Towards an Effective Privacy Impact and Risk Assessment Methodology: Risk Assessment

Majed Alshammari and Andrew Simpson

Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford OX1 3QD, UK
firstname.secondname@cs.ox.ac.uk

Abstract. It is increasingly recognised that Privacy Impact Assessments (PIAs) play a crucial role in providing privacy protection for data subjects and in supporting risk management for organisations. However, existing PIA processes are typically not accompanied with proper guidelines and/or methodologies that sufficiently support privacy risk assessments and illustrate precisely how the core part of the PIA—a risk assessment—can be conducted. We present an approach for assessing potential privacy risks built upon a privacy risk model that considers legal, organisational, societal and technical aspects. This approach has the potential to underpin a systematic and traceable privacy risk-assessment methodology that can complement PIA processes.

1 Introduction

In order to anticipate and prevent the processing operations that may lead to privacy violations or harms, the adverse impacts of these operations need to be proactively assessed in the early stages of the design process [6]. This has led to the emergence of the concept of a *Privacy Impact Assessment (PIA)*. PIAs tend to focus more on legal and organisational aspects than on social and technical ones [6, 7]. As such, it is necessary for PIA processes to be complemented by a privacy risk-assessment methodology that adopts an appropriate privacy risk model that considers organisational, legal, societal and technical aspects [2].

We present a methodical approach for assessing potential privacy risks built upon a privacy risk model that considers legal, organisational, societal and technical aspects. It illustrates the main steps of analysing and assessing the risk of privacy harms that may result from the processing of personal data. The approach has the potential to help underpin a systematic and traceable privacy risk-assessment methodology that can complement PIA processes in a holistic and effective fashion. We argue that this contribution lays the foundation for developing rigorous and systematic PIA methodologies.

2 Background and Motivation

Several countries have developed different PIA processes and/or methodologies [13]. The UK PIA handbook [12], for example, was based on extensive

reviews of existing PIA methodologies [3, 13]. In addition, the Privacy Impact Assessment Framework (PIAF) [14] is a project funded by the European Commission to develop a step-by-step guide to the PIA process. As another example, in [13] a 16-step guide to the PIA process is described.

The core of a PIA is a risk assessment, which involves risk identification and risk mitigation [9]. Although PIAs are expected to follow the same philosophy, existing PIA processes largely fall short in this respect [6, 9]. Existing processes cannot be applied easily: they are imprecise, lengthy or improperly structured [9]. They are typically not accompanied with proper guidelines, methodologies or risk models that sufficiently support privacy risk assessments or illustrate precisely how the technical part of the PIA can be conducted [6, 9]. For example, the steps of the process described by the UK PIA handbook [12] are generic and individual risks are not well-matched with corresponding controls. Importantly, it cannot be used as a process reference model [9].

Privacy risk assessments need to go beyond traditional security assessments to consider the nature of the risks arising from the processing of personal data, not least because the nature of privacy harms differs from the adverse impacts of security events; such impacts may extend beyond individuals to relatives, friends or wider society [10]. This necessitates adopting a risk model that defines the key risk factors that have impacts on the privacy of data subjects and establishes a conceptual relationship among these factors [6].

Typically, risk-assessment methodologies rely upon well-defined attributes of the key risk factors to determine levels of risk [8]. The specification of these factors, along with their attributes, requires an appropriate model that can be used to identify and analyse risks to the privacy of data subjects that may arise from the processing of personal data. To determine levels of risk, a risk assessment requires an assessment approach that establishes a set of assessment rules that specify the range of values the risk factors can assume [8]. Multiple threat scenarios need to be identified before assessing the severity and likelihood of privacy risks, which, in turn, require an analysis approach that describes how combinations of risk factors are identified and analysed to ensure adequate coverage of the problem space at a consistent level of detail [8]. To define a reasonable subset of all possible threat scenarios, it is useful to consider the nature of the relationships among these factors, the level at which these factors are characterised, and the dependencies between the attributes of these factors.

This gives rise to two questions: *how might one represent the relationships among the key risk factors in a way that is analytically useful for assessment?*; and *how might one use the attributes of the key risk factors identified in the risk model as inputs to determine the levels of risk in risk assessments?*

3 Building on Existing Approaches

3.1 Existing Approaches

With a focus on context-independent privacy-risk assessment approaches that may complement PIA methodologies, a number of privacy risk-management pro-

cesses, frameworks and methodologies have been proposed. We review, and subsequently build upon, two of these methodologies: [6] and [4]. We have chosen these as they explicitly define risk models that distinguish the key terms, assessable risk factors and relationships among these factors. With a focus on assessing the levels of privacy risks, we analyse these methodologies in terms of the risk model and the type of assessment approach.

