Selecting Privacy Enhancing Technologies for IoT-Based Services

Immanuel Kunz(✉), Christian Banse, and Philipp Stephanow

Fraunhofer AISEC, Garching, Germany
{immanuel.kunz,christian.banse,philipp.stephanow}@aisec.fraunhofer.de

Abstract. The rising number of IoT devices enables the provisioning of novel services in various domains, such as the automotive domain. This data, however, is often personal or otherwise sensitive. Providers of IoT-based services are confronted with the problem of collecting the necessary amount and quality of data, while at the same time protecting persons' privacy using privacy enhancing technologies (PETs). Selecting appropriate PETs is neither trivial, nor is it uncritical since applying an unsuitable PET can result in a violation of privacy rights, e.g. according to the GDPR. In this paper, we propose a process to select data-dependent PETs—i.e. technologies which manipulate data, e.g. by distorting values—for IoT-based services. The process takes into account two perspectives on the selection of PETs which both narrow down the number of potentially applicable PETs: First, a data-driven perspective which is based on the data's properties, e.g. its longevity and sequentiality; and second, a service-driven perspective which takes into account service requirements, e.g. the precision required to provide a particular service. We then show how the process can be applied for automotive services proposing a taxonomy for automotive data and present an exemplary application.

In this way, we aim at providing a reproducible method of selecting PETs that is more specific than existing approaches, and which can be applied both as a standalone process and complementary to existing ones.

Keywords: Privacy enhancing technologies · Automotive data · Privacy-Preserving IoT

1 Introduction

The increasing amount of IoT devices is enabling the provisioning of many novel services, like predictive maintenance and usage-based insurance. These services collect data from a multitude of IoT devices with the purpose of processing the data in a remote backend. Commercial solutions for implementing such services

N. Park et al. (Eds.): SecureComm 2020, LNICST 336, pp. 455–474, 2020.
https://doi.org/10.1007/978-3-030-63095-9_29

are already available provided, e.g., by Amazon Web Services[1] for general IoT services, and by BMW for the automotive domain[2].

Yet, IoT data often contains personal or otherwise sensitive information and it is in the interest of the user as well as the service provider to preserve users' privacy when collecting this data. The need for privacy protection results from legal requirements regarding the collection of personal data, e.g. the General Data Protection Regulation (GDPR). Two of its central principles are *data minimization* and *purpose limitation*, and it makes violations against these principles punishable by considerable fines. Also, the GDPR demands that personal data is secured appropriately and that certain rights are provided to the data subjects, e.g. to access and rectify their personal data. Therefore, minimizing processing and storage of personal data seems only reasonable when acting as an IoT service provider who wants to minimize this liability.

A multitude of general and specialized techniques to protect personal data, usually referred to as *Privacy Enhancing Technologies* (PETs), has been researched and developed in the past. Selecting suitable PETs for a given set of data, however, is not trivial. First, not all PETs are applicable to any type of data, e.g. to numeric or categorical data. Second, service providers usually require collected data to satisfy certain properties to conduct meaningful analysis, for instance, regarding its precision.

Research in the area of privacy requirements engineering focuses on risk-driven approaches [5,18] where data flows throughout a system are analyzed and privacy risks are identified which are then mitigated using appropriate PETs. Yet, there is not a clear mapping of privacy risks to mitigative PETs, so they can often only suggest large amounts of PETs that are potentially applicable, leaving the actual selection to the user. For some types of PETs it is still possible to make a more granular selection taking into account whether they can be applied in a *meaningful* manner, i.e. if a PET can practically be applied to the data type, and if the results are *usable* considering the service's purposes.

In this paper, we propose a process to select PETs for IoT-based services which aims at providing a more granular selection process. It focuses on data-dependent PETs, i.e. such technologies whose applicability depends on the characteristics of the data that is processed, and takes also into account service-specific requirements. As such, it complements existing risk-driven approaches. Its contribution is twofold:

- A general process for the selection of data-dependent PETs including data-driven and service-driven elicitation criteria, and
- a proposal for its application in the automotive domain including a taxonomy of automotive data that is mapped to applicable PETs.

The remainder of this paper is structured as follows: The following section describes the classes of PETs that we cover with our approach. Next, Sect. 3 describes the selection process including the relevant criteria we identified for a

[1] https://aws.amazon.com/iot/.
[2] https://aos.bmwgroup.com/.

data-driven and service-driven elicitation of PETs. Section 4 presents our application of the process to the automotive domain proposing a taxonomy for automotive data and presents an example. Section 5 discusses the process based on several requirements, such as reproducibility. Finally, Sect. 6 describes related work and Sect. 7 concludes the paper.

2 Background: Data-Dependent PETs

Many IoT-based services face considerable privacy challenges due to several reasons. First, IoT devices often generate large amounts of data that can be personal or otherwise sensitive, e.g. location data and usage statistics. Second, devices like the ones used in connected cars and smart homes, are often used throughout several years. And third, they are usually owned by the same person or group of persons over that time period which allows the creation of long-term tracking profiles. These conditions make it difficult to satisfy privacy goals like anonymity of users and unlinkability of individual data items.

A multitude of PETs has been proposed in the past to solve these problems. While different types of PETs exist, in this paper we focus on those whose applicability directly depends on their input data. This dependency results from the fact that these PETs manipulate data somehow, for instance by modifying, replacing or deleting values. Other types of PETs, e.g. based on encryption techniques or based on usage control, usually can be applied independently from the kind of data that is processed, and are not considered in this paper.

The following enumeration is extracted from an existing collection of PETs by Hundepool et al. [15]. They also map these PETs to two kinds of data they can process: continuous and categorical data. Continuous data is numeric data on which arithmetic operations can meaningfully be applied, while categorical data assumes values from a finite set on which arithmetic operations do not make sense. Note that categorical data can additionally be ordinal meaning that it can be sorted in a meaningful order. This enumeration and their mapping to continuous and categorical data forms the basis for our approach.

