



HAL
open science

Privacy Risk Analysis to Enable Informed Privacy Settings

Sourya Joyee De, Daniel Le Métayer

► **To cite this version:**

Sourya Joyee De, Daniel Le Métayer. Privacy Risk Analysis to Enable Informed Privacy Settings. IWPE 2018 – 4th IEEE International Workshop on Privacy Engineering, Apr 2018, London, United Kingdom. pp.1-8. hal-01939845

HAL Id: hal-01939845

<https://hal.science/hal-01939845>

Submitted on 29 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Privacy Risk Analysis to Enable Informed Privacy Settings

Sourya Joyee De and Daniel Le Métayer

Inria, Lyon and Université de Lyon

Email:sourya-joyee.de@inria.fr, daniel.le-metayer@inria.fr

Abstract—The work described in this paper is a contribution to enhancing individual control over personal data which is promoted, *inter alia*, by the new EU General Data Protection Regulation. We propose a method to enable better informed choices of privacy settings. The method relies on a privacy risk analysis parameterized by privacy settings. The user can express his choices, visualize their impact on the privacy risks through a user-friendly interface and, if needed, decide to revise them to reduce risks to an acceptable level.

Index Terms—privacy risk analysis, harm trees, privacy settings, quantified self, fitness tracking device

I. INTRODUCTION

Users reveal a lot of personal data to various websites and service providers. Even if data controllers must, in most cases, obtain their consent before collecting their data, this consent is more a formal right than a true protection. The main reason is that data subjects do not have the time and expertise to read and understand the general terms of use or privacy policies of the data controllers.

Ideally, users' choices should be based on a clear appraisal of the risks and benefits of the available options. On the legal side, this view is supported by the EU General Data Protection Regulation (GDPR) [11], which emphasizes control over personal data¹ and states that data subjects should be made aware of the risks related to personal data processing². In this paper, we propose a method, based on privacy risk analysis, to help users understand the privacy risks that may result from their choices of privacy settings. Our work relies on a privacy risk analysis methodology proposed in [9], [10]. The core of the approach is the construction and analysis of harm trees derived from information about the system, the personal data involved, the relevant risk sources, the feared events and their impacts in terms of privacy. The methodology is extended to take into account the privacy settings of the users and analyze their impact on the likelihood of privacy harms.

To illustrate our approach, we use as a case study a quantified self application. Quantified self is chosen both because of its fast growth and for the various privacy risks that such systems may pose to their users [12], [22], [16]. Fitness tracker devices (e.g., Fitbit) allow their users to track their number

of steps, strenuous activities, heart beats and location. They also provide users different types of derived information such as sleep patterns, calories burnt or goals achieved through a comprehensive dashboard. For the rest of this work, we consider a high-level specification of a fitness tracking system inspired by existing products, but we focus on a limited subset of functionalities for the sake of conciseness.

One of the desirable features of transparency enhancing technologies (TETs) targeted at data subjects is that they should be very user-friendly, with an easy-to-understand presentation of information about the privacy implications of different actions and possible choices [15]. To address these needs, we also propose a user interface through which users can easily communicate to the service provider their preferences and visualize their impact on the likelihood of privacy harms.

We describe the preliminaries on privacy risk analysis (illustrated with our case study) in Section II, and discuss user privacy preferences in Section III. In Section IV, we design a user-friendly, interactive interface that enables users to define their privacy settings and understand the resulting privacy risks. In Section V, we “lift the hood” and present the engine used to compute privacy risks. Finally, we discuss related works in Section VI and conclude with perspectives in Section VII.

II. PRELIMINARIES

In this section, we present the terminology used in the rest of the paper and illustrate it with our case study. We stress the fact that the technical terms and notions presented here (including harm trees) are useful to the reader but do not have to be known by the users. Users interact with the system only through the interface presented in Section IV, which hides all technicalities.

A. Definition of the System

A fitness tracking service consists of a fitness tracking device TD for each user i . This device collects fitness data fit_i and location data loc_i . This data is then forwarded to the service provider to be stored and processed. Apart from owning the device itself, the user generally needs to create a personal user account UA where he must provide identification ID_i and other information. The user must authenticate himself using his identification (ID_i) and password (pwd_i) to access his account. The fitness device owned by the user is linked to

¹For example, Recital 7 states that “Natural persons should have control of their own personal data”.

²For example, Recital 39 states that “Natural persons should be made aware of risks, rules, safeguards and rights in relation to the processing of personal data and how to exercise their rights in relation to such processing.”

UA. The system provides comprehensive information about the level of fitness of the user through a personalized dashboard accessible through the user account. The service provider uses fit_i and loc_i to derive this fitness related information $dfit_i$ (e.g. calories burnt, sleep pattern, active minutes and distance covered). Users can also maintain a list of contacts, share his data ($dfit_i$ and loc_i) with them and see the data shared by the contacts. For simplicity, we assume that the service provider manages the application server AS where all data processing takes place and the database server DS which stores all data.

B. Definition of data

We assume that some data such as fit_i and $dfit_i$ may or may not be associated with the identity ID_i and we use the notation $\overline{x_i}$ to denote the pair (x_i, ID_i) for conciseness. For example, the service provider may always store and process $\overline{fit_i}$ but give access to only fit_i to a third party.

