily related to engineering, information technology, and natural science, while ACM and DBLP publish and reference documents in the field of computer science, respectively. It was obvious that publications retrieved via these databases would most likely be biased towards the research culture of science. However, Section 8 will show that the genericity of the search terms used for the queries catches a multitude of disciplines beyond the research fields primarily represented by these databases.

8 Querying

All database queries were specifically limited to the title, abstract, and keywords[8] of the publications. By doing so, the reviewers hoped to keep the retrieval of unwanted publications that mentioned one of the search terms in the body text to a minimum. However, this attempt was only partially successful, as is discussed in Section 9.

Since the research topic spanned multiple disciplines, the subject-specific nomenclature of any one discipline was unlikely to appropriately reflect the subject matter as a whole. It was therefore inevitable to rely on generic synonyms to cover as many as possible usage contexts to constitute the thematic scope of the review.

[8] Including both, the authors' as well as the publisher's predefined keywords, if any.

Based on the definition of the subject matter in Section 3, the thematic domain d of the research represents the cojunction of the subdomains privacy, transparency, and usability:

$$d := \mathit{privacy} \wedge \mathit{transparency} \wedge \mathit{usability} \tag{1}$$

Relevant publications must therefore reside in the contextual intersection of all three domains. Further, let the following be defined, where monospaced text signifies literals:

$$\mathit{privacy} := \texttt{privacy} \tag{2}$$

$$\mathit{transparency} := \texttt{transparen}* \vee \texttt{polic}* \tag{3}$$

$$\mathit{usability} := \texttt{usab}* \vee \texttt{visual}* \tag{4}$$

The sub clause *privacy* is represented by a single literal expression (equ. 2) because no alternative spelling for the term exists, nor does a generally acknowledged synonym with equal or similar conciseness.

The clause *transparency* is split up into the two contextually related facets transparency and policy (equ. 3). Privacy policies denote codified statements about the legal or technical process underlying the storage and processing of personal data. As such, 'policy' is not considered a synonym for transparency in terms of semantics, but represents the normative ground truth used by a TET to make a justified statement as to whether a particular personal data items as legitimately being stored and processed by a data controller according to the underlying privacy policy. In order to catch syntactic variations, such as plurals and nominalised adjectives, the two sub components are represented by the wildcard-terminated literals '`transparen`$*$' and '`polic`$*$', respectively.

Similarly, the clause *usability* is subdivided into the two termini usability and visualisation (equ. 4). In human computer interaction (HCI), the term 'visualisation' refers not only to displaying graphical data, but more generally to the act of conveying information by employing the most appropriate, sometimes multi-modal means for the respective usage context [23]. In order for an application to become usable visualising relevant information is crucial regardless of the modality chosen for this task. Consequently, visualisation is introduced into the query as an equally important extension of usability. Again, wildcards are used to accommodate syntactical variations of the terms, yielding the literals '`usab`$*$' and '`visual`$*$', respectively.

Applied to domain d, the conjunction of the aforementioned disjunctions yields the following search query q:

$$q = \texttt{privacy} \wedge (\texttt{transparen}* \vee \texttt{polic}*) \wedge (\texttt{usab}* \vee \texttt{visual}*) \tag{5}$$

The query parsers of the selected databases expected different syntaxes for the logical operators, and, in the case of ACM and DBLP, were not self-explanatory. IEEE and Inspec used '`&`' or '`AND`' for conjunctions, and '`|`' or '`OR`' for disjunctions. The exact query strings used to express q for the four databases are listed in Table 2.

Database	Query
ACM	`+(privacy) +(transparen* polic*) +(usab* visual*)`
IEEE	`privacy AND (transparen* OR * polic*) AND (usab* OR visual*)`
Inspec	`privacy AND (transparen* OR * polic*) AND (usab* OR visual*)`
DBLP	`privacy$ (transparen | polic) (usab | visual)`

Table 2. Search query literals for the selected databases.

Date	ACM	IEEE	Inspec	DBLP
2016-11-15	565	289	28	17
2016-11-21	565	291	30	17
2016-11-28	565	292	30	17
2017-01-02	567	294	30	17
2017-01-09	568	294	30	17

Table 3. Number of results returned by various databases on specific dates.

It is worth noting that DBLP's query optimiser automatically extends all search terms to allow for the greatest possible variety of suffices. In order to achieve the exact same semantics for DBLP as for the three other databases, the literal '`privacy$`' was explicitly terminated by an end-of-word symbol (`$`).

The parentheses in q enforce explicit precedence of the two disjunctions over the conjunction. Contrary to formal logic, the query parsers of many databases evaluate disjunctions before conjunctions, which is why parentheses in q would not be strictly necessary. ACM and IEEE-Xplore emphasise the precedence of disjunctions by reformatting parenthesised search queries and removing obsolete parentheses from the query input field.

Table 3 lists the number of results returned by the four selected databases on different dates. The increasing number of results of ACM and IEEE emphasise the fact that considerable research is currently being conducted and published on the subject matter. In fact, throughout November and December 2016 the reviewers were notified by IEEE Xplore's search alert approximately once per week, indicating the availability of new publications that matched q.

9 Screening

During the screening process, all retrieved publications were matched against the scope specified in Section 4 and the criteria specified in Section 5. For this purpose, each retrieved bibliography record was first matched against the formal criteria in Section 5, such as the requirement of for all publications to be peer-reviewed. The majority of publications were either conference proceedings or journal articles and passed this test.

Next, the abstract of each publication was read in order to decide whether the publication matched the scope laid down in Section 4. In cases where a clear

Database	Retrieved	Relevant	Precision
ACM	565	8	0.01
IEEE	289	4	0.01

Table 4. Precision of q with regard to the selected databases as determined after the first stage of the screening process (date: 2016-11-15).

decision could not be made immediately, the publication was preliminarily set aside for later investigation. Publications that either passed this test or whose status was formerly qualified as unclear were each read in the following order of their sections: Abstract, conclusion, discussion, introduction, main part. Rereading these sections was repeated until a final decision in terms of relevance could be reached. Once compliance was affirmed, the publication was reread and notes were taken as regards the taxonomy specified in Section 12.

Only a minor fraction of the publications retrieved from the databases actually passed the first stage of the screening process. Even though the search was limited to the titles, abstracts, and keywords of the papers many publications bore little relationship with the subject matter. Table 4 lists the ascertained precision of query q as regards the number of publications that were found to be relevant for the review. Column 'Retrieved' denotes the number of publications retrieved via q, while the column 'Relevant' signifies the number of publications left after the screening. Table 4 lists only ACM and IEEE as sources for publications because the few relevant papers retrieved via Inspec and DBLP turned out to be references to papers that were actually published at ACM and IEEE.