The Methodology for Privacy Risk Management (CNIL).

The CNIL methodology [4] defines a risk model that illustrates feared events, threats, vulnerabilities and risk sources. A privacy risk is composed of one feared event and all the threats that make it possible. For a feared event to occur, one or more risk sources exploit, accidentally or deliberately, one or more vulnerabilities of supporting assets through different threats.[1]

A feared event describes both the adverse event and its potential impacts on subjects. It does not define the feared event by a set of attributes to help support the assessment approach; rather, it provides a set of the main types of feared events that affect the processing operations according to the types of primary assets. These focus on the availability, integrity and confidentiality of the primary assets. In addition, it takes into account the risks arising from the processing of personal data in a broader context by considering the potential impacts on identity, human rights, privacy and civil liberties. It does not define a privacy harm by a set of attributes to help support the assessment of its severity.

The risk level is assessed in terms of severity and likelihood, with levels of risk being based on two key risk factors: feared events are used to assess the severity, which depends on the level of identification of personal data and the prejudicial effect of the potential impacts; threats are used to assess the likelihood, which depends on the level of vulnerabilities of the supporting assets and the level of capabilities of the risk sources.

The prejudicial effect of the potential impacts is assessed on the level of consequences, the irreversibility of these consequences, and the level of difficulty with which these consequences can be overcome. The level of vulnerabilities is assessed on the ease of exploitation of the supporting assets. In particular, it focuses on threat actions that exploit the vulnerabilities of supporting assets rather than the primary assets.

The capabilities of risk sources are assessed based on their skills, time available, motivation, financial resources, etc. However, it does not explicitly consider the value of personal data to risk sources and the background knowledge when assessing their motivation.

The CNIL methodology uses a semi-quantitative approach that uses a fixed scale of levels (negligible, limited, significant, maximum), along with corresponding numbers. The risk levels are located on a risk map with severity and likelihood on its axes, with the aim of ordering and prioritising these risks.

[1] The PIA for smart grid and smart metering systems [5] is an example of a PIA that adopts the CNIL methodology to identify, analyse and assess potential privacy risks.

The Privacy Risk Analysis Methodology (PRIAM).

PRIAM [6] concretely defines a risk model that defines key risk factors with well-defined attributes: privacy harms, feared events, privacy weaknesses and risk sources. It also illustrates the relationships among these factors and describes the dependencies between their attributes. A privacy harm results from one or more feared events. Each feared event results from the exploitation of one or more privacy weaknesses by one or more risk sources.

The risk level is assessed in terms of severity and likelihood for each privacy harm. PRIAM estimates the severity of a privacy harm based on its intensity and victims, which are influenced by the 'irreversibility' and 'scale' attributes of the associated feared event respectively. The likelihood of a privacy harm is computed from the likelihood of its corresponding feared events derived from the likelihood of successful exploitation of associated privacy weaknesses, which depends on the capabilities of risk sources and the exploitability of privacy weaknesses. 'Harm trees' describe the many-to-many relationships among the key risk factors, representing possible exploitations as pairs of privacy weaknesses and risk sources.

The victims attribute is assessed according to the category of affected data subjects, whether they are 'individuals and their families', 'specific groups of individuals', or 'society as a whole'. Based on these categories, it is difficult to distinguish between the range of a privacy harm that affects data subjects only and the range of a privacy harm that affects data subjects along with their families. As such, it is useful to distinguish data subjects from their families when assessing the range of a privacy harm. In so doing, an additional level of assessment needs to be defined for comparison. This requires the establishment of an assessment rule that assesses whether the privacy harm affects data subjects only or data subjects together with their relatives, friends or colleagues.

The intensity of a privacy harm is assessed on its consequences, the irreversibility of these consequences, and the difficulty with which these consequences might be overcome. The capabilities of risk sources are assessed on their motivation (based on the value of the privacy breach and the incentives of the risk source) and capacity (based on the resources of the risk source and the exploitability of the relevant privacy weaknesses).

PRIAM uses a qualitative assessment approach involving various scales for assessing the severity of privacy harms; it also uses a semi-quantitative assessment approach that adopts a set of rules for assessing their likelihood.

3.2 The Problem Statement

The risk factors that have impacts on privacy risks and their contribution to the assessment of the overall risks vary between these approaches. This emphasises that these factors need to be defined in the context of data protection, and their contribution to the assessment of the overall risks needs to be defined at an appropriate level of detail. Further, the conceptual relationships between these factors need to be characterised by illustrating the dependencies between

the nominal and assessable attributes of each risk factor, and the dependencies between the nominal and assessable attributes of all risk factors. In addition, the assessment rules that specify the range of values the key risk factors can assume need to reflect the assessable attributes of these factors to facilitate their roles in risk assessments and their translation into qualitative terms for multiple stakeholders. Moreover, risk factors need to be represented in a way that is useful for analysis and assessment—with a view to developing a well-defined step-by-step guide can be developed to identify and assess potential privacy risks in a systematic and traceable manner.