The database server DS, which is a persistent storage, stores most types of data for user i in an encrypted format ($\overline{efit_i}$, $\overline{edfit_i}$, $\overline{eloc_i}$). It also stores the cryptographic keys k , k' and the passwords pwd_i . Password-protected data ($\overline{pdfit_i}$ and $\overline{ploc_i}$) are accessible through UA. Since the user knows the password, he can access $\overline{dfit_i}$ and $\overline{loc_i}$ through UA. He can access data that dates back to one year or one week (depending again on the retention time).

Data processing takes place in the application server AS. The tracking device TD usually stores data for a short period of time (for e.g., seven days). In the sequel, this type of storage is called *transient*. In both AS and TD, data is stored in encrypted format for a short duration.

The service provider ensures that all data are protected by encryption and access control mechanisms. He also ensures the protection of cryptographic keys and passwords.

C. Definition of the risk sources

Risk sources either intentionally or unintentionally, legally or illegally cause privacy breaches [9]. We consider the following risk sources for our case study³: the system owner or service provider (A.1), friends of the user (A.2), hackers (A.3), the general public (A.4) and business partners of the service provider (e.g., insurance providers) (A.5).

D. Definition of the privacy harms

Fitness service providers may sell identifiable data to third parties such as health insurance providers who may use the user's fitness data to increase health insurance premiums (H.1). User's personal habits or health conditions may also become accessible to the public (H.2) due to hackers or via other means. We refer to such negative impacts on the data subjects as privacy harms [9]. Other harms are also possible, but we do not discuss them here because of space limitations.

³Other risk sources such as governments could also be considered but they are not discussed here for the sake of conciseness.

E. Definition of the feared events

Harms result from the combination of one or more feared events [9] which are technical events of the system made possible by access to personal data (which we call "exploitation of data" here by analogy with the exploitation of vulnerabilities in computer security). Generally speaking, we distinguish three types of feared events resulting from, respectively, the access to personal data (FE.3), the use of personal data (FE.1), and the disclosure of personal data (FE.2).

F. Construction of the harm trees

A harm tree represents the relationships among privacy harms, feared events and the exploitation of personal data. The root node of a harm tree denotes a privacy harm. Leaf nodes represent the exploitation of data by the most likely risk source (for the root harm). They are represented as triples (personal data, system component, risk source). Intermediate nodes are feared events caused by risk sources. They can be seen as intermediate steps of privacy attacks. Child nodes are connected by an AND node if all of them are necessary to give rise to the parent node and by an OR node if any one of them is sufficient.

As an illustration, the harm trees pictured in Figure 1 and Figure 2 are assumed to result from a risk analysis (see Section V) conducted for our case study (for example in the context of an enhanced Data Protection Impact Assessment). Figure 1 shows that the harm increased health insurance premium (H.1) can be caused by the service provider disclosing to health insurance providers (FE.2) fitness related data (which may be done by disclosing either fitness data fit_i , or other data that can reveal fitness data such as $dfit_i$ in identified or de-identified form). Health insurance providers may use this data to increase health insurance premium (FE.1) for users who they deem unfit. To exploit de-identified data, the health insurance provider (A.5) must have access to identification (ID_i) information of the users as background information. Similarly, Figure 2 pictures the harm tree for H.2.

Some combinations of risk sources and exploitations are very unlikely in practice. For example, friends (A.2) of the user are very unlikely to attack servers to get access to the data. These combinations are thus left out of the harm trees.

ID_i may be obtained by a risk source either from a system component or as background information ("Bck" in harm trees). We assume that all other data elements can be obtained only from a system component (they are unlikely to be known as a background information by a risk source).

III. USER PRIVACY PREFERENCES

In this work, we assume that data subjects can specify their privacy preferences or privacy settings through *privacy parameters*. For the sake of conciseness, we consider only four *privacy parameters* for our case study:

- 1) The *retention duration* (Ret) of fitness ($dfit_i$, fit_i) and location data (loc_i) at the service provider's database (DS) and in the user account (UA). It can have two values: one year (L)

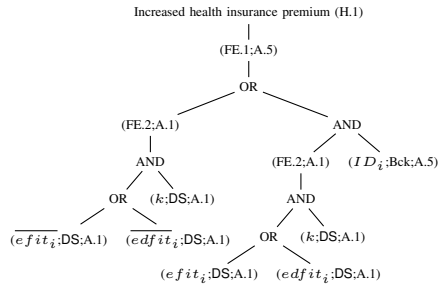


Fig. 1. Harm tree for “increased health insurance premium” (H.1)




Fig. 2. Harm tree for “undesirable access to personal habits by the public” (H.2)

and one week (S). The default value is: one week (S). The value of Ret for the components TD and AS are always short.

2) The *visibility* (Vis) of derived fitness ($dfit_i$) and location data (loc_i) from the user account (UA). These data can be made visible to the public (Pu) or friends (F) or kept private (Pr). The default value is: private (Pr).

3) The *recipients* (Rec) of fitness ($dfit_i$, fit_i) and location data (loc_i) from the service provider. The service provider may choose to disclose these data only to his sub-contractors (DA) essential to provide the service or to any third party (All) for different incentives. The default value is: sharing only with sub-contractors (DA).

4) The *form* ($Form$) in which the service provider discloses fitness ($dfit_i$, fit_i) and location data (loc_i) to their recipients. The service provider can disclose these data in an identified form (Id) or disclose only de-identified data ($deId$). The default value is: disclosure of de-identified data ($deId$).