This way, we do not analyze specific PET-algorithms, but treat a PET as a function taking either continuous- or categorical-valued inputs and generating a privacy-enhancing output. Our analysis therefore does not consider any specific PET algorithm but only classes of PETs. Still, it aims at preserving their semantics without loss of generality.

We further categorize the PETs into *deterministic distortion* PETs, i.e. PETs which manipulate data using deterministic values, and *randomized distortion* PETs, i.e. PETs which manipulate data using randomized values.

Deterministic Distortion—Recoding. Recoding techniques aim at reducing the precision of data. As such, they implement a form of discretization, for example by summarizing two categorical values or by rounding numeric values. Hence, they can be applied to both categorical and continuous data [15]. Special types of recoding techniques are top and bottom recoding where the top

and/or bottom values in a set are summarized into one category. These techniques require that the values can be ranked, i.e. they are either continuous, or they are categorical and ordinal.

Deterministic Distortion—Suppression. This technique means the removal of a value or a group of values and can be applied to both value types as well. One way to apply suppression is to use *local suppression* meaning that one particular value is suppressed, e.g. with the aim to eliminate an identifying combination of values (see Fischetti and Gonzáles [9]). Another possibility is to suppress certain attributes, e.g. when they contain identifying information, or certain records, e.g. when they represent easily identifiable outliers.

Deterministic Distortion—(Micro-)Aggregation. Aggregation techniques take a set of values and replace them by a statistic. Various micro-aggregation techniques are presented in [15]. When applying micro-aggregation to continuous data, records are clustered, i.e. a number of groups with a certain minimum size are built, and the values are replaced by their respective group average. Concerning categorical data, micro-aggregation can only be applied for ordinal values since only in this case can the data be assigned to groups of similar values (see Torra [27]).

Randomized Distortion—Swapping. With this technique, values of a continuous or categorical attribute are swapped between records. Hence, they are preserved within the data set and allow certain statistical analyses but cannot be traced back to a certain data subject. This is useful when sensitive attributes exist in the data set whose distribution is interesting but their connection to other entries of its original record is not necessary (see Moore [19]).

Randomized Distortion—Noise Masking. Noise masking aims at distorting data by adding a randomized value to, or multiplying it with, the original value. Different techniques exist that preserve certain statistical properties within a data set when applying noise masking, see for example Domingo-Ferrer et al. [7]. Practically, it can only be applied to continuous data since it requires the possibility to perform arithmetic operations.

Randomized Distortion—Post-randomization Method (PRAM). This method, proposed by Gouweleeuw et al. [4], can only be applied to categorical values. It changes each value independently with a certain probability. This way, its application results in a diminished certainty of the linkability between two attributes. Still, meaningful statistics can be computed afterwards—depending on the probability that has been applied in the method and the size of the data set.

Randomized Distortion—Synthetic Data. Three approaches can be distinguished to introduce synthetic data into a data set which are generally applicable to both continuous and categorical data. In the first approach, the data set is fully replaced by synthetic data, yet preserving certain properties, like the distribution of an attribute (see e.g. Rubin [24]). The second approach partially replaces original data with synthetic data, for example by only replacing sensitive values (see Little [17]). The third approach, combines original and synthetic data into a *hybrid* data set containing either more original data or more synthetic data (first proposed by Dandekar et al. [3]).

In the next section, a process for the selection of these PETs is proposed, considering how data types and service requirements influence their applicability.

3 A Process for the Selection of PETs in IoT-Based Services

The goal of the process proposed in this Section is to present a systematic way of eliciting a usable set of PETs for a given set of data. More concretely, we aim at fulfilling three requirements. First, the set of PETs it suggests for a given set of data and service requirements should be reproducible. Second, it should be applicable as a standalone process as well as be applicable to the results of other approaches, such as LINDDUN (see Sect. 6). Third, it should facilitate data minimization: it should guide users of the process towards selecting a set of PETs such that their application results in the minimal quantity and quality of data that is still sufficient to fulfill the service's purposes. We discuss the fulfillment of these requirements in Sect. 5.

The typical environment we assume for this process to be applied in is an IoT-based service, i.e. a service which continuously collects a pre-defined set of data from several data-generating devices. Such a service may, e.g., offer a predictive maintenance service using data generated by industrial IoT devices. Note that some of the criteria we propose here deterministically suggest or oppose a certain PET, while others only strengthen or weaken the applicability of certain PETs to some extent.

Figure 1 shows our selection process consisting of the following four steps:

1. *Service description:* The service is described including its actors and data required to provide the service.
2. *Data-driven PET-elicitation:* The applicability of the PETs presented in the last Section is assessed based on several data-dependent criteria. This way, *practically applicable* PETs are elicited. In Sect. 4.1, we will show how this can be done proposing a taxonomy for automotive data.
3. *Service-driven PET-elicitation:* The set of PET candidates is further refined based on the requirements of the automotive service towards the data. This step considers the required data precision as well as other requirements resulting in a *usable* set of PETs.
4. *PET-selection:* From the resulting set of usable PETs, a combination is selected.

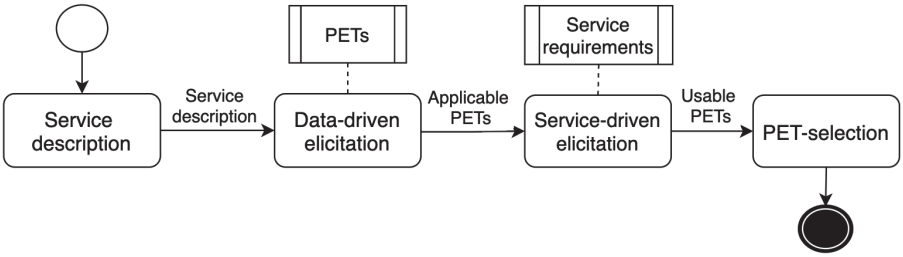

Fig. 1. The PET-selection process. It takes two inputs—a list of PET candidates and a list of service requirements—and provides a set of usable PETs.