3.3 The Contribution

To appropriately identify, assess and analyse the risk of privacy harms, a systematic and traceable privacy risk-assessment methodology—consisting of a well-defined risk model, an assessment approach, an analysis approach and an underlying process—is required. In [2], we defined a privacy risk model that supports the definition of the key risk factors (along with their attributes and conceptual relationships) by refining the risk model of [6]. In addition, we presented an analysis approach that describes how combinations of risk factors can be identified and analysed to ensure adequate coverage of the problem space at a sufficient level of detail. The results of the analysis approach can be used to develop and model threat scenarios that describe how the threat events that may result from the successful exploitation of primary assets' vulnerabilities by a set of threat sources can contribute to or cause privacy harms. In the following, we build upon those foundations by defining an assessment approach that refines the risk assessment approach of [6] in a number of ways.

First, it refines the harm tree approach by adding an additional level that separates threat sources from privacy vulnerabilities to represent the conceptual relationships among the key risk factors in a way that is analytically useful for analysis and assessment (addressing the first question of Section 2). Second, it adopts the risk model of [2] that: characterises the risk factors by well-defined attributes (nominal and assessable) to facilitate the identification, analysis and assessment of these factors in a systematic and traceable manner; and describes the dependences between the nominal and assessable attributes of the key risk factors, to refine how each risk factor can be used as an input to estimate the levels of risk (addressing the second question of Section 2). Third, it adopts the fixed levels of scale, along with the corresponding values of [4], with refined assessment rules that reflect the assessable attributes of the key risk factor (addressing the second question of Section 2).

4 A Risk-Assessment Approach

Our approach consists of four steps. The first step is to represent the conceptual relationships among the key risk factors for each privacy harm from which a reasonable set of all possible threat scenarios can be generated. The second

Table 1. Assessing the overall values from combinations of values, inspired by [4].

Sum of values	Overall values
< 5	1. Negligible
$= 5$	2. Limited
$= 6$	3. Significant
> 6	4. Maximum

step is to assess the severity of privacy harms. The third step is to assess the likelihood of occurrence. The fourth step is to assess the risk levels of privacy harms in terms of their severity and likelihood. We consider each in turn.

Our approach is built upon the risk model and the analysis approach of [2]. We consider the results of the analysis approach as relevant information that is necessary for determining the values of the attributes of key risk factors. In addition, we adopt the fixed scale of levels and the corresponding values of [4] (1. Negligible; 2. Limited; 3. Significant; and 4. Maximum) as assessment scales with refined and/or newly defined rules for assessing the key risk factors of the risk model of [2]. These scales can be easily translated for multiple stakeholders and allow relative comparisons between values in different scales or even within the same scale. Table 1 illustrates a set of rules for assessing overall values from combinations of values that can be applied to the key risk factors.

4.1 Step 1: The Construction of Harm Trees

In the risk model of [2], a privacy harm results from one or more threat events, each of which results from the successful exploitation of one or more vulnerabilities by one or more threat sources. Thus, it is useful to generate multiple threat scenarios describing how the threat events caused by the most likely threat sources can contribute to or cause a privacy harm.

In [6], a harm tree describes the relationship between a privacy harm (a root node) and all possible feared events (intermediate nodes) that exploit privacy weaknesses (leaf nodes) by the most likely risk sources, which are both represented as pairs of the form (privacy weakness, risk source). We slightly refine the harm tree by adding an additional level to present a privacy harm (a root node) and all possible threat events (intermediate nodes) that exploit the vulnerabilities of primary assets (intermediate nodes) by the most likely threat sources (leaf nodes), as illustrated in Figure 1. We use 'AND' and 'OR' connectors to combine child nodes and indicate whether all or some child nodes are necessary to enact the parent node. This refinement is to represent all possible exploitations of a vulnerability for each threat event by one or more threat sources. This helps analyse the exploitation of a vulnerability when there is collusion between two or more threat sources (when those threat sources are connected to a privacy vulnerability via 'AND'). In addition, it helps provide focus on the most important vulnerabilities that need to be addressed when a vulnerability is connected to several threat sources or its exploitation may lead to several threat events.