A *user privacy preference* is a conjunction of the values assigned to the *privacy parameters*. For example, the conjunction $(Ret = L) \wedge (Rec = DA) \wedge (Vis = F) \wedge (Form = deId)$ is a *user privacy preference*. The likelihood of the privacy harms may be affected by these preferences. Here, for simplicity, we assume that the user sets the same value for each parameter for all data elements. The default values of the *privacy parameters* are chosen such that they constitute the most protective privacy preference which is given by $(Ret = L) \wedge (Rec = DA) \wedge (Vis = F) \wedge (Form = deId)$.

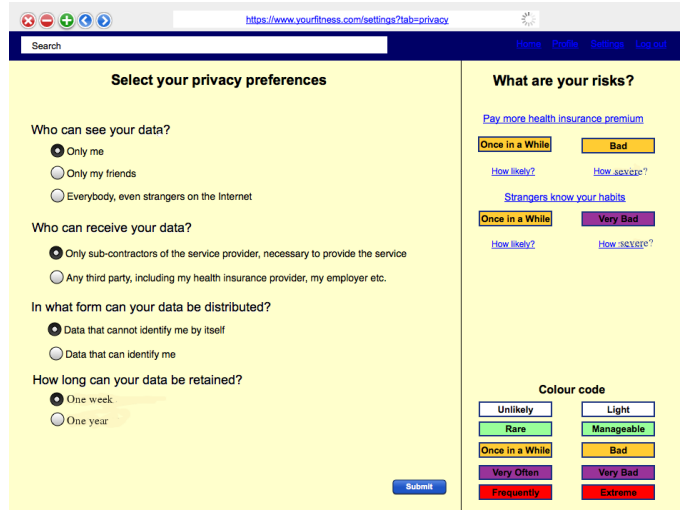


Fig. 3. First level screen showing the default *user privacy preference* ($Ret = S) \wedge (Rec = DA) \wedge (Vis = Pr) \wedge (Form = deId)$ and its risks

IV. USER INTERFACE DESIGN

The objective of this section is to show how users can be informed about the consequences of their privacy settings in a simple and intelligible way. Therefore, we focus on the user interface here and leave the presentation of the actual computation of the risks to the next section.

Users can express their *privacy preferences* through their account UA . The choices can be made when they initially open their account and set up their fitness tracking device.

Figure 3 shows the interactive screen using which users set their *privacy preferences*. *Privacy preferences* are displayed on the left pane, referred to as the *privacy preference pane*. For different *privacy preferences* selected by the user, the screen displays (prominently, on its right hand side) the risks that they may face. We refer to this pane as the *privacy risk pane*.

A. The Privacy Risk Pane

The right pane (see Figure 3) shows the users the risk levels corresponding to each harm. Each harm is presented using a short phrase that can be easily understood by the user. For example, the harm H.1 is presented as “Pay more health insurance premium”. Below each harm, the risk level is presented as the likelihood and the severity of the harm, using coloured buttons. The likelihood of the harm is dependent on the *user privacy preferences*. In contrast, the severity results only from the nature of the harm. To explain to the user what these buttons mean, we colour and label them using very short text, both indicative of their meaning and also caption them with phrases like “How likely?” (referring to the likelihood) and “How severe?” (referring to the severity).

When a user changes his privacy setting, the colours and texts inside the button representing likelihood also change (see Section IV-C).

B. The Privacy Preference Pane

On the left pane of the screen, users are asked a series of questions to determine their *privacy preferences*. The questions

are followed by alternatives which the users can select (by clicking on the corresponding radio buttons). The default selection is the most privacy preserving one. As Figure 3 shows, the most privacy preserving alternative is presented first. The questions and answer alternatives are as follows.

- 1) Who can see your data? (*Vis*).
 - Only me (*Pr*)
 - Only my friends (*F*)
 - Everybody, even strangers on the Internet (*Pu*)
- 2) Who can receive your data? (*Rec*)
 - Only sub-contractors of the service provider, necessary to provide service (*DA*)
 - Any third party (*All*) including my health insurance provider, my employer etc.
- 3) In what form can your data be distributed? (*Form*)
 - Data that cannot identify me on its own (*deId*)
 - Data that can identify me (*Id*)
- 4) How long can your data be retained? (*Ret*)
 - One week (*S*)
 - One year (*L*)

Whenever the question or the answer alternatives involve term(s) that the user may be unfamiliar with, suitable but short explanations and examples are used. For example, the user may not fully understand who a “third party” is. Therefore, examples are provided to make the user aware that a third party may mean his health insurance provider or even his employer.

C. Interaction between the panes

Initially, the privacy risk pane of the screen displays the risks to the users for the default alternatives. Whenever the user inputs a preference that is different from the default option, the risk pane displays the resulting changes in the harm likelihoods.

Thus, the interactive screen allows the user to observe the impact of the change he makes in the *privacy preferences* on the risk level. Based on these risk levels, he can decide on the most acceptable *privacy preference*. After he is satisfied with his selection, he can press the “*Submit*” button to communicate his preference to the service provider.