The reasons for individual publications being disregarded during the screening phase varied widely. Since ACM and IEEE are publishers rooted in a scientific background, all retrieved publications were peer-reviewed and all came with bibliography records that bore witness of their origin. Only two papers were dismissed because of being written in a language other than English. The most common reasons for failing to pass the screening were as follows:

Thematic scope. A considerable amount of the retrieved publications fell into the category of generic data 'visualisation,' and discussed various properties of graphical 'transparency.' In order to be found, these papers also mentioned the term 'privacy,' but were thematically unrelated to ex post-TETs. Conversely, it turned out that the majority of papers contained respective combinations of the search terms even though they originated from fields only partially related to the intended subject matter.

Generic apps. Many contemporary apps, specifically apps developed for mobile devices, are evaluated against the backdrop of usability. In the context of mobile computing, especially on platforms such as Android, the keywords 'privacy,' 'policies' and 'permissions' are sometimes used interchangeably, and the query seemed to have picked up on respective papers.

Security. Many publications dealt primarily with security issues, whereas the topics privacy and usability were only secondary aspects of the respective im-

plementation. A relatively large number of alternative approaches for access control infrastructures fell into this category.

Abstraction. Many publications were actually related to the subject matter, but introduced models of privacy-preserving infrastructures rather than concrete implementations. Such papers were not considered because they failed to meet the criteria specified in Section 5. It should be noted that some publications were entirely unspecific as regards the maturity of the implementations they discussed.

Ex ante. Many disregarded publications dealt with implementations that conceptually classified as ex ante-TETs, aiding users in their decision making process *before* they disclose any personal data. Some of these papers featured extensive user studies, both in terms of the test subjects' general understanding of privacy threats, as well as the comprehensibility of the visualisation of the respective implementation. Still, such papers were disregarded due to being out of the scope of this review.

10 Snowballing

Following the references of relevant publications backwards and forwards, a process sometimes referred to as 'snowballing,' serves the purpose of retrieving additional papers that are contextually or thematically related to the original document. In this review, snowballing as a means to retrieve additional publications was used after the results of the initial phase were available.

Abiding by the guidelines by Kitchenham et al. [19], Scopus was used to trace the references and citations.[9] Due to temporal constraints, only a single iteration of snowballing was conducted in both directions. When snowballing was performed on 2016-12-23, Scopus had no information on [29], presumably because the publication was relatively new at that date. As a result, neither references nor citations of this paper could be determined, and snowballing was performed only for the remaining papers.

The screening of the retrieved references and citations yielded a relatively small amount of additional results. However, this process did manage to retrieve several relevant papers published by ACM and IEEE that were missing in the initial set of retrieved publications because the search terms used for the query had not picked up on them. It also yielded a small amount of relevant publications from Springer.

11 Aggregation

During the aggregation phase, any duplicates retrieved via multiple databases were removed. Moreover, series of articles that described subsequent stages of

[9] Scopus calls backward references 'references,' whereas forward references are called 'citations.' Citations refer to the papers that cite the original paper, i. e. list the original paper as a reference.

Date	Stage	Quantity
2016-11-15	Querying	899
2016-12-19	Screening	12
2016-12-23	Extracting references	303
2016-12-23	Extracting citations	53
2017-01-02	Screening of references	6
2017-01-10	Screening of citations	3
2017-01-12	Aggregation	21

Table 5. Number of publications at the end of various stages of the review process.

the development of the same TET over a larger period of time were aggregated into the latest publication of the series. In such cases, the respective prototype gradually matured as regards functionality and usability, sometimes over the period of several years. For example, the TET 'Data Track' developed at Karlstad University underwent several changes throughout its published iterations, all of which are documented in Fischer-Hübner et al. [9].

The first aggregation phase coalesced relevant papers that successfully passed the initial screening after the database query. The second aggregation phase did the same for references and citations that were considered relevant after their respective screening. All papers that passed the two aggregation phases denoted either unique implementations, or described derivatives of implementations that stood out as distinctive products of their own. For example, the implementations of [18], [30] and [35] all relate to the same underlying technology, 'PeopleFinder,' but all three extend it with different distinctive features. The latter two publications also conduct independent user studies in order to evaluate different aspects of user behavior.

The number of publications extracted during various stages of the review process is documented in Table 5. The indicated quantities denote the number of publications yielded after the operations of the respective stage had been concluded. The indicated dates are the finishing dates of the respective operation, that is, they represent the date at which the indicated number of publications were available for the subsequent process step.

12 Review taxonomy

Based on Kitchenham et al.'s [19] recommendation for the quality assessment of reviewed literature, the reviewers laid down a taxonomy that was used to classify the publications included in the final result set. Unlike Kitchenham et al., however, the reviewers relied on a qualitative assessment rather than a quantitative scoring system. The taxonomy was chosen to provide a scheme for comparing individual publications rather than inferring a quantitative statement regarding the publication's absolute value for academic research.

In cases where not all criteria were fully applicable, the respective criterion is listed nonetheless, and a substitutional remark is provided instead of the actual

Year	Authors	Title (abbreviated)
2007	Hsieh, Tang, Low, et al. [15]	Field deployment of IMBuddy...
2008	Abdullah, Conti, Beyah [1]	A Visualization Framework...
2008	Kelley, Drielsma, et al. [18]	User-controllable Learning...
2009	Kolter, Kernchen, Pernul [20]	Collaborative Privacy...
2009	Sadeh, Hong, Cranor, et al. [30]	Understanding and capturing...
2009	Tsai, Kelley, Drielsma, et al. [35]	The impact of feedback...
2010	Kolter, Netter, Pernul [21]	Visualizing Past Personal Data...
2010	Toch, Cranshaw, et al. [33]	Empirical models of privacy...
2011	Schlegel, Kapadia, Lee [31]	Eying Your Exposure...
2012	Trabelsi, Sendor [34]	Sticky policies for data control...
2013	Balebako, Jung, Lu, et al. [2]	Little Brothers Watching You...
2013	Bilogrevic, Huguenin, et al. [4]	Adaptive Information-sharing...
2013	Louw, von Solms [22]	Personally Identifiable...
2013	Biswas, Aad, Perrucci [5]	Privacy Panel...
2013	Zavou, Pappas, Kemerlis [39]	Cloudopsy: An autopsy...
2014	Mun, Kim, Shilton, et al. [24]	PDVLoc...
2015	Xu, Zhu [38]	SemaDroid...
2015	Pistoia, Tripp, Centonze, et al. [27]	Labyrinth...
2016	Bier, Kühne, Beyerer [3]	PrivacyInsight...
2016	Fischer-Hübner, Angulo, et al. [9]	Transparency, Privacy...
2016	Riederer, Echickson, et al. [29]	FindYou...

Table 6. List of reviewed publications.

assessment. For example, in case a publication did not include a user study that fact is reported instead of indicating the exact nature of the study.

A compact list of all publications discussed in this review is presented in Table 6. This listing, as well as the summaries provided throughout the appendices, are given in order of the date of publishing. In case only the year of a publication but not its month could be extracted from its bibliography record, and several papers were published in the same year, such papers are sorted alphabetically using the first author's name as the key.

