Privacy Risk Assessment for Data Subject-aware Threat Modeling

Laurens Sion, Dimitri Van Landuyt, Kim Wuyts, Wouter Joosen
imec-DistriNet, KU Leuven
Heverlee, Belgium
{laurens.sion, dimitri.vanlanduyt, kim.wuyts, wouter.joosen}@cs.kuleuven.be

***Abstract*—Regulatory efforts such as the General Data Protection Regulation (GDPR) embody a notion of privacy risk that is centered around the fundamental rights of data subjects. This is, however, a fundamentally different notion of privacy risk than the one commonly used in threat modeling which is largely agnostic of involved data subjects. This mismatch hampers the applicability of privacy threat modeling approaches such as LINDDUN in a Data Protection by Design (DPbD) context.**

In this paper, we present a data subject-aware privacy risk assessment model in specific support of privacy threat modeling activities. This model allows the threat modeler to draw upon a more holistic understanding of privacy risk while assessing the relevance of specific privacy threats to the system under design. Additionally, we propose a number of improvements to privacy threat modeling, such as enriching Data Flow Diagram (DFD) system models with appropriate risk inputs (e.g., information on data types and involved data subjects). Incorporation of these risk inputs in DFDs, in combination with a risk estimation approach using Monte Carlo simulations, leads to a more comprehensive assessment of privacy risk.

The proposed risk model has been integrated in threat modeling tool prototype and validated in the context of a realistic eHealth application.

***Index Terms*—privacy, privacy by design, data protection by design, GDPR, threat modeling, risk assessment, privacy risk**

I. Introduction

The principle of Privacy by Design (PbD) is increasingly recognized as paramount for the realization of privacy-preserving software. Besides the growing awareness of privacy concerns due to increasingly impactful data breaches, its importance is also confirmed with the introduction of legislation and guidelines such as the EU's General Data Protection Regulation (GDPR) [1], the OECD Privacy Guideline [2], [3], and the Generally Accepted Privacy Principles (GAPP) [4], all of which advocate explicit privacy risk management. The GDPR [1] even imposes it, as it requires countermeasures proportional to the risk to the involved data subjects [1, Art. 32]. Hence, privacy risk assessment becomes an essential part of a comprehensive privacy engineering approach.

An important class of solutions for system analysis from a privacy perspective is threat modeling, which entails the systematic enumeration of misuse and attack vectors and considering their applicability in the system under design. Successful implementations of security threat modeling [5]–[8] have led to the conception of counterparts for eliciting privacy threats, the most notable methodology being LINDDUN [9].

It is in the prioritization of the uncovered privacy issues that threat modeling and privacy risk assessment are mutually reinforcing approaches. However, privacy is an inherently contested concept [10]; its risk can be approached from different perspectives: (i) legal risk involving data protection aspects; (ii) economic risk focusing on financial losses or reputational damage; (iii) societal risk in terms of fundamental rights of citizens or societal notions such as social cohesion; (iv) software engineering risk with approaches such as threat elicitation involving notions as attacker capabilities, threat feasibility, involved assets, countermeasure strengths, and engineering trade-offs; and so on.

Existing risk assessment methodologies in a privacy engineering context commonly focus on a narrow perspective, such as asset values, or are confined to a limited high-level assessment, but lack focus on data subjects. By creating a detailed risk decomposition, the involved risk factors are made explicit, leading to a more precise interpretation. Furthermore, a detailed decomposition provides support for a more fine-grained calculation of the resulting risk. Finally, retrieving the risk inputs from engineering models, automation can be supported, enabling an encompassing risk management approach that keeps track of the global reduction of privacy risk across multiple countermeasures and design iterations, allows for better traceability and auditability.

In this paper we (i) present a privacy risk decomposition to calculate privacy risk using Monte Carlo simulations, (ii) parameterize the risk to support different analysis scenarios, (iii) elaborate on the integration of the risk assessment model in a threat modeling context, (iv) implement the presented extensions in a prototype, and (v) apply it on an eHealth application illustrating its use in risk analysis scenarios.

This paper is structured as follows. Section II provides some background on privacy threat modeling. Section III presents the privacy risk assessment model. Section IV introduces the necessary threat modeling extensions for integrating privacy risk assessment. Section V validates the extension in a prototype and on an eHealth application case. Section VI provides a discussion and Section VII discusses related work. Finally, Section VIII concludes the paper.

This research is partially funded by the Research Fund KU Leuven and the PRiSE KU Leuven-C2 research project.

https://doi.org/10.1109/SPW.2019.00023

II. Background

This section first provides the necessary background on privacy threat modeling and then discusses different perspectives on privacy risk, leading up to the problem statement.

A. Privacy Threat Modeling

Privacy threat modeling methodologies, such as LINDDUN [9], [11], represent a class of architecture-level analysis methods, tools, and techniques that involve systematically assessing the applicability of known privacy-related issues (threats types) in the context of a specific system under design.