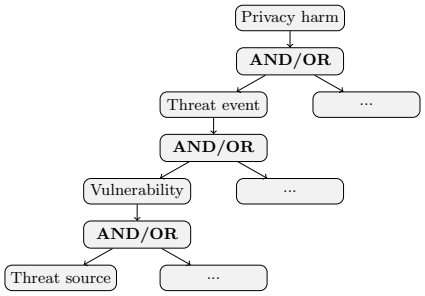

Fig. 1. The structure of the refined harm tree.

Table 2. Intensity of a privacy harm.

Values	**The affected data subjects encounter . . .**
1. Negligible	insignificant adverse consequences, which can be reversed without difficulty and last for a short time.
2. Limited	slight adverse consequences, which can be reversed with some difficulty and do not last for a long time.
3. Significant	serious adverse consequences, which can be reversed with great difficulty and last for a certain length of time.
4. Maximum	severe adverse consequences, which cannot be reversed at all and last for a long time.

4.2 Step 2: The Assessment of the Severity of a Privacy Harm

The *severity* of a privacy harm essentially depends on the intensity of adverse consequences of a threat event and the range of these consequences suffered by a variety of data subjects.

The *intensity* of a privacy harm represents the level of adverse consequences on the affected data subjects. It is based on the 'adverse consequences' attribute of the privacy harm, which indicates the duration of adverse consequences and their extent of damage. It is influenced by the 'nature', 'category' and 'scope' attributes of the corresponding threat event that reflect the irreversibility of the consequences. The factors that influence irreversibility vary between threat events according to their classification in the risk model of [2]. As such, the 'nature' attribute abstractly represents other specific attributes of the threat events that are classified according to the stages of the lifecycle model of [1]. These attributes, in turn, help assess the extent of damage caused by the adverse consequences and the difficulty with which these consequences can be reversed. Table 2 illustrates the rules for assessing the intensity value of a privacy harm.

The *range* of a privacy harm represents the scope of an adverse impact of a threat event encountered by a variety of subjects. It is based on the 'affected data subjects' attribute and influenced by the 'scope' attribute of the threat event that reflects the domain of the adverse event, assessed, perhaps, in terms of the number and categories of potential subjects whose personal data is impacted.

Table 3. Range of a privacy harm.

Values	Description
1. Negligible	Only specific individuals are affected.
2. Limited	Specific individuals and their relatives and/or friends are affected.
3. Significant	Specific categories of individuals are affected.
4. Maximum	The whole of society is affected.

Table 4. Exploitability of a vulnerability.

Values	The exploitation of the primary asset's vulnerability ...
1. Negligible	does not appear possible.
2. Limited	appears to be difficult.
3. Significant	appears to be possible.
4. Maximum	appears to be extremely easy.

Table 3 illustrates the rules for assessing the value of the range of a privacy harm. These are a refinement of the rules of [6, 7] to distinguish between the range of a privacy harm. The severity of a privacy harm is assessed by adding the assessed value of the intensity and the assessed value of the range, then selecting the overall value according to Table 1.

4.3 Step 3: The Assessment of the Likelihood of a Privacy Harm

The *likelihood* of a privacy harm is the highest value of the overall likelihood of occurrence of associated threat events. The overall likelihood of each threat event is a combination of the likelihood of the threat event occurrence and the likelihood of the threat event resulting in adverse impacts.

The *seriousness* represents the level of a vulnerability, which is based on its 'exploitability' and 'severity'. The 'exploitability' attribute indicates the level of exploitation of a primary asset's vulnerability, whereas the 'severity' attribute indicates the relative importance of mitigating a primary asset's vulnerability. Both are influenced by the attributes of the relevant element of context-relevant processing norms of [2]: personal data, data-processing activities, involved actors and processing principles. These include the attributes of personal data: 'category', 'type', 'sensitivity' and 'linkability'. All these attributes are used to estimate the degree to which the vulnerability of the primary asset can be exploited, and to enable the threat source to conduct adverse actions that breach these norms and violate contextual integrity.

Table 4 illustrates the rules for assessing the value of the exploitability of a vulnerability. Table 5 illustrates the rules for assessing the value of the severity of a vulnerability. The seriousness of the vulnerability is estimated by adding the estimated value of the exploitability and the estimated value of the severity, then selecting the overall value according to Table 1.

The *capability* of a threat source represents the motives, skills, resources or power that make them able to exploit the vulnerabilities of a primary asset.

Table 5. Severity of a vulnerability.

Values	**The successful exploitation of the primary asset's vulnerability . . .**
1. Negligible	leads to insignificant impacts.
2. Limited	leads to slight impacts.
3. Significant	leads to serious impacts.
4. Maximum	leads to severe impacts.

Table 6. Motivation of a threat source.