3.1 Step 1: Service Description

In this first step, the service is described, including its participants, the required data and its processing purposes. The participants include the service provider, the users and possibly subcontractors. This step captures whose personal data is processed and by whom. The data that is required to run the service represents an input to the data-driven elicitation (Step 2). The description of its purposes is needed as an input to the service-driven elicitation (Step 3) since they reveal the requirements that the service has towards the data. Finally, out of the required data, it is identified which data needs to be protected, e.g. anonymized or pseudonymized.

To decide this, it is helpful to distinguish between the categories *identifying* and *quasi-identifying* information as well as *confidential* and *non-confidential* information, as proposed by Hundepool et al. Note that these categories are not necessarily disjoint.

- **Identifiers:** This category includes directly identifying information like full name, exact address etc.
- **Quasi-identifiers:** quasi-identifiers are sets of variables that can identify a person if combined, possibly with external information (see Dalenius [2] and Samarati [25]). The problem that quasi-identifiers pose is that every variable is potentially part of a quasi-identifier and consequently they cannot be excluded from any dataset.
- **Confidential variables:** This is sensitive, but usually not identifying, information for instance someone's religion or health conditions.
- **Non-confidential variables:** This kind of information contains non-sensi tive information, for example the country of residence.

It is always use case-dependent whether certain data needs to be protected to preserve users' privacy. Still, the categorization above suggests that identifiers should be removed and confidential, i.e. sensitive, information should be unlinkable to a person. A more complex and use case-dependent task is to specify the quasi-identifiers and to decide which of their variables should also be removed or distorted.

3.2 Step 2: Data-Driven Elicitation

Step 2 of the process elicits PETs based on the data characteristics and proposes several criteria to that end. Every criterion either specifies that a certain class of PETs is applicable, has limited applicability to some extent, or is not applicable. Their relevance is evaluated for all data that was defined in step 1. This way, the list of possibly applicably PETs is narrowed down in a bottom-up elicitation. We also give some examples from the automotive domain to illustrate the criteria.

Continuous and Categorical Data. One simple criterion in the elicitation of data-dependent PETs is the given data's value type. In Sect. 2, we have described that PRAM only accepts categorical, and noise masking only accepts continuous data. Other PETs, in contrast, can be applied to both (i.e. recoding, suppression, aggregation, swapping, synthetic data).

Set Size. Set size refers to the number of values the data can assume. With a decreasing set size, the applicability of PETs in general decreases as well. The reason is that a small set size does not leave much possibility for distorting a value. For example, noise masking can only be applied using a small randomized value in the case of a small set size. So, either only a small distortion is applied—leaving a greater possibility to infer sensitive information—or a greater distortion is applied—resulting in a strongly decreased information content. For instance, appropriately distorting data that expresses a vehicle's engaged gear is more difficult than distorting data expressing the vehicle's horsepower. A further example is data that only expresses a binary value, e.g. a crash sensor which only measures if a crash has occurred or not.

Ordinal and Nominal Data. As seen in Sect. 2, among categorical data, ordinal and nominal data can be distinguished. *Nominal* categorical values cannot be compared to each other apart from equality. *Ordinal* categorical values can be ranked such that a minimum and maximum can be found and different values can be compared. Consider, for instance, a car which generates data about its driving mode, e.g. *eco* and *sport*, and about the seat positions. While the driving modes can only be compared for equality, different seat positions can be ranked and their differences can be quantified.

If the data is categorical, but not ordinal, aggregation PETs cannot be applied [15].

Data Longevity. This criterion expresses how often a specific data value changes. The longer a value lives, the more limited is the applicability of randomized distortion PETs since the resulting data may allow to infer properties of the raw data in the long-term. Applying, for instance, noise masking to a sensitive attribute with a constant distribution, may allow to infer the raw value in the long term. Also the appropriate parametrization of randomized distortion PETs becomes more difficult for long-lived data.

Value Sequences. Some kinds of data can only assume values in a specific sequence, like the state of a car's assistance system. In this case, the applicability of randomized distortion PETs, such as swapping, is limited since the resulting values can exhibit a non-valid sequence. This limitation also concerns suppression since suppressed values may be reconstructed using the expected value sequence.

Metadata and Identifiers. Data may transport identifiers and other metadata that is required for a correct transmission and linkability of data items. This data therefore does not transmit information that is directly required for the service's purposes.

When anonymizing identifiers, randomized distortion PETs are technically applicable, but not in a meaningful way. A MAC address, for example, contains an identifier for the device's manufacturer. Applying a randomized distortion PET like noise masking to this address would result in an anonymization of the manufacturer but would also result in falsely identifying a different manufacturer. The identity information, however, is the only information contained in this data and would effectively be erased by a randomized distortion. Other techniques, like recoding, can still be applied in a meaningful way, e.g. by masking a number of digits of an identifier.

In Sect. 4.1 we show how a taxonomy of the data can be built based on these criteria. This way, a direct mapping of data categories to applicable PETs can be established.

3.3 Step 3: Service-Driven Elicitation

The data-driven elicitation in step 2 results in a set of meaningfully applicable PETs. Yet, this set does not allow any statement about the usefulness of the candidate PETs considering the purposes of the service. For instance, an applicable PET may be suppression which, however, may not only suppress sensitive data, but also important information for the service's processing purposes. Therefore, also a service-driven elicitation—based on the service's requirements towards the data utility—can be conducted to further specify which PETs are useful. This step can be seen as a top-down mapping since it is based on the service's requirements.

Three criteria for the service-driven elicitation are examined in the following.