As shown in the pseudo-code below, the threat elicitation phase commonly involves four activities: (i) modeling the system (line 1), (ii) systematically iterating over the model elements (line 2), (iii) iterating over the known threat types (line 3), and (iv) based on the applicability of the threat type to the system element (line 4) and the perceived risk (line 5), documenting the identified privacy threats (line 6), which are then to be mitigated in later phases.

```
1 SystemModel systemModel
2 for each sc in systemModel:
3  for each tt in ThreatTypes:
4   if (tt.applicable(sc) &&
5       Risk(tt,sc) > threshold):
6    document(tt, sc)
```

Such an exhaustive threat elicitation approach is enumerative and therefore suffers from combinatorial explosion—the amount of threats to consider grows substantially with the number of system elements and the number of threat types to consider. In this context, privacy risk assessment is crucial to ensure the cost-effectiveness and efficiency of threat modeling approaches in general.

B. Perspectives on Privacy Risk

As explained earlier, privacy risk can be assessed from a wide range of different perspectives. This section elaborates on a number of risk perspectives that are relevant in the context of privacy threat modeling. Many existing risk assessment approaches focus on either technical failures (e.g., FMEA [12]) or the manifestation of security threats (e.g., FAIR [13], CORAS [14], security threat risk [15]). In these approaches, the risk impact depends on the value of business assets or the level of criticality of technical components or services.

These risk assessment models do not, however, include an assessment of the potential privacy impacts on data subjects. The GDPR and other applicable regulations dictate adopting a risk-based approach, and specifically advocate the execution of Data Protection Impact Assessments (DPIA), which fundamentally weigh the privacy impact against data subjects' fundamental rights. PRIAM [16] provides a much more detailed view on privacy risk to the data subjects, using privacy harm trees to assess the risk using the feared events, risk sources, and weaknesses. Other risk assessment models [17]–[19] focus on assessing the risk specifically from the point of view of a single data subject or user.

C. Problem statement

Threat modeling in practice is approached mainly from a security perspective, and despite many of the similarities between security and privacy as non-functional concerns, merely adopting security-centric risk assessment models (focused on factors such as assets, value, impact, technical feasibility) leads to an incomplete characterization of privacy risk: more notably, privacy threat modeling approaches lack awareness of the impact on the involved data subject types.

Furthermore, existing approaches are coarse-grained and provide limited support for traceability and repeatability of the resulting risk values (e.g., to find out the main contributing factors to a specific risk value). This, however, is essential for (i) calibration, e.g., to allowing analysis why different experts may reach different risk values in their assessment, (ii) strengthening the understanding of privacy risk, i.e. towards understanding which parameters (system context, type of attacker, involved data subjects, etc.) actually impact privacy risk the most visibly in a specific case, but also (iii) auditability and compliance reasons, i.e. to demonstrate that a suitable risk-based approach was taken.

III. Risk Assessment Model

This section elaborates on the proposed privacy risk assessment model that extends FAIR [13] with specific privacy and data subject risk factors. It is decomposed following the structure from Figure 1 from left to right and top to bottom. By decomposing privacy risk, the model unifies: (i) data subject risk by incorporating information on data subjects and types of their data being processed, (ii) technical risk originating from the system context and applicable security and privacy countermeasures, and (iii) the risk from an organizational perspective by scaling the impact according to the number data subjects and records involved.

A. Risk

The risk is decomposed into two underlying factors: (i) the *Loss Magnitude*, which represents the impact on the data subject(s); and (ii) the *Loss Event Frequency*, which represents the frequency of successful attacks from an adversary. These factors need to be multiplied to calculate the overall risk.

$$\begin{aligned}\boldsymbol{Risk} &= \boldsymbol{LM} \odot \boldsymbol{LEF} \\ &= [LM_1 \times LEF_1, LM_2 \times LEF_2, \ldots, LM_S \times LEF_S]\end{aligned}$$

B. Loss Magnitude (LM)

The *Loss Magnitude*, representing the impact, is decomposed of the following four factors: (i) *Data Type Sensitivity*, (ii) *Nbr. of Records*, (iii) *Data Subject Type*, and (iv) *Nbr. of Data Subjects*. These factors comprise both the impact derived from the involved data types as well as the involved data subjects. Each of these factors are discussed in more detail below. The loss magnitude can be obtained by multiplying them all together as follows:

$$LM_i = DTS_i \times NR_i \times DST_i \times NDS_i$$

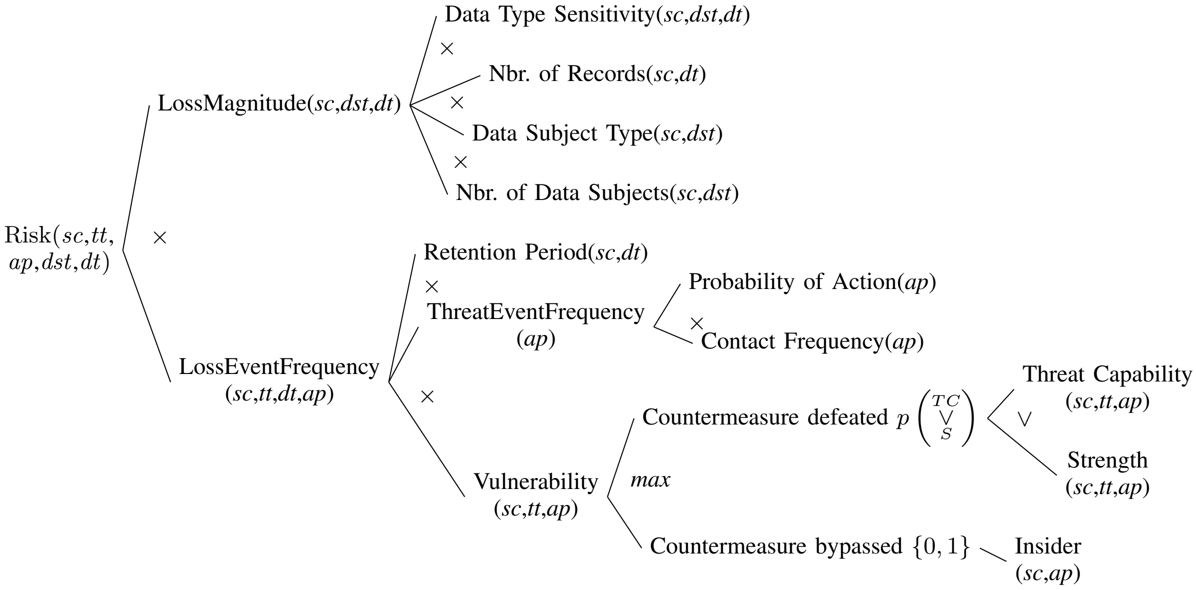

Fig. 1. Risk Decomposition

This figure shows the decomposition of how the risk can be calculated for a single system context (sc), a single threat type (tt), a single attacker profile (ap), a single data subject type (dst), and a single data type (dt). The risk values can be aggregated along these parameters as needed, which, for the total risk, would lead to: $\sum^{SC}\sum^{TT}\sum^{AP}\sum^{DST}\sum^{DT}\mathrm{Risk}(sc, tt, ap, dst, dt)$

Data Type Sensitivity (DTS) The *Data Type Sensitivity* represents the privacy risk inherent to the types of data that are being processed in the system. To distinguish between data types of different sensitivity levels, data types can be ordered on a numerical scale according to their sensitivity. This allows the risk assessment to factor in the impact of threats involving sensitive data such as medical information, in contrast to, for example, contact information such as home addresses.

Number of Records (NR) This factor represents the *number of records* of a certain data type for a certain data subject type. This value can be used for expressing two cases. First, if multiple records of a data type are being collected or processed, the risk value can be scaled appropriately with this factor. Second, if a value of this data type is only present for a fraction of the data subjects (e.g., only processed for half the data subjects), a fraction for this factor can be used to scale the risk value down accordingly.

Data Subject Type (DST) The *Data Subject Type* is used to specify the risk inherent to the type of data subject whose data are being processed. This factor is used to take into account special cases of vulnerable data subject types such as minors.

Nbr. of Data Subjects (NDS) This factor represents the *number of data subjects* of a certain type (i.e. the DST above). This is a scaling factor, analogous to the number of records, so the impact can be scaled according to the number of data subjects involved in the data processing operations.

C. Loss Event Frequency (LEF)

The second factor of the risk is the *Loss Event Frequency*. It represents the total frequency of successful attacks by an adversary. This frequency is obtained by combining the frequency of attacks (TEF) with the probability of a successful attack (V). For an attacker a, the LEF is calculated as follows:

$$\boldsymbol{LEF} = V \cdot (\boldsymbol{RP} \odot \boldsymbol{TEF})$$

D. Retention Period (RP)

The *Retention Period* represents the duration during which data is stored or processed and present for an adversary to be potentially exploited. This enables distinguishing between long-running processing operations that pose a higher risk versus very short-lived transactions after which the data is no longer retained. A threat event can only be successful when the data is available at the time of the attempted attack.

E. Threat Event Frequency (TEF)

The *Threat Event Frequency* represents the frequency of attempted attacks by an adversary. These attacks are not necessarily successful. The threat event frequency is further decomposed into: (i) the *probability of action*, representing the likelihood of an attempted attack, and (ii) the *contact frequency*, representing how frequently the adversary comes into contact with the system. The threat event frequency is obtained by multiplying these two factors together:

$$TEF_i = PoA_i \times CF_i$$

Probability of Action (PoA) The *Probability of Action* is used to determine the likelihood that an adversary will attempt to attack users' privacy. This probability will depend on the type of adversary (which in turn is based on its incentives, capabilities, and opportunities). For example, an external remote adversary could be more likely to attempt attack when coming into contact with the system, while an insider—an employee, for example—may be more or less likely to attempt an attack depending on the monitoring controls imposed on employees and possible repercussions when discovered.

Contact Frequency (CF) The *Contact Frequency* is the frequency with which an adversary comes into contact with the system. It again varies between different types of adversaries, allowing to make the distinction between, for example, external adversaries, users of the system, or insiders.

F. Vulnerability

The *Vulnerability* is the probability of a successful attack (taking into account the possibility of the adversary being an insider). The vulnerability is calculated as the maximum of: (i) the countermeasure being defeated (*CD*), and (ii) the countermeasure being bypassed (*CB*):

$$V = \max(CD, CB)$$

Countermeasure Defeated (CD) The *Countermeasure Defeated* factor is obtained by sampling from both the threat capability and strength distributions to calculate the fraction of successful attacks in which the adversary manages to defeat the countermeasures present. This calculation provides the probability of a successful attack. A single sample i is calculated as follows:

$$CD_i = f(TC_i, \mathrm{S}_i) \text{ with } f(x,y) = \begin{cases} 1 & x \geq y \\ 0 & x < y \end{cases}$$

Since we need the probability of the adversary defeating the countermeasure, S samples are aggregated as follows:

$$CD = \frac{\sum_{i=1}^{S} CD_i}{S}$$

Threat Capability (TC) The *Threat Capability* expresses the capability of the adversary in being able to defeat the technical security and privacy countermeasures.