Type of content. This property summarises the nature of the findings of the publication and the form of its actual scientific contribution. Examples of possible types are experiments, usability tests, or user studies. More than one type per publication is possible.

Description of aims. This property assesses whether the publication clearly and comprehensibly described the goals of the publication.

Description of context. This property assesses whether the authors clarify the academic, social, or technical context, in which the respective TET is supposed to be deployed, and which affect or outcome they hope to achieve.

Research method. Assesses whether the selected research method is deemed appropriate for the proposed solution.

Statement of findings. Assesses whether the results of the research described in the paper are presented in a clear and concise way. This property is

not tied to a certain nomenclature or formal structural sections, such as a Discussion or Conclusion section, since the structure of the publications can vary depending on the guidelines of the respective publisher and publishing context. The property instead assesses whether the authors discussed their results at the end of the paper, thereby recapitulating their findings.

Value for research or practise. Assesses whether the findings or the implementation are of any value for either further research or practical application, yielding the keywords 'theoretical,' or 'practical,' respectively.

Recruitment strategy. In papers that included user studies, this property states whether the authors described the recruitment process in such a way that the composition of the pool of test persons is transparent to the reader.

Control group. In case of user studies or empirical experiments, this property assesses whether a control group was evaluated.

Search step. This property indicates the search step during which the respective publication was retrieved. It gives evidence of whether the paper could be found using the search terms specified in Section 8. The three possible values are 'initial search,' 'forward snowballing,' or 'backward snowballing.' If a publication was retrieved during multiple steps, 'initial search' takes precedence over the two secondary search phases.

In the appendix of this review, each publication of the final result set (Table 6) is listed along with its bibliography data, a brief summary of its content, and a classification of the criteria described above.

13 Final remarks

This technical report documented the process of a systematic literature review conducted on usable ex post-TETs. It briefly touched upon the thematic context and specified the contextual scope as well as the selection criteria for prospective articles. The report elaborated on the methodology used to conduct the literature review and provided a big picture of the various processing steps required to yield a set of articles that meet the specified criteria. It motivated the selection of publication databases and search terms used to make the queries. It documented the screening of the retrieved articles, as well as their aggregation into a final set of papers that sufficed the formal selection criteria. Each paper of the final set of articles was summarised as regards its content, and categorised according to a taxonomy that underpinned various characteristics of the publication.

The literature review documented in this report is non-exhaustive, but systematic. It was non-exhaustive in that it started out with a finite set of databases and search terms, and in that it traced the references of the screened publications only one generation backwards and forwards. It was systematic in that its methodology was rigorous, and that it abided by clearly defined selection criteria for individual articles.

Appendix

A.1 Hsieh, Tang, Low, et al. (2007)

Bibliography

Title. *Field deployment of IMBuddy: A study of privacy control and feedback mechanisms for contextual IM* [15]
Authors. Hsieh, G.; Tang, K. P.; Low, W. Y.; Hong, J. I.
Journal. Lecture Notes in Computer Science
Year. 2007
Pages. 91–108

Summary

The authors present 'IMBuddy,' an instant messaging system that allows users not only to exchange personal messages, but also to share personal data in form of location data with selected individuals called buddies. On the one hand, IMBuddy allows users to specify their privacy preferences in terms of interruptibility for incoming messages and for sharing of location data. On the other, it allows users to request the location of other target users, e. g. in order to check for these users' immediate availability.

The paper forcues on the evaluation of various types of system feedback that enables users to achieve a higher level of transparency about the circumstances under which their buddies are requesting access to their data.

Two user studies are being conducted to evaluate the usability of IMBuddy, as well as the status awareness and personal comfort on part of the users of the service. The first study attests a high level of satisfaction on part of the participants. The test persons' statements on usability is carried over to serve as a basis for the design of the second generation of the prototype. The findings of the second study complement its predecessor's. In particular, multiple types of feedback that feed status information back to target persons about them being traced are appreciated by the test subjects.

Classification

Type. Implementation of client and server prototype, two user studies.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.

Control group. There were no control groups as such, but the lessons learned of the first user study were carried over to the design of the second, approving the findings of the previous study.
Search step. Backward snowballing.

A.2 Abdullah, Conti, Beyah (2008)

Bibliography

Title. *A Visualization Framework for Self-Monitoring of Web-Based Information Disclosure* [1]
Authors. Abdullah, K.; Conti, G.; Beyah, R.
Conference. IEEE International Conference on Communications 2008
Year. 2008
Pages. 1700–1707
DOI. 10.1109/ICC.2008.328
ISSN. 1550-3607

Summary

The authors develop visualisation techniques that allow various stakeholders, such as individuals, companies, or organisations, to self-monitor their network activities. The intended outcome of such monitoring is to enable the stakeholders to better understand the consequences of their doing, and to enable them to make better-informed decisions as regards planning future activities.

The authors discuss several possible scenarios that might lead to personal data being disclosed, and focus on query search strings of search engines. They propose (i) the user's browser cache, (ii) live tracking of network traffic via a browser plug-in, and (iii) a large set of search queries recently disclosed by AOL as possible sources for data analysis, discussing the respective pros and cons of each source.

The publication further discusses three specific usage contexts: (1) A time-sequence of personal information disclosure, (2) sorting similarly categorised information grouped by content and service provider, and (3) monitoring of the most frequently used search terms.

As regards visualising the data, the authors evaluate several different forms of presentation: (a) A search histogram depicting the user's activity over time, (b) a bubble chart for the quantitative analysis of the non-temporal aspects of online searching, (c) a 'file tree view' that allows for an hierarchical overview of the categories and sub categories of the topics that represent the user's personal interest, and (d) a visualisation of search strings based on Seesoft.

The authors conduct a user study in order to find out about the users' personal preferences as regards the visualisation in the context of specific use cases. Based on the results of this study, the authors create a prototype in form of a Firefox plug-in that implements the file tree view. They conduct a second study, asking users for feedback with regard to the perceived usefulness of the plug-in, as well as suggestions for improving it.

Classification

Type. Implementation of client-side prototype, two user studies.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Stated for the first study.
Control group. There were no control groups as such, but the findings of the first user study carried over to the implementation of a plug-in that, in turn, was evaluated.
Search step. Initial search.

A.3 Kelley, Drielsma, Sadeh, et al. (2008)

Bibliography

Title. *User-controllable Learning of Security and Privacy Policies* [18]
Authors. Kelley, P. G.; Drielsma, P. H.; Sadeh, N.; Cranor, L. F.
Conference. AISec '08
Year. 2008
DOI. 10.1145/1456377.1456380
ISBN. 978-1-60558-291-7

Summary

The authors introduce a machine learning model whose accuracy increases gradually by relying on incremental user feedback. The users' feedback, in turn, is based on the decisions made by the model on behalf of each user as it tries to accommodate a particular user's privacy preferences in the context of a location sharing platform. The feedback loop incrementally trains the model towards the actual preferences of each user.