Value Precision. One service-dependent criterion is the required precision of collected values since it determines the scope in which PETs may be used. Value precision can be changed, for example, by rounding values to a certain decimal place or by masking the last digits of an identifier. The precision of categorical data can only be changed by generalizing the values, i.e. by finding new, more general categories that comprise two or more original ones. Also, PETs that distort values using random values, like noise masking and PRAM, reduce the values' precision. In contrast, applying synthetic data, suppression or aggregation does not preserve the original values at all.

In case a service requires high precision of collected data, this criterion therefore limits the applicability of all PETs seen in Sect. 2 except swapping since swapping preserves all exact values and only swaps them between records. In contrast, a scenario in which a service does not require exact values but rather requires statistics about the data, endorses the application of randomized distortion PETs. Synthetic data, for instance, can be designed to manipulate data while preserving certain statistical properties.

Data Freshness. Data freshness expresses how long ago data has been generated. High requirements towards data freshness, e.g. requiring near real-time data, can limit the applicability of aggregation PETs since their application relies on collecting data from one or multiple data sources introducing a certain latency into the processing.

Attribute Dependency. A further criterion that influences the usefulness of PETs concerns the dependency between attributes. If a service provider needs to collect a certain combination of attributes, it may not be possible to obtain meaningful results when processing them independently from each other using different PETs. Consider, for example, a connected car generating distance and speed data. If the service's purpose is to infer a correlation between the two kinds of data, the data's usefulness may be reduced if the distance data would be perturbed using noise masking but the speed data would be recoded using generalization.

Such a set of data can also originate from multiple vehicles. Consider, for example, a fleet of vehicles whose average fuel consumption shall be determined. In this case, too, the statistic cannot be reliably computed anymore if the vehicles' consumption values have been processed using different PETs. In some case-specific, this problem can be avoided altogether if data is collected on a higher level of abstraction, i.e. the inference computations are done locally, inside the vehicle.

A special case of attribute dependency concerns the identifiability of users, i.e. the requirement to associate certain data items with a particular person or a group. If only the data values are needed, e.g. for statistical analysis, identification of a user is not necessary which suggests the suppression or pseudonymization of the identifier. Hence, if data needs to be traceable to a certain individual, the applicability of suppression and pseudonymization is limited. It may, however, be the case that identifying attributes need to be collected by a service together with other attributes. If there is no dependency between the two, swapping can be a useful PET since it can be used to break the link between identifying values and other, otherwise sensitive, values and still preserve them within the data set.

3.4 Step 4: PET-Selection

After the two elicitation steps, a set of usable classes of PETs results from which appropriate ones can be selected. This last selection step poses the question of

how to select PETs from a set of usable ones. In the following, we briefly describe criteria that may be relevant for this question. As it is highly application- and use case-specific, however, we consider a further investigation out of scope for this paper.

To select a final set of PETs, a service provider may consider the implementation cost and complexity of the PETs. If, for instance, suppression results as an option, it may be considered first, since it is relatively easy to implement and reliably protects the respective data. Also, the service provider may want the collected data to satisfy certain privacy metrics, like k-anonymity, that result from legal requirements or internal compliance requirements. These metrics may only be satisfiable by certain PETs.

Furthermore, it is possible that step 3 results in an empty set of useful PETs, i.e. there is no possibility to anonymize the given set of data such that it still meets the service's requirements. In this case, the consequence may be that the service provider needs to obtain the users' consent to the collection and processing of the raw personal data. Otherwise the service may simply not be implementable.

4 Selecting PETs for Automotive Services

In this section, we show how the process proposed above can be applied in the context of IoT-based automotive services. As seen in Sect. 3.2, the meaningful applicability of PETs depends on the data's characteristics, such as *set size*, *longevity*, and *value sequences*. In what follows, we propose a taxonomy for automotive data that is created based on the criteria we proposed, and we map the resulting categories to applicable PETs. This way, we aim at providing a domain-specific data-driven elicitation that can generally be applied to IoT-based services in the automotive domain. Finally, we show an exemplary application of the process which uses this taxonomy for the data-driven elicitation step.

4.1 A Taxonomy for Automotive Data

In the following, we present our taxonomy for automotive data and discuss to what extent the PETs described above are applicable for their processing. It comprises the categories of *Communication Metadata*, *Vehicle Attributes*, *Stream Data*, *State Variables*, *Event Data* and *Complex Data*. The mapping is summarized in Table 1 (at the end of this Subsection).

Communication Metadata. Communication metadata includes data that is needed for the transmission of messages, for instance: SIM card ID, IP address, MAC address, Bluetooth address and Vehicle Identification Number (VIN). While this data is numeric, it is actually categorical data since these identifiers assume values from a finite set and follow a fixed format that usually does not allow to perform arithmetic operations on them. This consideration influences the applicability of PETs as explained in Sect. 3.3.

Vehicle Attributes. Vehicle attributes include rather long-lived data that describes the vehicle, for example: Transmission information (e.g. number of gears), fuel type (e.g. gasoline, diesel, electric, hybrid), engine information (e.g. number of cylinders, battery type, construction type), vehicle dimensions (e.g. width, length, height), tire and wheel dimensions (e.g. diameter and width) and number of doors and seats.

Vehicle attributes are similar to communication metadata since they also assume values from a finite set and meaningful arithmetic operations cannot be performed on them, i.e. they pertain to the categorical value type. Yet, they also differ from communication data. First, they are often not numeric, and second, the application of randomized distortion PETs can yield meaningful results. In comparison to a MAC address, for example, the number of seats in a vehicle can be perturbed using a PET like PRAM. Still, the applicability of these PETs is limited due to the data's longevity and limited set size.

Vehicle attributes can, for instance, be used to fingerprint a vehicle for tracking purposes since they expose long-lived information.

Stream Data. This category contains data that is generated continuously by sensors, for example: Location, temperature (e.g. cabin, outside, motor), pressure (e.g. tire pressure), speed, distance (e.g. LiDAR, RADAR, ultrasound), battery status (e.g. charging cycles), steering angle and yaw rate.