Strength (S) This indicates the strength of a technical countermeasure in resisting an adversary. The strength of a countermeasure should be specified on the same scale as the capability of adversaries. More specifically, a countermeasure can resist an adversary if its strength is larger than the threat capability of the adversary (*S* > *TC*).

Countermeasure bypassed (CB) This factor indicates whether the adversary can bypass the measure as an insider, without needing the threat capability to technically defeat the measure. While the simplest representation of this factor is binary (0/1), it could also be represented as a probability of being able to bypass a countermeasure as insider.

G. Risk Factor Values

In order to facilitate the calculation of the risk for privacy threats, numeric values are required as inputs in the risk assessment calculation. Such a requirement raises the issue of determining the appropriate values, which can be difficult. Our approach explicitly supports and takes into account uncertainty about these values. Every numeric value used as an input for risk assessment is represented as an estimate with four parameters: the *minimum* value, the *maximum* value, the *most probable* value, and a *confidence* level value.

These four values define a modified PERT distribution [20], a distribution commonly used in risk management for managing the uncertainty in expert estimates.

By using this distribution, a wide range of values with differences in certainty can be expressed. For example, in case only the outer boundaries are known, the minimum and maximum value can be provided, and the confidence can be set to zero. This leads to a uniform distribution between the provided minimum and maximum values. When there is a high degree of certainty, closer values and high confidence lead to a distribution with a sharp peak.

H. Parameters

The previous sections elaborated on the individual risk factors. These factors cannot be determined for the system overall, as they depend on specific parameters. This section elaborates on the parameters that provide the necessary information for determining the factors for a single risk value.

System Context The first parameter that needs to be fixed is the system context. There can be large local differences in the system context. A localized risk value needs to take the precise context (such as local security or privacy countermeasures) into account. Afterwards, these specific risk values can be aggregated to obtain the risk for the whole system.

Threat Type Second, is the considered threat type. Different security or privacy threat types are not always applicable. Both the local system context and the threat type itself determine the applicability of a threat. Results can again be aggregated over all threat types per category or over all types in general.

Attacker Profile Third, the assessment has to be performed while considering a specific attacker profile to take into account different attackers with different capabilities. This is also essential for being able to make the distinction between external attackers and insiders. The attacker profile-specific risks can again be aggregated afterwards.

Data Subject Type Fourth, is the type of data subject. The risk is calculated for a single type of data subject. Again, multiple risk values can be aggregated to obtain a risk value for all types of data subjects.

Data Type Finally, the risk is calculated for a single data type (belonging to one or more data subjects). For invalid (data type, data subject type) parameter combinations (i.e. the data type does not belong to that data subject type), the resulting risk is zero. The risk values can again be aggregated over this parameter as well.

To obtain a total risk value, the resulting risk values can be aggregated over the previous five parameters. Other types of aggregation are also possible. By combining the risk values over all parameters *but* the data type, an overview of the data types and their associated risk can be obtained. Analogous analyses are possible for the data subject types, attacker profiles, etc.

IV. Impact on Threat Modeling

This section elaborates on the enhancements necessary to support the automated risk assessment in a threat modeling context, aligning these with the risk analysis parameters discussed in Section III-H: *System Context*, *Threat Type*, *Attacker Profile*, *Data Subject Type*, and *Data Type*.

A. DFD Model Extensions

The presented risk analysis model relies on a *System Context*. This parameter corresponds with the system in PRIAM [16]. For risk assessment in a threat modeling context, the same DFD system representation [21] used in traditional security and privacy threat modeling approaches [5], [7], [9], [11] can be relied upon. However, it does need to be extended with support for the representation of security and privacy countermeasures in DFDs [22] so that their effect can be incorporated as well.

B. Threat Types and Attacker Profiles

The *threat types* are already part of the STRIDE [7], [23] and LINDDUN [9], [11] threat modeling approaches. They are similar to the privacy weaknesses in PRIAM [16] (as they lead to privacy harms), and the vulnerabilities in CORAS [14] The traditional threat elicitation step already considers every threat type while iterating over the DFD model elements or interactions, no additional risk extensions are required.

In addition to the threat types, the risk assessment requires awareness of *attacker profiles*. These attacker profiles specify different types of adversaries against which to protect. They correspond with the risk sources from PRIAM [16] and the threats in CORAS [14]. The attacker profiles require additional inputs as discussed in Section III-G. The *Insider* property of an attacker profile is represented as a list of DFD elements for which the attacker can circumvent the countermeasures.

C. Data Subject Types and Data Types

The presented risk model is tailored for system-specific privacy threats. For assessing the *Loss Magnitude*, i.e. the impact on the involved *Data Subjects*, integration with a data protection perspective [24] is required. The data protection viewpoint [24] includes information on data subject types and data types, including the sensitivity of data types. By leveraging the correspondences between the data protection viewpoint and the DFD model, the relevant information can be extracted and used in the context of the privacy risk assessment.