The model is instrumented and tested against the backdrop of the participatory network 'PeopleFinder,' which allows users to share their locations with other participants of the service. The users' privacy preferences for sharing their location data are based on their personal relationship with a respective target user, as well as the current time. When the system makes a decision on behalf of a user whether to share that user's location, that decision is carried out and logged. It can be audited by the user at any time. By rating the decisions made by the system these users help gauging the quality of the original decision, thereby improving future decisions by relying on increasingly accurate predictions of the target user's preferences.

The authors state that the accuracy of the system's decision making process increases considerably over time.

Classification

Type. Implementation of a machine learning model, and a server-sided rating platform.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Stated.
Value. Practical, theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. No actual user study was conducted, but the feedback of the test subjects was fed back to the system, proving that the model can be trained to accommodate the users' actual preferences.

A.4 Kolter, Kernchen, Pernul (2009)

Bibliography

Title. *Collaborative Privacy – A Community-Based Privacy Infrastructure* [20]
Authors. Kolter, J.; Kernchen, T.; Pernul, G.
Journal. Emerging Challenges for Security, Privacy and Trust
Year. 2009
Pages. 226–236

Summary

The authors describe an architecture in form of a public information sharing platform that supports users in making informed decisions about which online service provider to choose based on publicly available information about various services. Respective recommendations and comments are provided by the users of the platform themselves.

The information and privacy policy sharing platform described by the authors qualifies as an ex ante-TET, and as such is not suitable for a review dealing exclusively with ex post-TETs. However, the TET also comprehends a 'Data Disclosure Log' on the client side whose purpose is to locally record data flows to the service providers that a user is relying on. These logs can be audited and used by the user to form a personal opinion on the trustworthiness of the respective service provider. The knowledge thus gained may then be shared publicly on the information sharing platform.

Classification

Type. Implementation of client and information sharing platform.
Description of aims. Yes.
Description of context. Yes.

Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Backward snowballing.

A.5 Sadeh, Hong, Cranor, et al. (2009)

Bibliography

Title. *Understanding and capturing people's privacy policies in a mobile social networking application* [30]
Authors. Sadeh, N.; Hong, J.; Cranor, L.; Fette, I.; Kelley, P.; Prabaker, M.; Rao, J.
Journal. Personal and Ubiquitous Computing
Year. 2009
Pages. 401–412
DOI. 10.1007/s00779-008-0214-3

Summary

The authors introduce 'PeopleFinder,' an online location tracking service that allows its users to retrieve other users' location, while permitting a specific target audience access to their own location. Permission to accessing a user's location is based on an opt-in choice, specified in detail by the personal privacy policy of the target user.

Access to a user's location can be refined and audited by that user at any time. A machine learning algorithm tries to predict respective adaption of the user's current policy based on ongoing decisions made during the use of the service, and offers hints as to adapt one's settings accordingly.

Lab experiments and a field study are conducted to better understand how people act and behave in a socio-technical location sharing environment. The authors state that the predictions of how users should share their data indicates a relatively high accuracy on part of the machine learning algorithm. According to the test subjects, the server-sided recommendation is considered useful in the context of their own decision making.

Classification

Type. Implementation of client and server prototype, lab and field studies.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.
Control group. n/a.
Search step. Backward snowballing.

A.6 Tsai, Kelley, Drielsma, et al. (2009)

Bibliography

Title. *Who's viewed you? The impact of feedback in a mobile location-sharing application* [35]
Authors. Tsai, J. Y.; Kelley, P.; Drielsma, P. H.; Cranor, L. F.; Hong, J.; Sadeh, N.
Conference. Conference on Human Factors in Computing Systems, Proceedings
Year. 2009
Pages. 2003–2012
DOI. 10.1145/1518701.1519005

Summary

The project 'Locyoution' builds on the 'PeopleFinder'-technology also employed by [18] and [30]. However, rather than being a stand-alone application, it relies on Facebook as the underlying online social network to leverage its user base.

The authors perform a user study in order to find out if and to what extend feedback changes the decision making process of users of a location sharing platform. To this end, the user study assigns half of its participants to a group whose members receive feedback on their sharing behavior, while the other half is assigned to a control group whose members don't receive any feedback.

Feedback is given automatically in form of an auditing log of all location requests along with status information about whether a particular request for accessing the user's location was granted based on the target user's privacy policy. The result of the study shows that users appreciate feedback from the system, increasing their level of comfort and encouraging them to share their location.

Classification

Type. Implementation of client and server side prototype, user study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.
Control group. Yes, studied to determine the impact of feedback.
Search step. Backward snowballing.

A.7 Kolter, Netter, Pernul (2010)

Bibliography

Title. *Visualizing Past Personal Data Disclosures* [21]
Authors. Kolter, J. and Netter, M. and Pernul, G.
Conference. ARES '10: Availability, Reliability, and Security.
Year. 2010
Pages. 131–139
DOI. 10.1109/ARES.2010.51

Summary

The authors introduce a 'Data Disclosure Log' architecture that records either deliberately or accidentally disclosed personal data, and allows for visualising the results. In that regard, the implementation provides an overview of a user's past transactions that involve the disclosure of personal data to a communication partner. The goal of the presented TET is to enable a user to perform an ex post-analysis of actions performed in the past.

The authors implement a prototype in form of a Firefox plug-in that parses outgoing network traffic, and scans the data for attribute values that are related to personal data. Logs of respective transactions along with contextually related meta data are being kept. Meta data comprehends the identity of the communication partner or service provider, the current time of day, as well as the modality or context of the transaction, such as a registration process, login, or purchase.

The introduced architecture consists of four major parts: (1) A module for data import, (2) a database, (3) a visualisation component, and (4) an export module.

The authors emphasise that various usability principles have been applied when the prototype was implemented. Visualisation of the collected data is implemented based on multiple forms of presentation: (a) The service provider overview, (b) a chronological view of activity performed over time, and (c) a hierarchical graph view that stacks multiple entities on top of each other.

The implementation is evaluated via a user study. The authors state that the test persons appreciated the usefulness of the tool.

Classification

Type. Model of the architecture, implementation of a prototype, user study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.
Control group. n/a.
Search step. Initial search.

A.8 Toch, Cranshaw, Drielsma, et al. (2010)

Bibliography

Title. *Empirical models of privacy in location sharing* [33]
Authors. Toch, E.; Cranshaw, J.; Drielsma, P. H.; Tsai, J. Y.; Kelley, P. G.; Springfield, J.; Cranor, L.; Hong, J.; Sadeh, N.
Conference. Proc. of ACM UbiComp
Year. 2010

Summary

The authors present 'Locaccino,' an add-on for Facebook that allows users to share their location among selected individuals and groups.

Locaccino consists of a sensing application for Symbian OS-devices and laptops on the client-side, and a web interface to the Facebook add-on on the server-side. The client-side app allows for the actual sharing of the users' locations. The web interface allows users to specify their privacy preferences and to review the history of location requests by other users.