Values	**The threat source . . .**
1. Negligible	does not have any specific motives based on the value of personal data.
2. Limited	has indistinct and unreasonable motives based on the value of personal data.
3. Significant	has definite and reasonable motives based on the value of personal data.
4. Maximum	has multiple, definite and strong motives based on the value of personal data.

It is mainly estimated based on the values of the 'motivation' and 'ability' as assessable attributes of the threat source, which are determined by the values of 'type', 'resources', 'role' and 'responsibility' as nominal attributes of the source.

The 'motivation' attribute indicates the value of personal data to threat sources that stimulates their motives to exploit vulnerabilities of the primary assets. In general, personal data has value both to data subjects and to the entities that collect and process it. In addition, it has a nuisance value when it is exploited for unfair or malicious purposes. The value of personal data is influenced by the attributes of the personal data as a primary asset: 'category', 'type', 'sensitivity' and 'linkability'. The motive is influenced by the 'type' attribute of the threat source. Table 6 illustrates the rules for assessing the value of the motivation of a threat source.

The 'ability' attribute indicates the level of resources by which a threat source is able to exploit the vulnerabilities of a primary asset. These resources include the skills, background knowledge, privileges, financial and technical resources. These assessable attributes are influenced by the 'type' attribute of the threat source. The 'privileges' is influenced by the roles of the threat source and assigned responsibilities, if any. Further, technical and financial resources are influenced by the 'type' attribute. The 'background knowledge' is influenced by other factors, such as the type of relationship with the data subject. Table 7 illustrates the rules for assessing the value of the ability of a threat source.

The capability of the threat source is estimated by adding the assessed values of motivation and ability, then selecting the overall value according to Table 1. The likelihood of occurrence of a threat event is estimated by adding the assessed values of capability and exploitability of the relevant primary assets' vulnerabilities, then selecting the overall value according to Table 1. The likelihood of the threat event resulting in adverse impacts is estimated by adding the assessed

Table 7. Ability of a threat source.

Values	The threat source ...
1. Negligible	does not appear to have specific skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.
2. Limited	has insufficient skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.
3. Significant	has real and significant skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.
4. Maximum	has definite and unlimited skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.

value of capability and the assessed value of the severity of the relevant primary assets' vulnerabilities, then selecting the overall value according to Table 1. The overall likelihood of occurrence of a threat event causing adverse impacts is estimated by adding the assessed value of the likelihood of occurrence of the threat event and the assessed value of the likelihood of the threat event resulting in adverse impacts, then selecting the overall value according to Table 1. The likelihood of occurrence of a privacy harm is the highest value of the overall likelihood of the corresponding threat events.

4.4 Step 4: The Assessment of the Risk Level of a Privacy Harm

The risk level of a privacy harm is assessed as a (likelihood, severity) pair. These pairs are used to determine the order in which the identified risks should be managed according to their severity and likelihood. We adopt the 'risk map' of [4] to locate the assessed risks according to their levels. The likelihood of a privacy risk is plotted on the X-axis; its severity is plotted on the Y-axis.

5 An Illustrative Example

The European Electronic Toll Service (EETS) [11] aims to support interoperability between Electronic Toll Pricing (ETP) systems at a European level to calculate and collect road-usage tolls. The main actors are users (individuals who subscribes to an EETS provider in order to get access to EETS), EETS providers (legal entities that grant access to EETS to road users), and toll chargers (public or private organisations that are responsible for levying tolls for the circulation of vehicles in an EETS domain). A user is required to provide a set of personal data specified by a responsible toll charger, as well as to be informed about the processing of their personal data. Accordingly, the EETS provider provides the user with an On-Board Unit (OBU) to be installed on-board a vehicle to collect, store, and remotely receive and transmit time, distance and location data over time. In this paper, we assess the values of the attributes of the risk factors identified in [2].

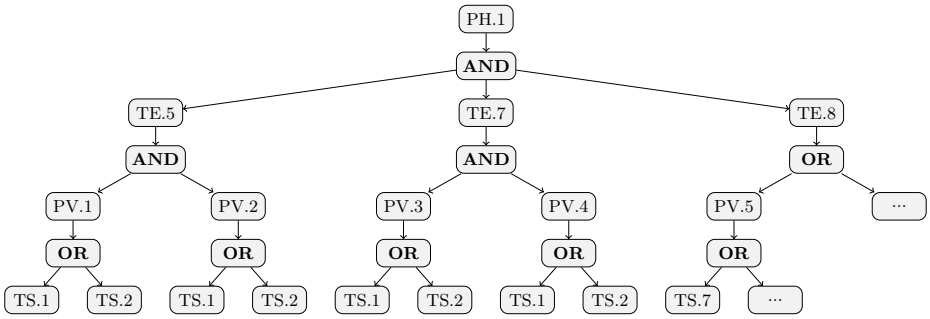

Fig. 2. The structure of the harm tree for the privacy harm PH.1.