While data subjects have no direct representation in PRIAM [16], they are included as victim in the privacy harm attributes and as stakeholder (although stakeholders also include controllers, third parties, etc.). CORAS [14] does not support data subjects. Personal data is supported in PRIAM [16] as part of the information gathering phase. It is, however, not an explicit factor in the risk assessment phase. CORAS [14] does not support data types, unless they are modeled as assets, which would still lack a link to the data subjects.

Finally, the threat modeling pseudo-code introduced in Section II-A can be extended to include the additional information from the presented parameters from Section III-H:

```
SystemModel systemModel
for each sc in systemModel:
 for each tt in ThreatTypes:
  for each ap in AttackerProfiles:
   for each dst in DataSubjectTypes:
    for each dt in DataTypes:
     if (tt.applicable(sc) &&
         Risk(sc,tt,ap,dst,dt) > threshold):
      document(sc,tt,ap,dst,dt)
```

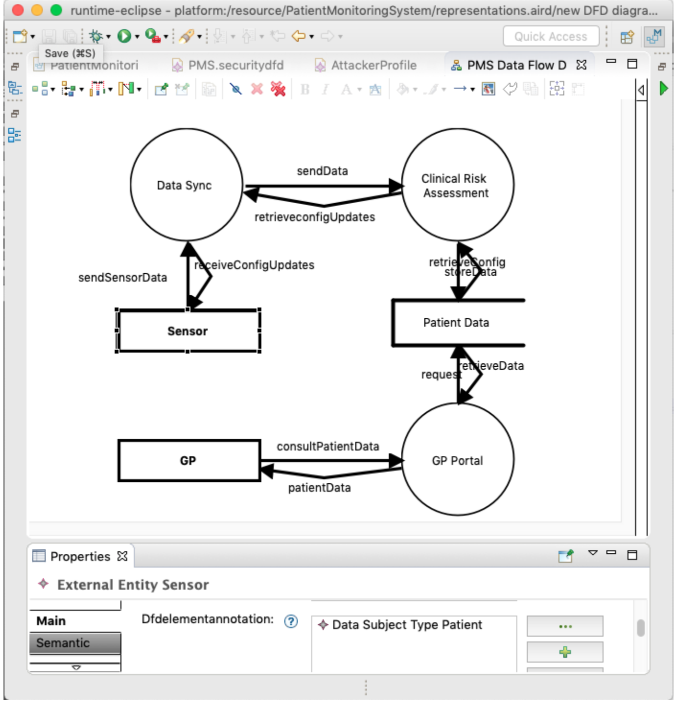

runtime-eclipse - platform:/resource/PatientMonitoringSystem/representations.aird/new DFD diagra...

Properties

External Entity Sensor

Dfdelementannotation: Data Subject Type Patient

Fig. 2. Screenshot of the prototype
This screenshot shows the DFD of the Patient Monitoring System. It illustrates how a traditional DFD lacks any information on data protection concepts. The properties pane shows how the extension links the Patient DataSubjectType *to the Sensor* External Entity *in the diagram. Similar links are present for all the other elements to capture which data types of which data subjects move through the system.*

V. Validation & Illustration

First, the prototype implementation of the risk assessment model is discussed. Next, the prototype implementation is used to apply the risk assessment on an eHealth application, followed by a description of potential risk analysis scenarios.

A. Prototype Implementation

To evaluate the feasibility of the presented risk assessment model, we implemented a proof of concept. Figure 2 shows a screenshot of the prototype implementation. The prototype implements: (i) the presented risk assessment model, (ii) the necessary threat modelling enrichments, and (iii) the integration with the data protection view [24] by extending previously developed tool support for security threat modeling [25].

The prototype uses Eclipse Ecore meta-models for representing the DFD model, threat types, attacker profiles, and the data protection view. They are extended with the necessary

TABLE I
EXCERPT OF INDIVIDUAL RISK ASSESSMENT RESULTS ACROSS THE FIVE PARAMETERS

System Context (DFD)	Threat Type	Attacker Profile	Data Subject Type	Data Type	Risk
storeData	Linkability	Motivated, limited capability	Patient	ECG Measurement	4.542
storeData	Identifiability	Motivated, limited capability	Patient	Risk level	6.632
storeData	Detectability	Opportunist	Patient	Risk level	10.409
patientData	Disclosure of Information	Opportunist	Patient	Body temp measurement	3.66
retrieveData	Linkability	Motivated, capable	Patient	Risk level	1.084
retrieveData	Detectability	Disgruntled employee	Patient	Risk level	0.035
GP	Detectability	Motivated, capable, organized	General Practitioner	Credentials	0.423
…	…	…	…	…	…
… *(6793 rows omitted)*					.

For a system with 16 *DFD elements,* 6 *threat types,* 5 *attacker profiles,* 2 *data subject types, and* 4 *data types. These entries can be aggregated across the different dimensions to gain insights into which parameters are the biggest contributors to the privacy risk.*

properties from Section IV. To perform the risk assessment, these models are queried with patterns. These model query patterns are defined in VIATRA and support querying concrete models for: (i) applicable threats, (ii) data subject types, (iii) data types, and (iv) the mapping from data types to the DFD elements where they are processed or stored.