This data is continuous, therefore all PETs except PRAM can be used to process it. A special kind of stream data is location data. It has a high privacy-relevance since it often reveals identifying and sensitive information about users. A number of specialized PETs exist for anonymizing location data, which can also be classified according to the PET classes included in our approach. For example, existing works have proposed recoding techniques [8,12,13], synthetic data [16] and aggregation PETs [6] for location data.

As an example, consider a predictive maintenance service which analyzes stream data to recognize anomalies in the generated data. It can use this data to predict when certain sensors or parts need to be replaced or maintained, for instance suggesting an oil or tire change.

State Variables. State variables express the state of a system, for example of an assistance system or the state of cabin settings. Examples are: Assistance systems (e.g. system state of cruise control, adaptive cruise control, lane-keeping assistance), cabin settings (e.g. positions of seats, mirrors, steering wheel), infotainment usage data (e.g. used media types and infotainment services), driving mode (e.g. eco, sport) and light settings (e.g. daytime running light, long distance light).

State variable data is categorical. Similar to communication metadata and vehicle attributes, it can be processed using recoding while the application of other techniques is either limited or not given at all. It still differs from other data categories since it is rather short-lived. Consider, for example, the light setting: It is often changed by a rotary switch where the setting can only be

changed to one of its neighboring settings. This way, a sequence for the light setting is implicitly established.

For example, this data may be used in a car-sharing service which automatically prepares the cabin settings, for instance seat and mirror positions, for the driver's preferences when she rents a vehicle.

Event Data. This data is sensor data that is generated irregularly indicating that a certain event has taken place, for example: Automatic safety belt tightening, diagnostic trouble codes, ESP intervention, lane departure warning, fatigue warning and emergency assist.

Event data can be treated as categorical data. Yet, it can only assume one value rendering most PETs non-applicable. Event data can neither be recoded—e.g. no generalized categories can be found—nor can the data be perturbed by noise masking, swapping or PRAM since the one possible value cannot be changed or replaced. Neither can aggregation be applied since no statistics of the data's values can be built. The only class of PETs which can be applicable without limits is synthetic data. Using synthetic data, events can be generated which preserve the original events' statistical properties, e.g. average time elapsed between occurrences, while at the same time masking their real occurrences.

For instance, crash detection data can be used by an application that triggers an automatic emergency call. The automatic safety belt tightenings can further be used to infer information about the driving behavior of the driver.

Complex Data. This category summarizes data which cannot be assigned to any of the above categories and is thus not included in Table 1. Consider, for example, combined data structures collected from various sources, e.g. represented as n-tuples, potentially requiring the application of a different PET to each n-th element. Further examples of complex data are image and sound data used for voice assistants and for inward-facing and outward-facing cameras. Their privacy-enhancement cannot be ensured by reducing precision since this does not necessarily ensure that persons recorded by a camera or a microphone are properly de-identified. Hence, specialized PETs are required for their anonymization, such as proposed in [1] and [23]. Other examples for complex data are: Installed applications, account information (e.g. user names, passwords, payment information), voice and video recordings, contact list, call history and data induced by the driver (e.g. query data for location-based services).

Table 1 summarizes the results of the mapping showing which classes of PETs (recoding, suppression, aggregation, swapping, noise masking, PRAM and synthetic data) are applicable to which automotive data categories (communication metadata, vehicle attributes, stream data, state variables, event data, location data). *X* represents *applicable*, *(X)* represents *limited applicability* and a gray box represents *not applicable*. Note that the general mapping between PET classes and value types is taken from [15] (as described in Sect. 2), and the granular

mapping of automotive data categories to PET classes is our contribution (as described in Sect. 4.1).

Table 1. Applicability of PET classes to automotive data categories. The mapping of the general applicability of value types (categorical, continuous) to PET classes is taken from Hundepool et al. [15]; our contribution extends this mapping to the automotive data categories and proposes a granular discussion of their grade of applicability.

Automotive data	Value type	PET classes						
		Det. Dist.			Rand. Dist.			
		Recod.	Suppr.	Aggr.	Swapping	Noise M.	PRAM	Synthet.
Metadata	Cat	X						
Vehicle Attr.	Cat	X	(X)	(X)	(X)	(X)	(X)	(X)
Stream data	Cont	X	X	X	X	X		X
State vars	Cat	X			(X)		(X)	(X)
Event data	Cat							X
Location	Cont	X	(X)	X	(X)	(X)		(X)

4.2 Example Process Application

In the following, we describe a fictitious example service to show how the process proposed above can be applied in a real-world scenario.

In our example, we consider an insurance company as the service provider who wants to collect data about the driving behavior of customers to offer a risk-based payment model. To calculate the risk-based price for the user, the service provider wants to collect a set of data from the customers' vehicles.

1) Service Description. In this first step, the required data as well as the actors in the service are described. The actors in this example include the insurance company as the service provider as well as the service users who are assumed to be the owners of the vehicles from which data is collected. To determine the driving behavior of the driver, the service provider wants to collect data about the vehicle's speed, the operating times (e.g. driving times in the morning or the evening) as well as data about the interventions of the Electronic Stability Control (ESC) assistance system. Also, the insurance company wants to know the vehicle's horsepower (HP) as vehicles with a higher number of HP are associated with a higher risk of causing accidents.

While this data is not directly identifying, it can be classified as sensitive as it reveals information about the concerned person's driving behavior.

2) Data-Driven Elicitation. The data-driven elicitation represents the second step of our selection process, identifying applicable PETs for the required

data. First, speed data can be categorized as stream data since it is generated continuously by the respective sensors. Hence, all considered classes of PETs except PRAM are applicable. The vehicle's operating times can be seen as a status variable that assumes values from the sequence of possible time values. It can therefore best be processed using recoding techniques. Third, the ESC interventions represent event data since they represent irregularly generated events. They can be processed using the PETs synthetic data or, applied to its timely occurrences, aggregation. Fourth, the vehicle's HP pertains to the category of vehicle attributes. Due to its longevity, only recoding PETs are applicable here.