D. Links for more information

The primary or first level screen leads to several linked or second level webpages. There are four types of links, all from the privacy risk pane: 1) from the privacy harms (“*Pay more health insurance premium*” and “*Strangers know your habits*”); 2) from the likelihoods (“*How likely?*”); 3) from the severities (“*How severe?*”) and 4) from the coloured buttons denoting the likelihood and severity levels. In all the second level webpages, we still retain the privacy risk pane so that the user can see the risk levels as he learns more about the different terms and colour codes. A “*Back*” button on the top left of these webpages allow the users to go back to the first level screen. Below, we discuss the design of the second level screens.

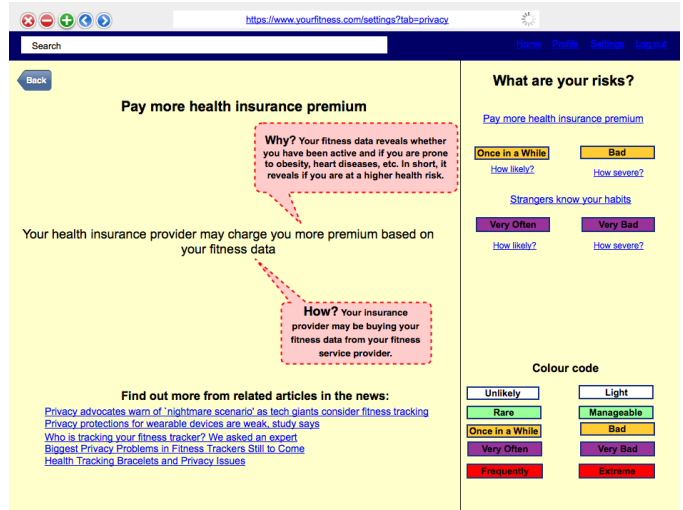


Fig. 4. Second level screen linked from the harm “*Pay more health insurance premium*” for the user privacy preference ($\text{Ret} = S$) \wedge ($\text{Rec} = DA$) \wedge ($\text{Vis} = F$) \wedge ($\text{Form} = deId$)

1) *Links from privacy harms*: The privacy harms on the privacy risk pane are linked to further screens that explain what these harms mean. Figure 4 shows the screen obtained by clicking the link on the text “*Pay more health insurance premium*” on the privacy risk pane. Each screen states the harm in a comprehensible sentence and also answers potential questions that the user may have after reading it. For example, in Figure 4, we see the explanation “*Your health insurance provider may charge you more premium based on your fitness data.*” along with two red bubbles answering two important questions the user may ask: 1) why a health insurance provider could charge more premium based on fitness data and 2) how the health insurance provider could obtain such data.

To substantiate our claims that the harms considered are indeed legitimate scenarios and to educate the user further, we also provide links to relevant news articles at the bottom of the screens. These news articles present scenarios where such harms have occurred or may occur.

2) *Links from likelihoods*: The privacy risk pane shows the likelihoods of the different privacy harms. The likelihood of each harm is labelled as “*How likely?*” which also links to a screen that shows what leads to the current level of likelihood. Figure 5 shows this screen corresponding to the likelihood of H.1. In this screen, we highlight to the user what contributes to the likelihood of the harm, based on our analysis using the harm trees. For example, Figure 5 shows that the likelihood of the harm H.1 is only “*Once in a While*” mainly because of two positive (indicated by green button with “+”) factors: 1) the service provider can only disclose de-identified data to his sub-contractors (since $\text{Form} = deId$ and $\text{Rec} = DA$) and 2) the sub-contractors are bound legally by the service provider not to re-identify the data disclosed to them.

3) *Links from severity*: The privacy risk pane shows the severities of the different privacy harms. The severity of each harm is labelled as “*How severe?*” which also links to a screen that shows what leads to the level of severity. Figure 6 shows

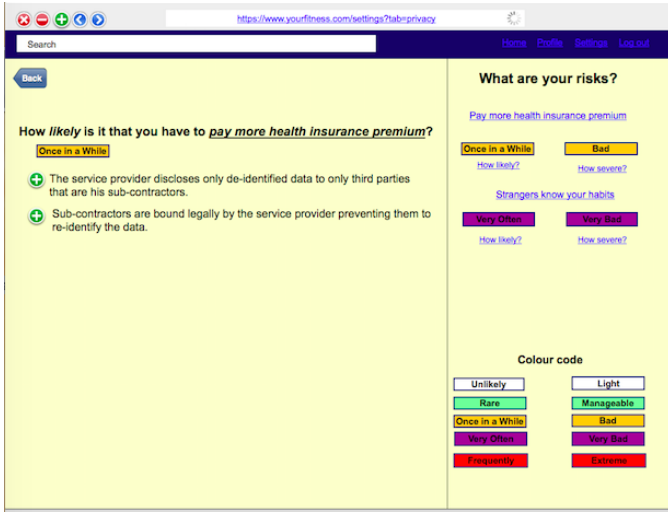


Fig. 5. Second level screen linked from “How likely?” for H.1 for the user privacy preference $(Ret = S) \wedge (Rec = DA) \wedge (Vis = F) \wedge (Form = deId)$