B. Application on an eHealth Application

The resulting prototype implementation is used to apply the risk assessment on a concrete eHealth case. The eHealth application case is a Patient Monitoring System for monitoring cardiovascular disease patients. Patients are equipped with wearable sensors that measure health parameters, such as body temperature and ECG. Those health parameters are communicated via a mobile app to the back-end, which will perform a clinical risk assessment based on the received information. The analysis results are subsequently made available to a general practitioner (GP) via the GP Portal. Figure 2 shows a screenshot of the application prototype with the DFD of the patient monitoring system. Besides the (visualized) DFD model, there is a corresponding data protection model containing the information on the data subject types and the data types, including links to all the DFD elements where the corresponding data types are being processed or stored.

Running the privacy assessment on this application case results in Table I. Each row of this table corresponds with risk documentation step in line 9 of the pseudo-code in Section IV. The table provides a very fine-grained view on the risk associated with each combination of parameters.

Once the data in Table I is calculated, it can be aggregated in multiple different ways to offer interesting insights into where high risk is situated and which parameters are the biggest contributors to such high risk.

Figure 3 provides a visualization of such an analysis activity. By aggregating the risk per system context (column 1 in Table I) and data subject type (column 4 in Table I), i.e. by aggregating for every combination of these two parameters, an overview is obtained of where the risk is the highest for every (system context, data subject type)-pair. This aggregated information can be visualized as a heatmap, by overlaying a 2d density plot on top of the DFD, making it easily detectable which system

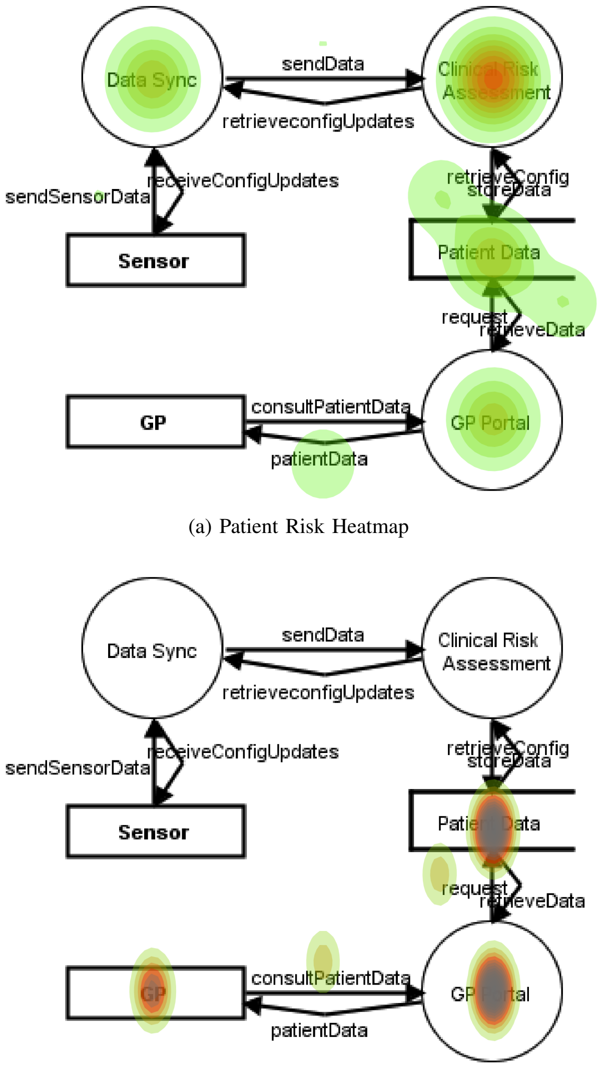

(a) Patient Risk Heatmap

(b) General Practioner Risk Heatmap

Fig. 3. Heatmaps of data subject type risks
The two images illustrate the distribution of risk in the DFD for different data subject types. The heatmaps are constructed by overlaying a 2d density plot on top of the DFD. Figure 3a shows the distribution of the patient risk. Note that the sendData *and* sendSensorData *data flows in this diagram do have a non-zero risk value, but it is very small. Figure 3b shows the distribution of risk of the general practioner.*

elements have the highest risk associated with them for each data subject type. This can assist in prioritizing privacy efforts.

While the heatmaps in Figure 3 visualize the risk aggregates for the different data subject types, such visualizations could also be created for any other combination of parameters. This supports other analyses of the risk distribution for, for example, different data types, attacker profiles, etc.

VI. Discussion and Future Work

This section discusses the implications of the presented privacy risk assessment model and outlines our future work in this context. First, Section VI-A outlines approaches towards accurate estimation and calibration of the individual risk factors. Then, Section VI-B discusses the extent to which the underlying assumption of independence of the involved risk factors holds in a realistic context. Finally, Section VI-C discusses the value and relation of the present risk model in the context of risk assessment activities specifically aimed at legal compliance.

A. Estimation and calibration of risk factors

Any risk assessment model that involves estimation of individual risk factors depends highly on the correctness and accuracy of these input values. While this is no different in the proposed privacy risk assessment model, the presented approach does explicitly take into account uncertainty by representing input values as estimates to parameterize a distribution from which to sample. This explicitly supports taking into account various ranges of input values for the factors.

Furthermore, the numerical values can also be used to rank elements for a relative ordering. For example, different data types can have ordered sensitivity values associated with them. Such assignments could be reused by collecting them in a data type catalog. These catalogs can contain assignments These could be provided in a data type catalog, with values assigned according to, for example, GDPR sensitivity interpretations, to allow easy reuse across multiple models, or be constructed from the severity scale from the CNIL PIA knowledge bases [26].