5.1 The Construction of Harm Trees

Figure 2 shows the harm tree for the privacy harm PH.1: Increased car insurance premium. It occurs when a EETS provider (TS.1) or a toll charger (TS.2) makes excessive data inference to derive driving patterns (TE.5) for EETS users and shares these driving patterns with car insurance companies (TE.7). An insurance provider (TS.7) makes excessive data inference to re-identify its current or potential customers (TE.8) by linking the derived data to particular drivers to find out their health conditions and vehicle use, or to discover whether a policy holder is where they claim to have been at any point in time. An 'improper data model' (PV.1) and a 'lack of data minimisation' (PV.2) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.5. An 'improper purpose specification' (PV.3) and a 'lack of logs and audit trails' (PV.6) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.7. A 'weak anonymisation technique' (PV.5) that is exploited by TS.7 leads to the occurrence of TE.8.

5.2 Severity Assessment

The intensity of PH.1 is based on the extent of damage caused by, the irreversibility of and the duration of, the adverse consequences of associated threat events TE.5, TE.7 and TE.8. TE.5 is categorised as a type of 'identification'. It is characterised as unanticipated by EETS users; it is also characterised as extensive. TE.7 is categorised as a type of 'disclosure'. It is characterised as extensive and accurate. TE.8 is categorised as a type of 'identification'. It is characterised as unanticipated by EETS users; it is also characterised as extensive. The adverse consequences of these threat events are identified as a series of related impacts started by discovering private facts about EETS users, then revealing sensitive information beyond expected boundaries and ended by charging higher rates of insurance premium. They are assessed as slight because calculating insurance premium depends on other factors, including address, occupation, claims history, etc. The duration of these consequences may last for a certain length of time: the period of cover. However, they may last for longer than the period of cover when the disclosed data is used as 'driving history' by insurance providers to calculate

car insurance quotes. Once profiles are created, disclosed to insurance providers and sensitive information is inferred, it is technically difficult to reverse these consequences. As such, the intensity of PH.1 is assessed as '2. Limited'.

PH.1 may affect specific EETS users based on their driving patterns. They may also affect specific categories of EETS users based on their health conditions. Together, these may impact insurance premiums. Thus, the range of this privacy harm is assessed as '3. Significant'. The severity of PH.1 is assessed as '2. Limited' by adding the assessed value of the intensity and the assessed value of the range, and selecting the overall value according to Table 1.

5.3 Likelihood Assessment

The likelihood of occurrence of TE.5, TE.7 and TE.8 is assessed on the capability of TS.1, TS.2 and TS.7 and the exploitability of PV.1–PV.5 for each possible exploit.

The motivations of TS.1, TS.2 and TS.7 are based on the values of location-related data to those sources and their motives according to these values. The utility of 'identification and contact data' and 'location data' makes such data highly valuable to EETS providers, toll chargers and insurance providers. It also has a nuisance value when it is exploited by insurance providers. Thus, the motivation of TS.1 and TS.2 are assessed as '3. Significant' and the motivation of TS.7 is assessed as '4. Maximum'.

The ability of those sources is based on their skills, background knowledge, privileges, and technical and financial resources. According to the 'type' attribute of TS.1 and TS.2, they are insiders and institutions. This implies that they have technical skills and detailed background knowledge about conceptual, logical and physical data models, as well as about the processing operations. It also implies that they have legitimate privileges to collect and process location-related data according to their roles and responsibilities. EETS providers and toll charges play the roles of data processors and data controllers respectively. Based on these, they have access rights to both the 'fine-grained location data' and 'identification and contact data'. In addition, they have both technical and financial resources to benefit from the values of the collected data by creating comprehensive and identifiable profiles. As such, the abilities of TS.1 and TS.2 are assessed as '4. Maximum' and '3. Significant' respectively.

According to the 'type' attribute of TS.7, they are outsiders and institutions. Insurance providers are third parties that do not have direct roles with respect to the processing of personal data. Based on this, they do not have access rights to the 'fine-grained location data' and 'identification and contact data'; rather, they can legally process this data when it is anonymised. In addition, they have both technical and financial resources to benefit from the values of the disclosed data by making excessive inference. As such, the ability is assessed as '2. Limited'. The capability of TS.1 is assessed as '4. Maximum'; those of TS.2 and TS.7 are assessed as '3. Significant'.