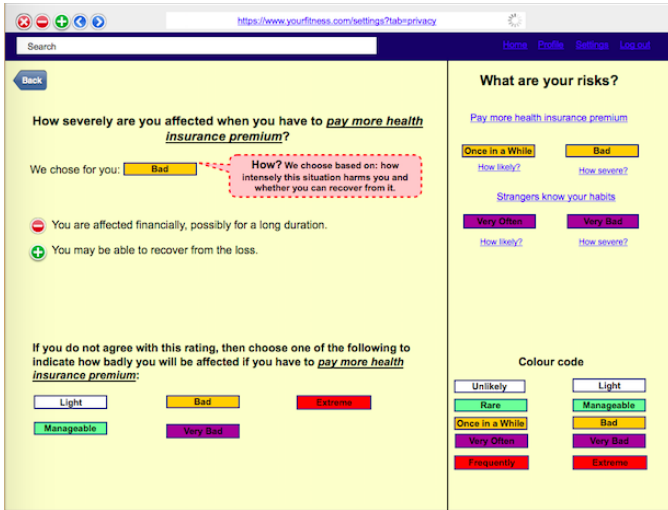


Fig. 6. Second level screen linked from “How severely?” for H.1 for the user privacy preference $(Ret = S) \wedge (Rec = DA) \wedge (Vis = F) \wedge (Form = deId)$ this screen corresponding to the severity of H.1.

We assign severity values to harms based on two factors used by the CNIL [7]: 1) how much inconveniences or difficulties are faced by the data subject and 2) whether or how easy it is for the data subject to recover from the harm. We remind the user of these two factors in a red bubble following the level of severity (see Figure 6) and explain to him our reasons for assigning a certain severity level to a harm.

If the user does not agree with the severity level assigned to a harm, he can select the severity level (from the same scale). This severity level is then displayed in the privacy risk pane.

4) *Links from coloured buttons:* The coloured buttons on the privacy risk pane of the first level screen also link to another screen which explains in detail the colour scheme and the texts inside the buttons.

E. Usability features

Great care has been taken to enhance the usability of the user interface. In particular, we have: 1) Used illustrative examples and avoided the use of technical terms. 2) Presented the privacy harms in a simple language so that users can relate them to harmful consequences in their own lives. 3) Ensured readability by restricting the amount of information presented in each screen and using appropriate fonts and colours. 4) Ensured that the process of selecting privacy preferences is not confusing or too time consuming to users by including only one basic window for the purpose and limiting the number of questions and answer options. 5) Allowed users to choose the severity level of the privacy harms if they do not agree with the default choice, but kept this as an option only at a later stage to avoid confusion and too much time consumption for the average user who may just agree with the default values. 6) Presented more information through hyperlinks about the privacy harms and their severity and likelihood and the colour codes for inquisitive users who may want to educate themselves further. We have included news articles related to each privacy harm so that users can refer to them to improve their awareness. 7) Kept the default privacy preference to be the most privacy preserving one, so that users who skip this selection step can still benefit from the highest level of privacy (“privacy by default” principle).

V. RISK ANALYSIS WITH PRIVACY PARAMETERS

In Section IV, we presented the interactions with users to allow them express their privacy preferences and to inform them about privacy risks, but without explaining the actual computation of these risks. Here, we focus on the risk analysis itself, based on the methodology introduced in [9], [10], enhanced with facilities to deal with privacy parameters. The primary objective of a risk analysis is to identify the privacy harms for a system in a given context and to assess the associated risks, generally measured in terms of likelihood and severity.

Several factors can influence the likelihoods of the privacy harms. The *exploitability* of personal data can be characterized by the resources (e.g., technical resources, access rights, background knowledge) needed by a risk source to exploit them. The dual notion is the *capacity* of a risk source which is defined by its resources (e.g., technical resources, access rights, background knowledge). The *motivation* represents the incentives and disincentives of a risk source to cause a feared event or a harm. The values of the *privacy parameters* influence the values of *exploitability* of data, the *capacity* and the *motivation* of risk sources in certain cases. In the next subsections, we study this influence and show how it can be taken into account in the privacy risk analysis process.

A. Exploitability of Data

Some data may be accessible to certain risk sources legitimately. It may be either because the risk source controls a component storing or processing the data or because the risk source has legitimate access rights to the data (e.g., when

Component	UPref	Data	Exploitability
DS	Ret = L	$edfit_i, efit_i,$ $eloc_i$	Transience
DS	Ret = S	$edfit_i, efit_i,$ $eloc_{i,t}$	Persistence
UA	Ret = L	$pdfit_i, ploc_i$	Transience
UA	Ret = S	$pdfit_i, ploc_{i,t}$	Persistence
AS/TD	× (always stored for short time)	$edfit_i, efit_i,$ $eloc_i$	Persistence

TABLE I

EXPLOITABILITY VALUES OF DATA AFFECTED BY RETENTION (RET)

$Vis = F$). The control over data allows a risk source to use that data in any way. In our case study, the service provider (A.1) has full control over the database server (DS) and the application server (AS) and hence can access all necessary data in these components without attacking them or UA.

Risk sources that have no control over a piece of data have to exploit (or attack) it, persistently or transiently. To this aim, they need resources that may or may not be available to them. By transient exploitation, we mean an exploitation for a short period of time or infrequent exploitation; by persistent exploitation we mean an exploitation for a long period of time (e.g., for several days or months). When the retention duration of a data is long, i.e., a year (Ret = L), it becomes vulnerable to transient exploitation. In contrast, if the retention duration is short, i.e., a week (Ret = S), it is only vulnerable to persistent exploitation. For example, data (such as $edfit_i$ and $pdfit_i$), if retained for a short duration in DS or in UA, require persistent exploitation. On the other hand, when retained for a long duration, transient exploitation is sufficient. We assume that data is stored in AS and TD for a short time. Hence, these data require persistent exploitation.