B. Independence of the Risk Factors

As discussed in Section III-H, the current risk assessment model is rooted upon the assumption that the risk factors are independent from each other. While such an assumption greatly simplifies the risk calculation, the reality is, unfortunately, more complex. Below, we provide some example illustrations of dependencies between these factors:

ThreatType–AttackerProfile: The capability and attack frequency can vary depending on the considered threat type. An external adversary may be more likely to attempt a linkability attack, while an insider may be more likely to identify users as some countermeasures could be bypassed by this adversary.

ThreatType–DataTypeSensitivity: The impact of a certain threat type manifesting itself may depend on both the threat type and the data type sensitivity. For example, the information disclosure of a certain data type may have a bigger impact (on the data subject) than a detectability threat. In other cases, the reverse may be true. For example, the result of a medical test may be negative (with limited information disclosure impact), while detecting that this information is in a database of test results for certain medical condition may have a bigger impact.

AttackerProfile–DataTypeSensitivity: The sensitivity of data types may vary depending on the adversary. For example, medical data may be considered more sensitive when the external adversary is the insurance company compared to, for example, a doctor at a hospital who is not authorized to look at other patients' records.

The independence of the factors reduces the amount of information that is (and needs to be) available for each factor. Necessarily, the resulting risk score will be less precise. This issue can be partially mitigated by choosing the boundaries of the provided estimates in such a way that the variation (because of the other factors) is still captured.

The amount of detail in the risk factors involves a necessary trade-off exercise. Each of the risk analysis parameters (Section III-H) could be moved completely down to every risk component (Section III). This would, however, require end-users to enter a prohibitively large amount of information as all combinations must be considered for every factor.

In future work, we intend to model the causality of these (and potentially other underlying) factors, to evaluate whether a different and independent set of risk factors can be constructed to improve the precision of the risk assessment without sacrificing usability in the number of required inputs.

C. Compliance Checks

The current privacy risk decomposition is very suitable for the risk assessment of the 'hard privacy' threats in LINDDUN (i.e. LINDD); it is much less suitable for assessing the risk of the 'soft privacy' threats (*Non-compliance* or *Unawareness*) as these require very different types of inputs.

For example, assessing the non-compliance risk closely aligns to Data Protection Impact Assessment (DPIA) exercises. Given the integration of the engineering view (Section IV-A) with a data protection view (Section IV-C) [24], the information in that view on data subjects and data types can be leveraged for conducting compliance assessment activities. Repeating such assessment activities for every part of the system enables a localized non-compliance risk assessment.

VII. Related Work

Beckers [27] compared multiple privacy requirements engineering approaches. None of the considered approaches support the notion of risk. Risk is, however, explicitly required by privacy regulations such as the GDPR [1].

Heckle and Holden's [28] findings suggest that neither privacy impact assessments (PIAs) nor classic risk analysis models are sufficient for privacy risk assessment in the context of voting systems. Abu-Nimeh and Mead [29] propose combining them by the IRS PIA [30] in Security Quality Requirements Engineering (SQUARE) [31]. While such a PIA [30] supports a detailed assessment of the realization of privacy-by-policy in the framework of Spiekermann and Cranor [32] given its focus on assessing compliance with privacy principles, a set of questions may not be the best approach. Alshammari and Simpson [33] make the case for a model-based approach for privacy compliance checking. The incorporated data protection view [24] supports such a model-based compliance assessment. Furthermore, its integration in the risk assessment provides support for assessing the risk of privacy threats such as identifiability and linkability, supporting the realization of privacy-by-architecture [32].

PRIAM [16] provides a very detailed description of information that needs to be collected for the privacy risk assessment. The risk assessment itself requires the construction of harm trees, in which the risk is assessed with the combination of privacy weaknesses and risk sources for feared events which can lead to the harm at the top of the tree. Our approach can be considered a kind of instantiation of this approach, but explicitly requires the assignment of numerical estimates for the risk factors. By requiring such numerical assignments, a completely automated assessment can be performed.

Hong et al. [34] presented a privacy risk model specifically developed for ubiquitous computing systems, focusing on the selective disclosure of personal information (*personal privacy*). Similar as the IRS PIA [30], a set of questions is used for the privacy risk analysis, after which the risks are prioritized.

VIII. Conclusion

In this paper, we presented a privacy risk assessment model that is firmly embedded in a privacy threat modeling context. It thus assumes a software construction point of view yet involves extensive analysis of the privacy implications imposed on data subjects. As such, the privacy risk assessment model allows for a more comprehensive, systematic, and data subject-aware privacy threat assessment. By enriching elicited privacy threats with risk analysis information, privacy engineering efforts can be prioritized and appropriate countermeasures, in line with the risk posed to data subjects, can be determined.

The focus on data subjects and ensuing privacy risk implications is essential to align threat modeling activities with compliance requirements imposed by regulations such as the GDPR. Explicit breakdown of the overall risk involved in a data processing effort allows for a more fine-grained risk assessment, sensitivity analysis of the impact of various parameters on the resulting privacy risk, follow-up and management of overall privacy risk, not only at development or system construction time, but also in the context of system operation and evolution.

References

[1] European Union, “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016,” *Official Journal of the European Union*, vol. 59, no. L 119, pp. 1–88, may 2016.

[2] OECD, “The OECD Privacy Framework,” *Organisation for Economic Co-Operation and Development*, pp. 1–154, 2013. [Online]. Available: http://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf

[3] ——, “OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,” *OECD*, 1980.