3) Service-Driven Elicitation. In the third process step, the service requirements towards the data are considered to further specify the set of applicable PETs. The first criterion concerns the value precision the service provider requires to be able to provision the service. Regarding the speed data, the service provider does not need to know exact speed values but rather speed intervals in which the vehicle is operated. We assume that the required precision for the speed data equals to intervals of 20 km/h. Similarly, the HP is only needed in intervals of 20 HP to evaluate the risk that the car model induces. The operating times of the vehicle are required only to infer the time of day, i.e. day- or night-time. Lastly, the ESC interventions are required as the exact number of occurrences in the preceding month.

Data freshness is the second criterion of the service-driven elicitation considering how fast the service provider needs the data after it has been generated. In our example, the service provider does not have any freshness requirements since the risk assessment can also be conducted on old data. Also, the third criterion, value dependency, does not further narrow down the set of applicable PETs.

4) PET-Selection. From the set of meaningfully applicable PETs that result from the service-driven elicitation, in step 4 it can be selected as follows.

All speed data can be rounded down to the nearest multiple of 20 km/h to meet the service provider's minimal requirement towards the precision of the data, i.e. a recoding technique can be applied. Applying randomized distortion PETs is possible as well, e.g. by applying additive noise masking with a randomized value between −10 and +10.

The only meaningfully applicable class of PETs for the horsepower value is recoding PETs. Since the required precision equals intervals of 20 km/h, a rounding PETs with this granularity can be applied.

Concerning the status variable of operating times, the service provider is especially interested in the information at which times of the day the vehicle is operated. Therefore, a generalization is suggested to the values *day time* and *night time*. Other PETs are not usable here. For example, applying a randomization with PRAM would result in an easily reversible sequence where the variable can be matched with the actual time of day when the data was received.

The event data showing ESC interventions can be processed using an aggregation of the events, e.g. building a statistic of the average occurrences in a

certain time frame. Alternatively, synthetic data could be generated to mimic the real occurrences of the events preserving relevant statistics.

In this Section we have shown a case study of the application of our process building a taxonomy of automotive data and using it in a concrete example. While this case study is specific to the automotive domain, we expect our process to be equally well applicable to other IoT-related domains. Considering the proposed taxonomy, it can be transferred, for instance, to smart home appliances as well as these devices also generate stream data (e.g. a fridge's temperature), event data (e.g. a finished washing machine), state variable data (e.g. the lighting state) and metadata and they may as well possess long-lived attributes.

5 Discussion

In what follows, we discuss to what extent the proposed process can satisfy the three requirements we defined in Sect. 3, namely its *reproducibility*, its *applicability* as a standalone process as well as complementary to other approaches, and its ability to facilitate *data minimization*.

5.1 Reproducibility

Since the reproducibility of the process depends on the reproducibility of the two elicitation steps 2 and 3, we discuss this requirement for these two steps. Meanwhile, step 1 only provides the service description and the selection in step 4 is not examined in detail in this paper.

Step 2 provides data-driven criteria which indicate to what extent certain PETs are applicable. On the one hand, these criteria are not completely unambiguous since in some cases, they only indicate a "limited" applicability which does not allow a deterministic decision. On the other hand, we have described the application-specific criteria which determine the applicability if it is limited, e.g. regarding the PETs' parametrization. For example, if the data values at hand follow a specific sequence, swapping can only be applied if the resulting states still contain sufficient randomness. The results of step 2 are therefore reproducible to the extent where a limited applicability might be evaluated differently by different users.

Step 3 includes service-driven criteria, like precision and freshness of the data. These criteria are more use case-dependent. For instance, a high value precision limits the applicability of most PETs. It is, however, use case-dependent how the required value precision influences the usability of a PET. The service-driven criteria therefore rather provide guidelines for service providers that have to be evaluated using the specific purpose and data at hand. As such, their reproducibility depends on how a user of the process evaluates the degree of precision and freshness the service requires.

In summary, the process steps are reproducible to the point where application-specific criteria need to be considered. Given the same set of input data, step 2 results in a reproducible set of applicable and non-applicable PETs,

while a result of limited applicability may be interpreted differently by different users. Step 3 includes two criteria that also may be evaluated differently by different users since their evaluation depends on the specifics of the service.

To further enhance the reproducibility of the process, it may be beneficial to make the results of the single process steps more comparable. To that end, appropriate metrics could be defined for the PETs to make the selection and parametrization of a certain PET comparable to other use cases. For instance, identifying a metric that measures the randomness in a state variable sequence would make the evaluation of the applicability of noise-masking PETs more comparable and its application more reproducible given a certain set of requirements.

5.2 Standalone and Complementary Applicability

The proposed process is applicable as a standalone process given a service description which includes the required data and its purposes. Applying the process standalone can be especially useful if it is integrated into the requirements engineering process during a system design. The process can be applied early on in the system design to assess whether useful PETs can be found for the required data and to assess to what extent the data collection can be minimized. Since its applicability does not require a complete system design or data flow analysis, it may be applied many times for such an assessment during a system design.

The process can, however, also be applied complementary to other existing approaches (see Sect. 6). As such it may be utilized to narrow down a list of potentially applicable PETs that have been elicited from a risk analysis. In the following, we first discuss how our process compares to existing approaches and then explain how a complementary application can be conducted.

One difference between risk-driven approaches and our process is that risk-based approaches suggest a list of PETs based on privacy threats. As such, they do not elicit *practically applicable* PETs—which is ensured in our approach by the data-driven elicitation—but suggest PETs that generally may apply to the identified threats. Furthermore, they leave the service-driven elicitation to the user applying the approach.