Cryptographic keys and passwords are securely stored by the service provider. So, control on the component storing them (DS) is required to exploit them (this is the highest level of protection or the lowest level of exploitability). We also assume that the service provider has taken enough measures to prevent the disclosure of passwords through UA. So, control is required to exploit password-protected data by obtaining pwd_i from UA.

The exploitability values of the different data types for different components and for different retention durations are shown in Table I⁴. The exploitability values that are not affected by retention durations are shown in Table II.

Background information is not a part of the data stored in the system. So, it does not have any exploitability value.

B. Capacity and Motivation of Risk Sources

A risk source can possess the capacity for transient or persistent exploitation or may control one or more data elements or one or more components. The highest capacity of any risk source with respect to a data element or a component is to have control over that data element or component. For example, the service provider (A.1) controls AS and DS. The least capacity

⁴In Table I, the exploitability value of data stored in TD or in AS are affected by the retention time. However, in our case, we have assumed that the service provider has reasonably decided to store the data in these components for a short duration. If he had made the other choice or different choices for these two components, the exploitability values would have been different.

Component	UPref	Data	Exploitability
UA	× (irrespective of retention time)	$dfit_i, fit_i, loc_i$	Control
DS	× (irrespective of retention time)	k, k', pwd_i	Control
UA	× (no retention time)	pwd_i	Control

TABLE II

EXPLOITABILITY VALUES OF DATA NOT AFFECTED BY RETENTION (RET)

of any risk source is the inability to perform any exploitation as in the case of the user's friend (A.2) when $Vis = Pr$. A.3 and A.5 have persistent and transient capacities respectively.

The control over data can be influenced by the value of *visibility*. For example, under the default value for *visibility* (i.e., $Vis = Pr$), a friend (A.2) of the user or the public (A.4) have no control over any data and no technical resources to exploit them. However, when the user allows his data to be visible to his friends (i.e., $Vis = F$), they gain control over this data (similarly for the public in general when $Vis = Pu$).

The availability of background information to some risk sources is also considered as a part of their capacity. We consider that a risk source has a "high" capacity if it possesses the background information relevant for an exploitation, and "low" otherwise. The only background information considered here is ID_i . As shown in Figure 1, A.5 must be able to exploit $dfit_i, fit_i$ and/or loc_i which can be provided by A.1 in de-identified form. We assume that A.5 has a "low" chance of possessing this background information.

The *privacy parameters recipients* and *form* influence the value of the motivation of the service provider (A.1) to perform an exploitation. We assume that the motivation of the service provider to comply with the privacy settings of the user is always "high". If the user specifies a less privacy preserving option, then the motivation of the service provider to choose the option which gives him more incentive is always "high" and that of the option which gives him less incentive is "low". For example, the motivation of the service provider to disclose identified data is "low" when the user limits this disclosure to de-identified data, due to the fear of legal sanctions and the loss of consumer trust. Otherwise, the motivation of disclosing identified data is "high" due to financial incentives.

For other risk sources, the motivations (see Table III) are not affected by the privacy parameters. The motivation of a friend (A.2) to access (FE.3) and disclose the user's personal data (FE.2) when it is legitimately accessible to him (i.e., $Vis = F$) is "medium" because such an access is generally easy but these actions may lead to embarrassments. On the other hand, the motivation for a friend (A.2) to access (FE.3) and disclose personal data (FE.2) through an attack to know the password (A.2 is not a hacker) is "low" because it may lead to loss of friendships. Hackers, on the contrary, have a "high" motivation to access (FE.3) and disclose (FE.2) any data. Any other member of the public (A.4) has a "medium" motivation to access the user data (FE.3) (they may seek quick monetary gains, but also fear getting caught). Third parties, other than sub-contractors of the service provider, have a "high" motivation (considering worst case) to disclose and

Data	Risk Source	Feared Event	Motivation
$dfit_i, fit_i, loc_i$	A.2	FE.2 and FE.3	Medium
$pwd_i, pdfit_i, ploc_i$	A.2	FE.2 and FE.3	Low
$dfit_i, fit_i, loc_i, pwd_i, pdfit_i, ploc_i$	A.2	FE.1	×
$dfit_i, fit_i, loc_i$	A.4	FE.3	Medium
$dfit_i, fit_i, loc_i$	A.4	FE.1, FE.2	×
$dfit_i, fit_i, loc_i, dfit_i, fit_i, loc_i, ID$	A.5 (except DA)	FE.2 and/or FE.1	High
$dfit_i, fit_i, loc_i, dfit_i, fit_i, loc_i, ID$	A.5 (DA)	FE.2 and/or FE.1	Low
$dfit_i, fit_i, loc_i, dfit_i, fit_i, loc_i, ID$	A.5	FE.3	×
$pwd_i, pdfit_i, ploc_i, k, k', edfit_i, efit_i, eloc_i$	A.3	FE.1 and/or FE.2 and/or FE.3	High

TABLE III

MOTIVATION OF RISK SOURCES (EXCEPT THE SERVICE PROVIDER (A.1))

use data for unauthorized purposes (FE.2, FE.1) and to make use of any identifying information (*ID*) available to them as background information. Sub-contractors of the service provider, however, have “low” motivation as they are legally bound by the service provider not to disclose, misuse or re-identify data disclosed to them by the service provider.