The authors conduct an extended user study to evaluate the behavior of the participants, specifically their willingness to share their personal location data related to the nature of the place they found themselves in. The study provides insight as regards the users' sharing behavior in relation to the entropy of places, and the general busyness of users in terms of their general mobility. The authors state that most users feel comfortable using the location sharing infrastructure.

Classification

Type. Implementation of client-side prototypes for two platforms, server-side Facebook add-on, field study, pre- and post-study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.
Control group. n/a.
Search step. Backward snowballing.

A.9 Schlegel, Kapadia, Lee (2011)

Title. *Eyeing Your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy* [31]
Authors. Schlegel, Roman and Kapadia, Apu and Lee, Adam J.
Conference. SOUPS '11
Year. 2011
Pages. 14:1–14:14
DOI. 10.1145/2078827.2078846
ISBN. 978-1-4503-0911-0

Summary

The paper describes a participatory sharing infrastructure that allows users to share personal information with other users based on their relationship with these users in terms of their respective role as members of distinctive social groups. The authors present a prototype of an Android app that allows users to do two things: (1) Manage their privacy preferences with regard to their group affiliations and the time of day. (2) Visualise to what extend queries of current their status from the designated parties actually take place.

Requests are visualised on the target user's device via an intuitive display that uses a pair of gazing eyes as a metaphor for surveillance, the size of the eyes changing in relation to the number of requests. The purpose of this somewhat unusual display is to provide users with an ambient, semantically enhanced feedback about queries, and to methodically signal that such queries have an impact on their privacy.

The authors conduct two user studies, a preliminary and a main study, both of which confirm the tool's effectiveness of conveying the intended information. Based on the combined results of both studies, the researchers are able to identify certain patterns regarding the users' level of comfort in relation to the amount of requests per time unit, and the affiliation of the requesting party. Most test subjects also state that they consider the functionality of the app useful.

Classification

Type. Implementation of client-side and server-side prototypes, two user studies.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.
Control group. Yes. Members of the control group received more detailed feedback than members of the regular group. This allowed for a comparison regarding the awareness and sharing behavior of the members of both groups.
Search step. Initial search.

A.10 Trabelsi, Sendor (2012)

Bibliography

Title. *Sticky policies for data control in the cloud* [34]
Authors. Trabelsi, S.; Sendor, J.
Conference. Privacy, Security and Trust (PST) 2012
Year. 2012
Pages. 75–80
DOI. 10.1109/PST.2012.6297922

Summary

The authors introduce an access control and management abstraction layer that assists users in maintaining and auditing requests to their personal data in one central place, called "Sticky Policy Access Control Engine" (SPACE). SPACE stores all sensitive personal data in a 'Safe Zone' that allows data subjects to control and maintain the access policies to their data, and third parties to access it, provided they have been permitted by the data subject to do so.

When the personal data of a data subject, called data owner, are accessed by a third party, that event is propagated, and can be visualised by the owner. Moreover, data owners can visualise access to their data at any later date in order to review and audit previous data requests.

The authors provide multiple forms of presentation to visualise access to data, specifically in form of (1) a hierarchical 'Access Control Map,' (2) a temporal 'History Map,' and (3) a 'Data Geolocation Map' that provides insight in the location of a requesting party.

Since no user study was conducted, it is unclear which degree of maturity the discussed prototype actually has.

Classification

Type. Implementation of client-side and server-side prototype.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical (?), theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Initial search.

A.11 Balebako, Jung, Lu, et al. (2013)

Bibliography

Title. *"Little Brothers Watching You": Raising Awareness of Data Leaks on Smartphones* [2]
Authors. Balebako, Rebecca; Jung, Jaeyeon; Lu, Wei; Cranor, Lorrie Faith; Nguyen, Carolyn
Conference. SOUPS '13
Year. 2013
Pages. 12:1–12:11
DOI. 10.1145/2501604.2501616
ISBN. 978-1-4503-2319-2

Summary

The authors present the results of a user study based on the evaluation of 'Privacy Leaks,' a prototype for the Android platform. The purpose of the app is to raise the user's awareness about data leakage as regards personal data being leaked by apps installed on the device.

For this purpose, Privacy Leaks offers two modes of operation to notify users about data leakage. The 'just-in-time' mode notifies users in real time, that is, at the time the device is actually leaking the data to a remote communication partner, such as an online service. That notification manifests in form of a textual message conveyed via Android's notification system, as well as a momentary vibration of the device.

The second mode allows users to examine the number of leakages retrospectively. Privacy Leaks features a tabular grid layout that maps the type of data being leaked to the respective app that leaked the data, using color schemes to denote the severity of the disclosure based on the total number of leakages that occurred.

During the extensive user study, the test subjects confirm that Privacy Leaks raises their awareness regarding the leakage of their personal data. As regards usability, the authors discuss several points that the participants raised during the study.

Classification

Type. Implementation of client-side prototype, user study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Clearly stated.
Control group. n/a.
Search step. Initial search.

A.12 Bilogrevic, Huguenin, Agir, et al. (2013)

Bibliography

Title. *Adaptive Information-sharing for Privacy-aware Mobile Social Networks* [4]
Authors. Bilogrevic, Igor; Huguenin, Kévin; Agir, Berker; Jadliwala, Murtuza; Hubaux, Jean-Pierre
Conference. International Joint Conference on Pervasive and Ubiquitous Computing 2013
Year. 2013
Pages. 657–666
DOI. 10.1145/2493432.2493510
ISBN. 978-1-4503-1770-2

Summary

The authors introduce 'SPISM,' a client-server architecture for sharing personal data among the participants of a mobile online network. The kind of personal data considered in the paper is the users' current location and their availability at certain dates.

Users act either in their role as requesters of an other user's data, or as persons targeted by such requests. A targeted user's mobile device takes into account various factors associated with her or his current circumstances, including but not limited to the current time, location, the target's relationship with the requesting user, as well as decisions made previously as to whether and how personal data was shared in the past. These parameters act as an input vector for a trained machine learning classifier that then provides a recommendation of whether and with what confidence the request should be answered.

The results of the user study conducted by the authors suggest that the predictions made by the classifier used for SPISM allows for high accuracy.

Classification

Type. Implementation prototype of mobile app and directory server, user study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. Stated. The user study served the purpose of evaluating the machine learning classifier, not the usability of the app.
Control group. n/a.
Search step. Initial search.

A.13 Biswas, Aad, Perrucci (2013)

Bibliography

Title. *Privacy Panel: Usable and Quantifiable Mobile Privacy* [5]
Authors. Biswas, D. and Aad, I. and Perrucci, G. P.
Conference. Availability, Reliability and Security (ARES) 2013
Year. 2013
Pages. 218–223
DOI. 10.1109/ARES.2013.29

Summary

The authors present 'Privacy Panel (PP),' a client-side infrastructure for mobile devices that tracks other applications' disclosure of the user's personal data.