The LINDDUN methodology, proposed by Deng et al. [5], defines six steps to systematically approach the elicitation of privacy requirements. LINDDUN is an acronym for the considered privacy threats which are linkability, identifiability, non-repudiation, detectability, disclosure of information, unawareness of users as well as policy and consent non-compliance. Following this methodology, first a Data Flow Diagram (DFD) is created and privacy threats are identified. Then, threats are prioritized, mitigation strategies are established and in the last step PETs are selected. Note that we do not discuss policy and consent compliance here since they do not influence the selection of data-dependent PETs.

Consider again the example application described above, where data about operating times, ESP interventions and the vehicle's HP is collected. When identifying threats for these data flows, none of the threats described above can be excluded. Two trips may, for instance, be linkable if the driver has consistently

high numbers of ESP interventions due to her driving style. Apart from this, a vehicle may always be identifiable by the metadata that is included in a message. Consequently, a risk-driven approach generates a considerable number of threats for the data flows and will suggest a large number of potentially applicable PETs. Here, the proposed process can narrow down the list of PETs to the ones which are practical and usable.

A complementary application can be conducted based on the data flows that are identified in the risk-based approach. LINDDUN, for example, includes the creation of a data flow diagram. This diagram can also be used as the service description which corresponds to the result of step 1 of our process. Furthermore the process takes two inputs as illustrated in Fig. 1: PET candidates and the service requirements. As PET candidates, the results from the risk-based analysis are used. Note, however, that only data-dependent PETs can be taken into account here. The service requirements are needed as input to step 3. We assume these to be known since a service description including necessary data flows is already required for the risk-based analysis.

5.3 Data Minimization

The proposed process goes beyond the decision of collecting certain data items or not, but aims at selecting data-dependent PETs that can limit the quantity of data—e.g. suggesting suppression—as well as the quality of data—e.g. suggesting randomized distortion PETs—to what is necessary for the service's purposes. It does so by applying data-driven as well as service-driven criteria to the PET elicitation. The process is, however, limited to the PETs it is scoped to (see Sect. 2). Also, the appropriate minimization of data collection depends on the parametrization of the chosen PETs. Still, we would argue that it provides many important criteria towards data minimization as it is required by various data protection regulations, such as the GDPR which requires the collection of personal data to be "limited to what is necessary in relation to the purposes for which they are processed" (Article 5).

6 Related Work

6.1 Privacy Requirements Engineering

Existing work on privacy requirements engineering focuses on risk-based approaches. In Sect. 5, we have already discussed LINDDUN.

A further approach by Luna et al. [18] is called *Quantitative Threat Modeling Methodologies*. Here, similar to LINDDUN, first a DFD is created and it is mapped to privacy and security threats. Next, possible misuse case scenarios are examined, i.e. use cases from the perspective of an attacker. In a fourth step, the threats are quantified using attack trees and finally, mitigation strategies are chosen. The methodologies of Deng et al. and Luna et al. are both based on the threat modeling approach developed by Microsoft called STRIDE [22].

Oliver [21] proposes a privacy requirements framework. It describes ontological structures for the description of data and system properties, like information types (e.g. location or identifier) and usage types (e.g. advertising or profiling). From these, privacy requirements can be generated.

Spiekermann and Cranor [26] identify three types of system activities from which privacy requirements can be deduced: data transfer, data storage and data processing. Additionally, they include user expectations and behavior, as well as the threat model—which can vary between users—in their requirements analysis. Furthermore, they differentiate between two basic approaches in engineering privacy requirements, namely *privacy by architecture* and *privacy by policy*.

In comparison to these mostly risk-driven approaches, our selection process is data- and service-driven. As such, it can complement risk-driven approaches since it does not target specific privacy threats but aims at providing a more granular selection. Following the classification by Notario et al. [20], our process is a *goal-oriented* approach rather than a *risk-based* one.

6.2 Categorizing Automotive Data

A further contribution of this paper is a taxonomy of automotive data. Existing proposals focus on building taxonomies according to the functional parts of a car rather than for the applicability of PETs.

One similar approach to categorizing automotive data was proposed by Hornung [14] who identifies four categories: local data (e.g. Bluetooth ID, MAC addresses), environment data (e.g. concerning the infrastructure and the weather), third-party data (e.g. about installed applications) and personally-identifiable information (e.g. biometric data and voice recordings). He also suggests other categories based on the complexity of data which are: vehicle attributes (e.g. model), communication data (e.g. Car-2-Car communication), sensor data (e.g. temperature), processed data for the driver (e.g. navigation data) and infotainment data.

The GENIVI Alliance[3] suggests a list of automotive data specifying the value type and unit in a tree structure[4]. Version 2.0 of this specification defines as top-level branches vehicle parts, like drive train and chassis, but also concrete signals like drive time and ambient air temperature which do not directly pertain to any of these vehicle parts. Hence, these categories are especially useful in the design of on-board automotive software systems where it is important to have a well-arranged structure that provides access to the required data.

Concerning concrete sensor data, Fleming [10,11] lists automotive sensors and categorizes them based on their functionality or the respective vehicle part power train, chassis or body.

In comparison to these approaches, our taxonomy proposed in Sect. 4.1 is based on data-driven criteria for the applicability of PETs.

[3] https://www.genivi.org.
[4] https://github.com/GENIVI/vehicle_signal_specification.

7 Conclusion and Future Work

In this paper, we have proposed a process for the selection of PETs, focusing on two subproblems of the PET selection problem. First, we have focused on specific classes of PETs, namely PETs which manipulate data. Second, we have focused on IoT-based services establishing data-driven and service-driven criteria for the elicitation of PETs. We have shown that within this scope the applicability of PETs can be elicited more granularly than in previously proposed approaches. We have then shown how the process may be applied specifically for the automotive domain proposing a novel taxonomy for automotive data. We have furthermore discussed to what extent the proposed process satisfies reproducibility, standalone and complementary applicability, as well as data minimization.