[4] American Institute of Certified Public Accountants Inc. and Canadian Institute of Chartered Accountants, “Generally Accepted Privacy Principles,” no. August, pp. 1–84, 2009.

[5] M. Howard and S. Lipner, *The Security Development Lifecycle*, 2006.

[6] A. Shostack, “Experiences threat modeling at Microsoft,” in *Modeling Security Workshop. Dept. of Computing, Lancaster University, UK*, 2008.

[7] ——, *Threat Modeling: Designing for Security*, 2014.

[8] D. Dhillon, “Developer-Driven Threat Modeling: Lessons Learned in the Trenches,” *IEEE Security Privacy*, vol. 9, no. 4, pp. 41–47, jul 2011.

[9] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, “A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements,” *Requirements Engineering*, 2011.

[10] D. K. Mulligan, C. Koopman, and N. Doty, “Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 374, no. 2083, p. 20160118, 2016.

[11] K. Wuyts, “Privacy Threats in Software Architectures,” Ph.D. dissertation, KU Leuven, jan 2015.

[12] N. R. Tague *et al.*, *The quality toolbox*. ASQ Quality Press Milwaukee, WI, 2005, vol. 600.

[13] J. Freund and J. Jones, *Measuring and managing information risk: a FAIR approach*. Butterworth-Heinemann, 2014.

[14] M. S. Lund, B. Solhaug, and K. Stølen, *Model-driven risk analysis: the CORAS approach*. Springer Science & Business Media, 2010.

[15] L. Sion, K. Yskout, D. Van Landuyt, and W. Joosen, “Risk-based design security analysis,” in *Proceedings of the 1st International Workshop on Security Awareness from Design to Deployment*, 2018, pp. 11–18.

[16] S. J. De and D. Le Métayer, “PRIAM: a privacy risk analysis methodology,” in *Data Privacy Management and Security Assurance*. Springer, 2016, pp. 221–229.

[17] S. J. De and D. L. Metayer, “Privacy risk analysis to enable informed privacy settings,” in *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS PW)*, April 2018, pp. 95–102.

[18] S. J. De and A. Imine, “To reveal or not to reveal: balancing user-centric social benefit and privacy in online social networks,” in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*. ACM, 2018.

[19] F. Karegar, N. Gerber, M. Volkamer, and S. Fischer-Hübner, “Helping john to make informed decisions on using social login,” in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, 2018.

[20] D. Vose, *Risk analysis: a quantitative guide*. John Wiley & Sons, 2008.

[21] T. DeMarco, *Structured Analysis and System Specification*, 1979.

[22] L. Sion, K. Yskout, D. Van Landuyt, and W. Joosen, “Solution-aware Data Flow Diagrams for Security Threat Modelling,” in *Proceedings of SAC 2018: The 6th track on Software Architecture: Theory, Technology, and Applications (SA-TTA)*, 2018.

[23] S. Hernan, S. Lambert, T. Ostwald, and A. Shostack, “Threat Modeling: Uncover Security Design Flaws Using The STRIDE Approach,” *MSDN Magazine*, vol. 6, nov 2006.

[24] L. Sion, P. Dewitte, D. Van Landuyt, K. Wuyts, I. Emanuilov, P. Valcke, and W. Joosen, “An Architectural View for Data Protection by Design,” in *2019 IEEE International Conference on Software Architecture (ICSA)*, mar 2019, p. to appear.

[25] L. Sion, D. Van Landuyt, K. Yskout, and W. Joosen, “SPARTA: Security & privacy architecture through risk-driven threat assessment,” 2018.

[26] CNIL, “Privacy Impact Assessment (PIA) 3: Knowledge Bases,” CNIL, Tech. Rep., 2018.

[27] K. Beckers, “Comparing privacy requirements engineering approaches,” *Proceedings - 2012 7th International Conference on Availability, Reliability and Security, ARES 2012*, pp. 574–581, 2012.

[28] R. R. Heckle and S. H. Holden, “Analytical tools for privacy risks: Assessing efficacy on vote verification technologies,” in *Symposium On Usable Privacy and Security*, 2006.

[29] S. Abu-Nimeh and N. R. Mead, “Privacy risk assessment in privacy requirements engineering,” *2009 2nd International Workshop on Requirements Engineering and Law, RELAW 2009*, pp. 17–18, 2009.

[30] Internal Revenue Service, “Internal Revenue Service Model Information Technology Privacy Impact Assessment,” IRS, Tech. Rep., 1996.

[31] N. R. Mead and T. Stehney, “Security Quality Requirements Engineering (SQUARE) Methodology,” *SIGSOFT Softw. Eng. Notes*, 2005.

[32] S. Spiekermann and L. F. Cranor, “Engineering privacy,” *IEEE Transactions on Software Engineering*, vol. 35, no. 1, pp. 67–82, 2009.

[33] M. Alshammari and A. Simpson, “A model-based approach to support privacy compliance,” *Information & Computer Security*, 2018.

[34] J. I. Hong, J. D. Ng, S. Lederer, and J. A. Landay, “Privacy risk models for designing privacy-sensitive ubiquitous computing systems,” *Proceedings of the 2004 conference on Designing interactive systems processes, practices, methods, and techniques - DIS ’04*, p. 91, 2004.