The seriousness of PV.1–PV.5 is based on the exploitability and severity of each exploit. The ease of the exploitation of PV.1 is influenced by the rele-

Table 8. The seriousness of PV.1, PV.2, PV.3, PV4 and PV.5.

Vulnerability	Exploitability	Severity	Seriousness
PV.1	3. Significant	3. Significant	3. Significant
PV.2	3. Significant	3. Significant	3. Significant
PV.3	3. Significant	3. Significant	3. Significant
PV.4	3. Significant	3. Significant	3. Significant
PV.5	3. Significant	4. Maximum	4. Maximum

Table 9. The overall likelihood of occurrence of TE.5, TE.7, TE.8, TE.4 and TE.5.

Code	Exploit	Threat event	L.O.	L.R.	O.L.
EX.1	TE.5 - PV.1 - TS.1	TE.5	4. Maximum	4. Maximum	4. Maximum
EX.2	TE.5 - PV.1 - TS.2		3. Significant	3. Significant	3. Significant
EX.3	TE.5 - PV.2 - TS.1		4. Maximum	4. Maximum	4. Maximum
EX.4	TE.5 - PV.2 - TS.2		3. Significant	3. Significant	3. Significant
EX.5	TE.7 - PV.3 - TS.1	TE.7	4. Maximum	4. Maximum	4. Maximum
EX.6	TE.7 - PV.3 - TS.2		3. Significant	3. Significant	3. Significant
EX.7	TE.7 - PV.4 - TS.1		4. Maximum	4. Maximum	4. Maximum
EX.8	TE.7 - PV.4 - TS.2		3. Significant	3. Significant	3. Significant
EX.9	TE.8 - PV.5 - TS.7	TE.8	3. Significant	4. Maximum	4. Maximum

L.O.: Likelihood of occurrence. **L.R.:** Likelihood of resulting in adverse impacts.
O.L.: The overall likelihood.

vant element of context-relevant processing norms. The relevant element of the processing norms is 'attributes', which refers to personal data. In this context, personal data is classified into two types: 'identification and contact data' and 'location data'. Both types are categorised as 'collected data'. These types of data are not sensitive in themselves; rather, they are valuable and can be used to derive sensitive data, such as driving history or patterns, and health conditions. The fine-grained location data can easily be linked with 'identification and contact data' with reasonable effort as they are modelled, collected and processed by an EETS provider and accessed by a toll charger. The vulnerability of 'improper data model' can be easily exploited by EETS providers and toll chargers. Thus, the exploitability of PV.1 is assessed as '3. Significant'.

The severity of PV.1 is influenced by the relevant element of context-relevant processing norms. The relevant element of the processing norms is 'attributes', which refers to personal data. Thus, PV.1 enables TS.1 and TS.2 to breach the processing norms and violate the contextual integrity by making unjustified inference (TE.5) with the aim of deriving driving patterns for EETS users based on their deriving history. This type of aggregation is a threat event that may lead to the privacy harm PH.1. As such, the severity of PV.1 is assessed as '3. Significant'. The seriousness of PV.2, PV.3, PV.4 and PV.5 are similarly assessed based on the corresponding exploitability and severity, as per Table 8.

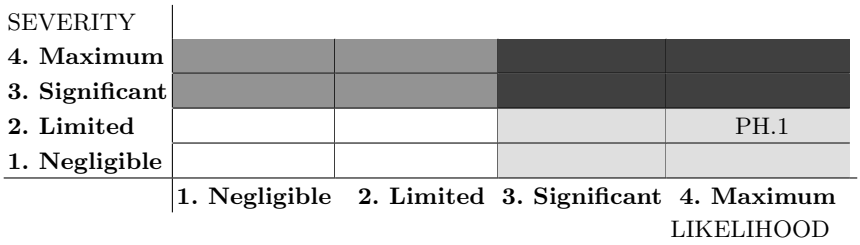

Fig. 3. Risk map in the context of EETS.

The threat event TE.5 may result from the exploitation of PV.1 and PV.2: both vulnerabilities are necessary for the occurrence of TE.5. The vulnerability PV.1 may be exploited by two different threat sources: TS.1 or TS.2. With reference to the harm tree, either one of those sources is sufficient to exploit the vulnerability PV.1. This leads to four possible exploits: EX.1–EX.4. According to the 'AND' connector, two of them are necessary for the occurrence of TE.5. In this case, the highest value of the overall likelihoods of EX.1 and EX.2 is taken: '4. Maximum'. Similarly, the highest value of the overall likelihoods of EX.3 and EX.4 is taken: '4. Maximum'. The overall likelihood of the threat event TE.5 is assessed as '4. Maximum' The overall likelihood of occurrence of all possible exploitations of TE.7 and TE.8 are similarly assessed, as per Table 9. As such, the likelihood of occurrence of PH.1 is assessed as '4. Maximum' by taking the highest value of the overall likelihood of occurrence of TE.5, TE.7 and TE.8.