Some combinations of feared events and risk sources do not make sense. The corresponding rows are marked with ‘×’ in Table III. For example, the friends (A.2) of the user are not given access to data for any particular purpose.

Similarly, the public (A.4) is not provided access to data with any specific purpose nor is there any intention of the user to hide some data when he allows its disclosure to the public. Third parties (A.5) do not perform unintended access to data (FE.3) because this would qualify them as hackers (A.3).

The motivation of business partners of the service provider (A.5) to use background information is “high” due to potential financial incentives.

C. Computation of Likelihoods

The computation of the likelihoods of the harms based on the harm trees shown in Section II-F can be carried out in two steps. The first step is the assessment of the likelihoods of the leaves of the harm trees (likelihood of exploitation of personal data) from the *motivation* and the *capability* of the relevant *risk sources* using Table IV. The capability of the risk source to perform an exploitation is derived by comparing the value of the exploitability of the data and the capacity of the risk source. A risk source has a “high” capability when its capacity satisfies the desired conditions (w.r.t. control, persistent and transient access) for exploitability, otherwise it has a “low” capability. This assessment is based on Section V-A and Section V-B. To be consistent with other leaf nodes, the leaf nodes corresponding to background information (for which there are no exploitability) are directly assigned a likelihood value based on a “high” capability (since background information, when available, is easily usable by risk sources) and the motivation of the risk source to use it. The second step is the computation of the likelihood of each harm according to the following rules (applied bottom-up), where P_i is the likelihood of the i th child node: R1) AND node with independent child nodes: $\prod_i P_i$.

Likelihood of exploitation	Risk source capability	Motivation
Negligible	Low	Low
Limited	High	
Negligible	Low	Medium
Significant	High	
Limited	Low	High
Maximum	High	

TABLE IV

MEASUREMENT RULE FOR LIKELIHOOD OF EXPLOITATION

Scale used for computation	How likely?
Negligible	Unlikely (white)
Limited	Rare (green)
Intermediate	Once in a While (amber)
Significant	Very Often (purple)
Maximum	Frequently (red)

TABLE V

MAPPING OF SCALES FOR LIKELIHOOD

R2) AND node with dependent child nodes⁵: $Min(P_i)$, R3) OR node with independent child nodes: $1 - \prod_i (1 - P_i)$. R4) OR node with dependent child nodes⁶: $Min(1, \sum_i P_i)$.

To perform the computations of the second step, it is necessary to translate the symbolic likelihood values of Table IV into numerical values. This transformation has to be made by the privacy expert in collaboration with the owner and should be documented. In this paper, we use as an illustration the following correspondence for the likelihood values (p): 1) *Negligible (N)*: $p < 0.01\%$; 2) *Limited (L)*: $0.01\% \leq p < 0.1\%$; 3) *Intermediate (I)*: $0.1\% \leq p < 1\%$; 4) *Significant (S)*: $1\% \leq p < 10\%$; 5) *Maximum (M)*: $p \geq 10\%$.

D. Choice of Privacy Preferences

Table VI shows that the default user setting leads to the lowest level of risk, i.e., the likelihoods of both H.1 and H.2 are “intermediate”. We also observe that changing the default values of *Ret* to *L* does not increase harms (i.e., their likelihood values remain unchanged). Changing the default values of *Rec* to *All* or *Vis* to *Pu* or *Form* to *Id* make the harms riskier (i.e., their likelihood values increase).

The results of the previous sections can help the user to decide upon an acceptable likelihood for each harm, given their severity. Based on Table VI, and the acceptable threshold, he can then decide which values for the privacy parameters he prefers. Let us assume that the user decides that the acceptability threshold for a harm with “very bad” severity (for example H.2) is “intermediate” and that of a harm with “bad” severity (for example, H.1) is “significant”. Then, he may choose any privacy preference (from Table VI) other than the ones in which $Vis = F$ or $Vis = Pu$.

Table V shows how the colour coding and the texts used in the screens for “How likely?” in Section IV map to the scale for likelihood used in Section V.

VI. RELATED WORKS

The communication of privacy policies to users in a comprehensible form has been an important focus of privacy

⁵In order to err on the safe side in terms of privacy protection, we consider dependent nodes such that one node may imply the other nodes.

⁶In order to err on the safe side in terms of privacy protection, we consider dependent nodes such that each node may exclude the other nodes.

User Preference	Likelihood for H.1	Likelihood for H.2
$(Ret = S) \wedge (Rec = DA) \wedge (Vis = Pr) \wedge (Form = deId)$ (Default preference)	Intermediate	Intermediate
$(Ret = L) \wedge (Rec = DA) \wedge (Vis = Pr) \wedge (Form = deId)$	Intermediate	Intermediate
$(Ret = S) \wedge (Rec = All) \wedge (Vis = Pr) \wedge (Form = deId)$	Significant	Intermediate
$(Ret = S) \wedge (Rec = DA) \wedge (Vis = F) \wedge (Form = deId)$	Intermediate	Significant
$(Ret = S) \wedge (Rec = DA) \wedge (Vis = Pu) \wedge (Form = deId)$	Intermediate	Significant
$(Ret = S) \wedge (Rec = DA) \wedge (Vis = Pr) \wedge (Form = Id)$	Significant	Intermediate
$(Ret = S) \wedge (Rec = All) \wedge (Vis = Pr) \wedge (Form = Id)$	Significant	Intermediate

TABLE VI

SUMMARY OF HARM LIKELIHOODS FOR DIFFERENT USER PREFERENCES

research. The ToS;DR project [23] aims to classify, through a transparent and peer-reviewed process, the terms of service and the privacy policies into six colour-coded classes. Privacy icons [17] or privacy policy icons [13] are simplified pictures used to visualize elements of privacy policies. Inspired by nutrition labeling etc., Kelley et al. [19], [20] propose the privacy nutrition label to improve the accessibility, readability and understanding of privacy policies among users.