For that purpose, PP defines three modalities that installed apps might disclose personal data about: (1) location (the physical geolocation of the device), (2) contacts (based on the user's address book), and (3) contents (files and media stored in the file system of the device). The authors' goal is to provide a metric to quantify the disclosure of personal data via any of these modalities. Using quantified values, PP visualises the impact on the user's privacy in form of readable, meaningful ratings.

In case of location, the rating is based on the access frequency and precision of the localisation, as well as the type and nature of the current environment. For the user's contacts, the level of intrusiveness depends on the access frequency and the number of details disclosed by an app. The impact of disclosed file system contents is based on the access frequency, as well as the type and ownership of the disclosed data or media file, respectively.

In order to provide this functionality, PP intercepts respective library calls when apps access the three modalities during run-time, and logs the event for later auditing. Users can set threshold values for all three modalities, and have themselves notified by PP should these thresholds be exceeded.

The authors discuss the impact on the CPU-load of the device as regards the overhead caused by running PP.

Classification

Type. Implementation of client-side prototype.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical, theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Initial search.

A.14 Louw, von Solms (2013)

Bibliography

Title. *Personally Identifiable Information Leakage Through Online Social Networks* [22]
Authors. Louw, Candice; von Solms, Sebastiaan
Conference. South African Institute for Computer Scientists and Information Technologists Conference (SAICSIT '13)
Year. 2013
Pages. 68–71
DOI. 10.1145/2513456.2513467
ISBN. 978-1-4503-2112-9

Summary

The authors introduce a model and tool to visualise personal information published on the online social network Facebook. The visualisation deals with personal data explicitly published by a user, as well as transitive information that can be derived from other information, such as a user's gender being inferred from his or her name or profile photo.

In terms of possible recipients of that data, friends, friends of friends, and the general public are considered as stakeholders.

Classification

Type. Implementation of a 'prototype software model' [22]. It remains unclear whether an usable implementation actually exists.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Stated.
Value. Theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Initial search.

A.15 Zavou, Pappas, Kemerlis, et al. (2013)

Bibliography

Title. *Cloudopsy: An autopsy of data flows in the cloud* [39]
Authors. Zavou, A.; Pappas, V.; Kemerlis, V. P.; Polychronakis, M.; Portokalidis, G.; Keromytis, A. D.
Journal. Lecture Notes in Computer Science (LNCS 8030)
Year. 2013
Pages. 366–375
DOI. 10.1007/978-3-642-39345-7-39

Summary

The authors introduce 'Cloudospy,' a cloud-based service intended to increase the transparency with regard to interactions and data flows between multiple cloud services. The TET allows for data auditing and visualisation of logged transactions for end users as well as service providers who provide services based on the infrastructure of a third party cloud service provider.

For that purpose, Cloudospy, which is hosted by a central cloud service provider, keeps extensive append-only logs of all transactions between two or more service providers that occur as the result of an action or request triggered by a particular end user.

Cloudospy provides two modes of visualising data: For the end user, it displays data flows of his or her personal data within the administrative domain of the cloud service provider. For the online service provider, it provides an overview of the data flows of all customers.

Classification

Type. Implementation of third party-sided (central cloud service) prototype.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Backward snowballing.

A.16 Mun, Kim, Shilton, et al. (2014)

Bibliography

Title. *PDVLoc: A personal data vault for controlled location data sharing* [24]
Authors. Mun, M. Y.; Kim, D. H.; Shilton, K.; Estrin, D.; Hansen, M.; Govindan, R.
Journal. ACM Transactions on Sensor Networks, volume 10, number 4
Year. 2014
DOI. 10.1145/2523820

Summary

The paper discusses the infrastructure of a personal data vault (PDV), an intermediate service between an user on the one hand, and a requesting party on the other.

The user acts as a continuous producer of personal data, such as location or health data, and subsequently acts as the owner of that data. The data are stored in a PDV, which is provided by a trusted third party. Read access to the data on part of requesters is manageable and auditable only by the owner of the data.

The prototype aids users in three ways: (1) The Granular Access Control Lists (ACL) supports users in setting up and reviewing the appropriate granularity necessary for each requester. (2) The 'Rule-Recommender' recommends refined settings for the ACL based on heuristics found in existing data. (3) The 'Trace-Audit' allows auditors to review the actual access to the data.

A user study shows that users feel encouraged to rely on the support of PDVoc, and would use it in the future.

Classification

Type. Implementation prototype, performance evaluation, user study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Theoretical.
Recruitment strategy. Stated.
Control group. n/a.
Search step. Forward snowballing.

A.17 Xu, Zhu (2015)

Bibliography

Title. *SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones* [38]
Authors. Xu, Zhi; Zhu, Sencun
Conference. 5th ACM Conference on Data and Application Security and Privacy (CODASPY '15)
Year. 2015
Pages. 61–72
DOI. 10.1145/2699026.2699114
ISBN. 978-1-4503-3191-3

Summary

The authors introduce 'SemaDroid,' a framework for Android that allows for managing access to the sensors of a mobile device on a fine grained basis.

To this end, SemaDroid intercepts system calls of third party apps via the sensor-API, and redirects them to its own sensor manager. Depending on the user's policy, that manager can either permit or deny the call. Optionally, it can also channel faked or sanitised data back to the requesting app.

Apart from being able to specify fine-grained policies for individual sensors and apps, users are enabled to review the actual access of the sensors. The reviewed information is available in textual form, including the exact times, durations, and accuracy used for providing the sensor data.

Classification

Type. Implementation of a client-side prototype.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Stated.
Value. Practical, theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Initial search.

A.18 Pistoia, Tripp, Centonze, et al. (2015)

Bibliography

Title. *Labyrinth: Visually Configurable Data-Leakage Detection in Mobile Applications* [27]
Authors. Pistoia, M.; Tripp, O.; Centonze, P.; Ligman, J. W.
Conference. 16th IEEE International Conference on Mobile Data Management
Year. 2015
Pages. 279–286
DOI. 10.1109/MDM.2015.69
ISBN. 1551-6245

Summary

In order to counter the uncertainty on part of users as to personal data being disclosed unintentionally by an app installed on a mobile device, the authors introduce the TET 'Labyrinth.'

Labyrinth scrutinises one Android or iOS application at a time, and offers a visual configuration UI that allows users to change their privacy settings. A man-in-the-middle framework captures data that is being exchanged with remote servers, and logs these events locally for later review. Labyrinth relies on an analyser that parses and gauges transmitted data, recognising patterns that might relate to personal data.

The authors state that Labyrinth causes only little overhead as regards system resources. The TET is portable across the major mobile platforms, and features a simple configuration via a GUI that is supposed to be usable even by non-developers. Content-wise, the publication focuses mostly on the technical aspects of the implementation.

Classification

Type. Implementation of client-side prototype.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Stated.
Value. Theoretical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Initial search.