Future work includes the development of criteria for the elicitation of other types of PETs, e.g. PETs based on encryption-techniques or based on usage-control, and further criteria for the selection of PETs in step 4 of the process. Also, we will examine the application of our selection process to other domains, such as smart homes and medical devices. We further plan to extend the process integrating the quantification of privacy using metrics like k-anonymity (as discussed in Sect. 5), and improve its usability by transforming it to an iterative approach.

Acknowledgment. This work was partly funded by the Bavarian Ministry of Economic Affairs and Media, Energy and Technology, within the project Bayern-Cloud.

References

1. Birnstill, P., Ren, D., Beyerer, J.: A user study on anonymization techniques for smart video surveillance. In: 12th International Conference on Advanced Video and Signal Based Surveillance (AVSS), pp. 1–6. IEEE (2015)
2. Dalenius, T.: Finding a needle in a haystack or identifying anonymous census records. J. Off. Stat. **2**(3), 329 (1986)
3. Dandekar, R.A., Domingo-Ferrer, J., Sebé, F.: LHS-based hybrid microdata vs rank swapping and microaggregation for numeric microdata protection. In: Domingo-Ferrer, J. (ed.) Inference Control in Statistical Databases. LNCS, vol. 2316, pp. 153–162. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-47804-3_12
4. De Wolf, P.P., Gouweleeuw, J., Kooiman, P., Willenborg, L., et al.: Reflections on PRAM. In: Statistical Data Protection, pp. 337–349. Citeseer (1998)
5. Deng, M., Wuyts, K., Scandariato, R., Preneel, B., Joosen, W.: A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. Requir. Eng. **16**(1), 3–32 (2011). https://doi.org/10.1007/s00766-010-0115-7
6. Domingo-Ferrer, J.: Microaggregation for database and location privacy. In: Etzion, O., Kuflik, T., Motro, A. (eds.) NGITS 2006. LNCS, vol. 4032, pp. 106–116. Springer, Heidelberg (2006). https://doi.org/10.1007/11780991_10
7. Domingo-Ferrer, J., Sebé, F., Castellà-Roca, J.: On the security of noise addition for privacy in statistical databases. In: Domingo-Ferrer, J., Torra, V. (eds.) PSD 2004. LNCS, vol. 3050, pp. 149–161. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-25955-8_12

8. Duckham, M., Kulik, L.: A formal model of obfuscation and negotiation for location privacy. In: Gellersen, H.-W., Want, R., Schmidt, A. (eds.) Pervasive 2005. LNCS, vol. 3468, pp. 152–170. Springer, Heidelberg (2005). https://doi.org/10.1007/11428572_10
9. Fischetti, M., González, J.J.S.: Models and algorithms for optimizing cell suppression in tabular data with linear constraints. J. Am. Stat. Assoc. **95**(451), 916–928 (2000)
10. Fleming, W.J.: Overview of automotive sensors. Sens. J. **1**(4), 296–308 (2001)
11. Fleming, W.J.: New automotive sensors—A review. Sens. J. **8**(11), 1900–1921 (2008)
12. Hoh, B., Gruteser, M., Xiong, H., Alrabady, A.: Enhancing security and privacy in traffic-monitoring systems. Pervasive Comput. **5**(4), 38–46 (2006)
13. Hoh, B., Gruteser, M., Xiong, H., Alrabady, A.: Preserving privacy in GPS traces via uncertainty-aware path cloaking. In: Proceedings of the 14th ACM Conference on Computer and Communications Security, pp. 161–171. ACM (2007)
14. Hornung, G.: Verfügungsrechte an fahrzeugbezogenen Daten [Rights of disposition for vehicle-related data]. Datenschutz und Datensicherheit-DuD **39**(6), 359–366 (2015)
15. Hundepool, A., et al.: Statistical Disclosure Control. Wiley, Hoboken (2012)
16. Kido, H., Yanagisawa, Y., Satoh, T.: An anonymous communication technique using dummies for location-based services. In: Proceedings of the 2005 International Conference on Pervasive Services, ICPS 2005, pp. 88–97. IEEE (2005)
17. Little, R.J.: Statistical analysis of masked data. J. Off. Stat. **9**(2), 407 (1993)
18. Luna, J., Suri, N., Krontiris, I.: Privacy-by-design based on quantitative threat modeling. In: 7th International Conference on Risks and Security of Internet and Systems (CRiSIS), pp. 1–8. IEEE (2012)
19. Moore, R.: Controlled data-swapping techniques for masking public use microdata sets. US Census Bureau [Custodian] (1996)
20. Notario, N., et al.: PRIPARE: integrating privacy best practices into a privacy engineering methodology. In: Security and Privacy Workshops, pp. 151–158. IEEE (2015)
21. Oliver, I.: Experiences in the development and usage of a privacy requirements framework. In: 24th International Requirements Engineering Conference (RE), pp. 293–302. IEEE (2016)
22. Potter, B.: Microsoft SDL threat modelling tool. Netw. Secur. **1**, 15–18 (2009)
23. Qian, J., et al.: VoiceMask: anonymize and sanitize voice input on mobile devices. arXiv preprint arXiv:1711.11460 (2017)
24. Rubin, D.B.: Statistical disclosure limitation. J. Off. Stat. **9**(2), 461–468 (1993)
25. Samarati, P.: Protecting respondents identities in microdata release. Trans. Knowl. Data Eng. **13**(6), 1010–1027 (2001)
26. Spiekermann, S., Cranor, L.F.: Engineering privacy. Trans. Softw. Eng. **35**(1), 67–82 (2009)
27. Torra, V.: Microaggregation for categorical variables: a median based approach. In: Domingo-Ferrer, J., Torra, V. (eds.) PSD 2004. LNCS, vol. 3050, pp. 162–174. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-25955-8_13