Poor, confusing interface design, permissive default settings, limited visual feedback etc. can often lead to the underutilization of available privacy options [14], [21].

Different types of Transparency Enhancing Tools (TETs) have also been proposed to allow users to express their privacy choices or to inform them about the privacy policies of service providers. PrivacyBird [1] can compare user privacy preferences with P3P policies and help the user decide whether to reveal data to the website [8]. Other similar TETs provide insight about the privacy implications of potential or past data disclosures (Mozilla Privacy Icons [2], PrimeLife's Privacy Dashboard [3], Google Dashboard [4], Privacyscore [18]). Still others provide insights into third party tracking (Lightbeam [5]) or promote privacy awareness and education about privacy problems among users and (Me & My Shadow [6]). The work described in this paper is complementary to the above proposals. Unlike previous work in this area, our objective is to help users in the definition of their privacy settings, based on a privacy risk analysis. It is also complementary to previous papers by the authors [9], [10] which introduced a methodology for privacy risk analysis and its application to smart metering but did not address the use of privacy risk analysis to enable informed privacy settings.

VII. CONCLUSIONS AND FUTURE DIRECTIONS

The concepts presented in this paper can form the basis of a full fledged privacy tool enabling data subjects to make informed choices about their privacy settings. It could also form the core of an education tool to increase awareness about privacy [24]. An avenue for further research is the extension to dynamic risk analysis to take into account the personal data disclosure history of the user. It would be useful, for example, to analyze the impact on privacy risks of the disclosure of new personal data to a third party that has already collected data on the subject. Another interesting research direction is the integration of alternative actions (such as disclosure of anonymized data, less precise data, or even fake data) and their

consequences. A better understanding of these options would further enhance the control of individuals on their personal data.

VIII. ACKNOWLEDGEMENTS

This work has been partially funded by the CHIST-ERA project UPRISE-IoT.

REFERENCES

- [1] <http://www.privacybird.org>.
- [2] <https://disconnect.me/icons>.
- [3] <http://primelife.ercim.eu/results/opensource/76-dashboard>.
- [4] <https://support.google.com/accounts/answer/162744?hl=en>.
- [5] <https://www.mozilla.org/en-US/lightbeam/>.
- [6] <https://myshadow.org>.
- [7] Commission Nationale de l'Informatique et des Libertés (CNIL). Privacy Impact Assessment (PIA) Tools (templates and knowledge bases), 2015.
- [8] L. F. Cranor, P. Guduru, and M. Arjula. User Interfaces for Privacy Agents. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 13(2):135–178, 2006.
- [9] S. J. De and D. Le Métayer. PRIAM: A Privacy Risk Analysis Methodology. In *11th International Workshop on Data Privacy Management (DPM)*. IEEE, 2016.
- [10] S. J. De and D. Le Métayer. Privacy Harm Analysis: A Case Study on Smart Grids. In *International Workshop on Privacy Engineering (IWPE)*. IEEE, 2016.
- [11] European Commission. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016.
- [12] E. Field. Biggest Privacy Problems in Fitness Trackers Still To Come. *Law360*, 2015.
- [13] S. Fischer-Hübner, J. Angulo, and T. Pulls. 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*, pages 77–92. Springer, 2013.
- [14] R. Gross and A. Acquisti. Information Revelation and Privacy in Online Social Networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 71–80. ACM, 2005.
- [15] H. Hedbom. A Survey on Transparency Tools for Enhancing Privacy. In *IFIP Summer School on the Future of Identity in the Information Society*, pages 67–82. Springer, 2008.
- [16] K. Hill. Fitbit Moves Quickly After Users' Sex Stats Exposed. *Forbes*, 2011.
- [17] L.-E. Holtz, K. Nocun, and M. Hansen. Towards Displaying Privacy Information with Icons. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 338–348. Springer, 2010.
- [18] M. Janic, J. P. Wjibenga, and T. Veugen. Transparency Enhancing Tools (TETs): An Overview. In *2013 Third Workshop on Socio-Technical Aspects in Security and Trust*, pages 18–25. IEEE, 2013.
- [19] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder. A Nutrition Label for Privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, page 4. ACM, 2009.
- [20] P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pages 1573–1582. ACM, 2010.
- [21] H. R. Lipford, A. Besmer, and J. Watson. Understanding privacy settings in facebook with an audience view. *UPSEC*, 8:1–8, 2008.
- [22] I. Oskolkov. Your fitness is their business. Nothing personal. <https://blog.kaspersky.com/fitness-trackers-privacy/6480/>, 2014.
- [23] H. Roy, d. M. Jong, J.-C. Borchardt, I. McGowan, J. Stout, and S. Azmayesh. Terms of Service Didn't Read. <https://tosdr.org>, 2012.
- [24] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A Design Space for Effective Privacy Notices. In *11th Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17. USENIX Association, 2015.