A.19 Bier, Kühne, Beyerer (2016)

Bibliography

Title. *PrivacyInsight: The next generation privacy dashboard* [3]
Authors. Bier, C.; Kühne, K.; Beyerer, J.
Journal. Lecture Notes in Computer Science (9857 LNCS)
Year. 2016
DOI. 10.1007/978-3-319-44760-5-9

Summary

Based on the jurisdiction of the GDPR [8], the authors define formal requirements for the implementation of a centralised software framework called 'PrivacyInsight' that provides transparency for the data subject as regards the auditing process of disclosed personal data. PrivacyInsight allows for visualising personal data, and for exercising the data subject's rights to rectify and erase personal data.

The visualisation considers the actual data, as well as the data flows between the stakeholders that were involved in the process of storing and processing the user's personal data, including subordinate organisational units within an organisation or company.

The authors conduct a user study to evaluate PrivacyInsight's usability by comparing it with two other TETs that allow for an ex post-analysis of personal data: GenomSynlig, an online visualisation TET created by Karlstad University, and a JSON-document provided by a service provider. The JSON-document comprehends all personal data gathered by the service provider about the data subject, and represents that user's explicit request for a transcript. The study shows that most users favour the use of PrivacyInsight over the two other alternatives.

PrivacyInsight allows data subjects to exercise their legal rights to rectify their personal data, and to have it erased by the data controller.

Classification

Type. Implementation of prototype, user study.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Theoretical.
Recruitment strategy. Clearly stated.
Control group. None, but the prototype was compared with GenomSynlig, as well as a JSON-file.
Search step. Forward snowballing.

A.20 Fischer-Hübner, Angulo, Karegar, et al. (2016)

Bibliography

Title. *Transparency, privacy and trust – Technology for tracking and controlling my data disclosures: Does this work?* [9]
Authors. Fischer-Hübner, S. and Angulo, J. and Karegar, F. and Pulls, T.
Journal. IFIP Advances in Information and Communication Technology
Year. 2016
Pages. 3–14
DOI. 10.1007/978-3-319-41354-9-1

Summary

The latest paper of a series that documents multiple iterations of 'Data Track,' an ex post-TET developed at Karlstad University ([26], [11], [10]), the publication discusses the technical and legal foundations that Data Track is built upon, and summarises the latest findings with regard to the development of the TET.

Data Track allows users to analyse and visualise their personal data downloaded from online social networks such as Facebook. For that purpose, users are guided through the process of exporting and downloading their data from the service, and subsequently import and visualise it via Data Track.

Data Track allows for three different modes of visualising personal data: (1) The 'Trace view' illustrates the entities involved in storing and processing the user's personal data. In addition to the communication partner's identity, the trace view offers details about the actual information being stored by that entity. It also allows users to rectify and delete such data. (2) The 'Timeline view' lists events that denote the processing of the user's personal data. Contents are listed in temporal order, and allow for auditing time-based events of the past. (3) The 'Map view' visualises location-based services on a navigable geographical map, allowing users to pinpoint past events to specific locations.

Throughout its iterations, Data Track underwent several user studies and usability tests whose purpose was to ascertain its usability.

Classification

Type. Implementation of a client-side prototype, several user studies.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Clearly stated.
Value. Practical.
Recruitment strategy. Clearly stated.
Control group. n/a.
Search step. The reviewed version was retrieved via forward snowballing, while the referenced version was retrieved during the initial search.

A.21 Riederer, Echickson, Huang, et al. (2016)

Bibliography

Title. *FindYou: A Personal Location Privacy Auditing Tool* [29]
Authors. Riederer, Christopher; Echickson, Daniel; Huang, Stephanie; Chaintreau, Augustin
Conference. 25th International Conference Companion on World Wide Web
Year. 2016
Pages. 243–246
DOI. 10.1145/2872518.2890546
ISBN. 978-1-4503-4144-8

Summary

The authors briefly present 'FindYou,' a site that allows users to upload their personal data that they have downloaded from the social online networks FourSquare, Instagram, and Twitter. Alternatively, users can enter their data manually.

One the one hand, FindYou extracts location data and maps it to census tracts, that is, partitions assigned by the United States Census. On the other hand, it tries to infer additional secondary information from the original location data, and tries to predict some of the user's ethnographical traits.

The purpose of the tool is two-fold: First, users should be enabled to learn what kind of information might be inferred from the data traces they left behind in online social networks. Second, users can anonymously donate their data for scientific research.

Classification

Type. Implementation of web site with individual feedback.
Description of aims. Yes.
Description of context. Yes.
Research method. Appropriate.
Findings. Partially stated.
Value. Practical.
Recruitment strategy. n/a.
Control group. n/a.
Search step. Initial search.

Glossary

GDPR. General Data Protection Regulation [8].
ISO. International Standards Organization.
Personal data. Refers to personally identifiable information.
PET. Privacy enhancing technology.
Privacy. Refers to the term 'information privacy' or 'data privacy' as discussed by [14].
TET. Transparency enhancing technology/tool.
Transparency. The principle as stipulated in GDPR, Chap. III, Art. 12 et seq.
Usability. Refers to the definition of usability in ISO 9241-11 [16].
User. Refers to the individual that uses a TET. That person may be the data subject whose data is being reviewed, or a legal representative or guardian.
UI. User interface.

References

1. Abdullah, K., Conti, G., Beyah, R.: A Visualization Framework for Self-Monitoring of Web-Based Information Disclosure. In: 2008 IEEE International Conference on Communications. pp. 1700–1707 (May 2008)
2. Balebako, R., Jung, J., Lu, W., Cranor, L.F., Nguyen, C.: "Little Brothers Watching You:" Raising Awareness of Data Leaks on Smartphones. In: Proceedings of the Ninth Symposium on Usable Privacy and Security. pp. 12:1–12:11. SOUPS '13, ACM, New York, NY, USA (2013), http://doi.acm.org/10.1145/2501604.2501616
3. Bier, C., Kühne, K., Beyerer, J.: PrivacyInsight: The Next Generation Privacy Dashboard. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 9857 LNCS, 135–152 (2016)
4. Bilogrevic, I., Huguenin, K., Agir, B., Jadliwala, M., Hubaux, J.P.: Adaptive Information-sharing for Privacy-aware Mobile Social Networks. In: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing. pp. 657–666. UbiComp '13, ACM, New York, NY, USA (2013), http://doi.acm.org/10.1145/2493432.2493510
5. Biswas, D., Aad, I., Perrucci, G.P.: Privacy Panel: Usable and Quantifiable Mobile Privacy. In: Availability, Reliability and Security (ARES), 2013 Eighth International Conference on. pp. 218–223 (Sept 2013)
6. Cooper, H.M.: Organizing knowledge syntheses: A taxonomy of literature reviews. Knowledge in Society 1(1), 104–126 (1988)
7. The European Parliament and the Council of the European Union: Directive 95/46/EC of the European Parliament and of the Council (October 1995)
8. The European Parliament and the Council of the European Union: Regulation (EU) 2016/679 of the European Parliament and of the Council, Art 12 (8) (April 2016)
9. Fischer-Hübner, S., Angulo, J., Karegar, F., Pulls, T.: Transparency, Privacy and Trust – Technology for Tracking and Controlling My Data Disclosures: Does This Work? In: IFIP International Conference on Trust Management. pp. 3–14. Springer (2016)