5.4 Risk Level Assessment

The risk level of PH.1 is assessed as a pair of likelihood and severity: (4. Maximum, 2. Limited). Figure 3 shows that PH.1 corresponds to 'limited risks'.

6 Conclusions

We have presented a privacy risk-assessment approach that expands the PRIAM methodology in a number of ways. It assesses the level of a vulnerability by analysing its exploitability and severity, which are influenced by the characteristics of the main elements of context-relative processing norms. It also assesses the motivation of a threat source by analysing the value of personal data and the motives behind conducting data-processing activities that may lead to privacy harms in the given context. Third, it assesses the irreversibility of a threat event. Finally, it uses fixed levels of scale that for multiple stakeholders.

We will build upon this work in a number of ways. First, we will use additional case studies to further validate the approach and highlight its usefulness and practical impact in various domains. Second, we intend to identify a risk-assessment process to underpin a privacy-risk assessment methodology that can complement PIA processes. We also plan to use such a methodology as a means for managing the assessed privacy risks in a structured manner.

Acknowledgments. The authors would like to thank the reviewers for their constructive comments.

References

1. Alshammari, M., Simpson, A.C.: Personal data management: An abstract personal data lifecycle model. In: Teniente, E., Weidlich, M. (eds.) Proceedings of the 2017 International Conference on Business Process Management (SPBP 2017). Lecture Notes in Business Information Processing, vol. 308, pp. 685–697. Springer (2017)
2. Alshammari, M., Simpson, A.C.: Towards an effective PIA-based risk analysis: An approach for analysing potential privacy risks. `http://www.cs.ox.ac.uk/publications/publication11663-abstract.html` (2017)
3. Clarke, R.: An evaluation of privacy impact assessment guidance documents. International Data Privacy Law 1(2), 111–120 (2011)
4. Commission Nationale de l´Informatique et des Libertés: Methodology for Privacy Risk Management (How to implement the Data Protection Act). `https://www.cnil.fr/sites/default/files/typo/document/CNIL-ManagingPrivacyRisks-Methodology.pdf` (2012)
5. European Commission: Data Protection Impact Assessment Template for Smart Grid and Smart Metering systems. `http://ec.europa.eu/energy/sites/ener/files/documents/2014_dpia_smart_grids_forces.pdf` (2014)
6. Joyee De, S., Le Métayer, D.: PRIAM: A privacy risk analysis methodology. In: Livraga, G., Torra, V., Aldini, A., Martinelli, F., Suri, N. (eds.) Proceedings of the 11th International Workshop on Data Privacy Management (DPM 2016) and 5th International Workshop on Security Assurance (QSA 2016). Lecture Notes in Computer Science, vol. 9963, pp. 221–229. Springer (2016)
7. Joyee De, S., Le Métayer, D.: A refinement approach for the reuse of privacy risk analysis results. In: Schweighofer, E., Leitold, H., Mitrakas, A., Rannenberg, K. (eds.) Proceedings of the 5th Annual Privacy Forum (APF 2017). Lecture Notes in Computer Science, vol. 10518, pp. 52–83. Springer (2017)
8. National Institute of Standards and Technology (NIST): Privacy Impact Assessment (PIA). `https://www.nist.gov/document/nist-tip-pia-consolidatedpdf` (2012)
9. Oetzel, M.C., Spiekermann, S.: A systematic methodology for privacy impact assessments: A design science approach. European Journal of Information Systems 23(2), 126–150 (2014)
10. Solove, D.J.: A taxonomy of privacy. University of Pennsylvania Law Review 154(3), 477–564 (2006)
11. The European Commission: The European Electronic Toll Service (EETS): 2011 Guide for the Application of the Directive on the Interoperability of Electronic Road Toll Systems. `http://ec.europa.eu/transport/themes/its/road/application_areas/electronic_pricing_and_payment_en` (2011)
12. UK Information Commissioner's Office (ICO): Privacy Impact Assessment Handbook. `www.adls.ac.uk/wp-content/uploads/2011/08/PIA-handbook.pdf` (2009)
13. Wright, D.: Making Privacy Impact Assessment more effective. The Information Society 29(5), 307–315 (2013)
14. Wright, D. and Wadhwa, K.: A Privacy Impact Assessment framework for data protection and privacy rights. `http://www.vub.ac.be/LSTS/pub/Dehert/507.pdf` (2011)