10. Fischer-Hübner, S., Angulo, J., Pulls, T.: How can cloud users be supported in deciding on, tracking and controlling how their data are used? In: IFIP PrimeLife International Summer School on Privacy and Identity Management for Life. pp. 77–92. Springer (2013)
11. Fischer-Hübner, S., Hedbom, H., Wästlund, E.: Trust and assurance hci. In: Privacy and Identity Management for Life, pp. 245–260. Springer (2011)
12. Hedbom, H.: A Survey on Transparency Tools for Enhancing Privacy, pp. 67–82. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
13. Hildebrandt, M.: Behavioural biometric profiling and transparency enhancing tools. FIDIS Deliverable 7 (2009)
14. van den Hoven, J., Blaauw, M., Pieters, W., Warnier, M.: Privacy and Information Technology. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy. Spring 2016 edn. (2016)
15. Hsieh, G., Tang, K.P., Low, W.Y., Hong, J.I.: Field deployment of IMBuddy: A study of privacy control and feedback mechanisms for contextual IM. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 4717 LNCS, 91–108 (2007)
16. International Organization for Standardization: Guidance on usability. Tech. Rep. ISO 9241-11:1998(E), ISO (March 1998)
17. Janic, M., Wijbenga, J.P., Veugen, T.: Transparency Enhancing Tools (TETs): An Overview. In: 2013 Third Workshop on Socio-Technical Aspects in Security and Trust. pp. 18–25 (June 2013)
18. Kelley, P.G., Hankes Drielsma, P., Sadeh, N., Cranor, L.F.: User-controllable Learning of Security and Privacy Policies. In: Proceedings of the 1st ACM Workshop on Workshop on AISec. pp. 11–18. AISec '08, ACM, New York, NY, USA (2008)
19. Kitchenham, B., Brereton, P.: A systematic review of systematic review process research in software engineering. Information and software technology 55(12), 2049–2075 (2013)
20. Kolter, J., Kernchen, T., Pernul, G.: Collaborative Privacy – A Community-Based Privacy Infrastructure. Emerging Challenges for Security, Privacy and Trust 297, 226–236 (2009)
21. Kolter, J., Netter, M., Pernul, G.: Visualizing Past Personal Data Disclosures. In: Availability, Reliability, and Security, 2010. ARES '10 International Conference on. pp. 131–139 (Feb 2010)
22. Louw, C., von Solms, S.: Personally Identifiable Information Leakage Through Online Social Networks. In: Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference. pp. 68–71. SAICSIT '13, ACM, New York, NY, USA (2013), http://doi.acm.org/10.1145/2513456.2513467
23. Moere, A.V.: Beyond the Tyranny of the Pixel: Exploring the Physicality of Information Visualization. In: 2008 12th International Conference Information Visualisation. pp. 469–474 (July 2008)
24. Mun, M.Y., Kim, D.H., Shilton, K., Estrin, D., Hansen, M., Govindan, R.: PDVLoc: A personal data vault for controlled location data sharing. ACM Transactions on Sensor Networks 10(4) (2014)
25. Nielsen, J.: Usability Engineering. Elsevier (1994)
26. Pettersson, J.S., Fischer-Hübner, S., Bergmann, M.: Outlining "Data Track": Privacy-friendly data maintenance for end-users. Advances in Information Systems Development 1, 215–226 (2006)

27. Pistoia, M., Tripp, O., Centonze, P., Ligman, J.W.: Labyrinth: Visually Configurable Data-Leakage Detection in Mobile Applications. In: 2015 16th IEEE International Conference on Mobile Data Management. vol. 1, pp. 279–286 (June 2015)
28. Randolph, J.J.: A guide to writing the dissertation literature review. Practical Assessment, Research & Evaluation 14(13), 1–13 (2009)
29. Riederer, C., Echickson, D., Huang, S., Chaintreau, A.: FindYou: A Personal Location Privacy Auditing Tool. In: Proceedings of the 25th International Conference Companion on World Wide Web. pp. 243–246. WWW '16 Companion, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland (2016)
30. Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M., Rao, J.: Understanding and capturing people's privacy policies in a mobile social networking application. Personal and Ubiquitous Computing 13(6), 401–412 (2009), cited By 121
31. Schlegel, R., Kapadia, A., Lee, A.J.: Eyeing Your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy. In: Proceedings of the Seventh Symposium on Usable Privacy and Security. pp. 14:1–14:14. SOUPS '11, ACM, New York, NY, USA (2011), http://doi.acm.org/10.1145/2078827.2078846
32. Shneiderman, B., Plaisant, C., Cohen, M., Jacobs, S., Elmqvist, N.: Designing the User Interface: Strategies for Effective Human-Computer Interaction. Pearson, 6th edn. (2016)
33. Toch, E., Cranshaw, J., Drielsma, P.H., Tsai, J.Y., Kelley, P.G., Springfield, J., Cranor, L., Hong, J., Sadeh, N.: Empirical models of privacy in location sharing. Proc. of ACM UbiComp (2010)
34. Trabelsi, S., Sendor, J.: Sticky policies for data control in the cloud. In: Privacy, Security and Trust (PST), 2012 Tenth Annual International Conference on. pp. 75–80 (July 2012)
35. Tsai, J.Y., Kelley, P., Drielsma, P., Cranor, L.F., Hong, J., Sadeh, N.: Who's viewed you? The impact of feedback in a mobile location-sharing application. pp. 2003–2012 (2009)
36. Vom Brocke, J., Simons, A., Niehaves, B., Riemer, K., Plattfaut, R., Cleven, A., et al.: Reconstructing the giant: On the importance of rigour in documenting the literature search process. In: ECIS. vol. 9, pp. 2206–2217 (2009)
37. Webster, J., Watson, R.T.: Analyzing the past to prepare for the future: Writing a literature review (2002)
38. Xu, Z., Zhu, S.: SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones. In: Proceedings of the 5th ACM Conference on Data and Application Security and Privacy. pp. 61–72. CODASPY '15, ACM, New York, NY, USA (2015), http://doi.acm.org/10.1145/2699026.2699114
39. Zavou, A., Pappas, V., Kemerlis, V.P., Polychronakis, M., Portokalidis, G., Keromytis, A.D.: Cloudopsy: An autopsy of data flows in the cloud. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 8030 LNCS, 366–375 (2013)

Usable Transparency Enhancing Tools

This technical report documents the procedure of a literature review conducted on usable ex post-transparency enhancing tools (TETs). The review of scientific literature serves the purpose of providing insight into the characteristics of existing implementations of usable TETs. By providing a concise summary of existing implementations, the report aims to facilitate future research on the subject matter.

ISBN 978-91-7063-804-6 (pdf)

Faculty of Health, Science and Technology

WORKING PAPER | July 2017