



A systematic overview on methods to protect sensitive data provided for various analyses

Matthias Templ¹ · Murat Sariyar²

Published online: 18 August 2022
© The Author(s) 2022

Abstract

In view of the various methodological developments regarding the protection of sensitive data, especially with respect to privacy-preserving computation and federated learning, a conceptual categorization and comparison between various methods stemming from different fields is often desired. More concretely, it is important to provide guidance for the practice, which lacks an overview over suitable approaches for certain scenarios, whether it is differential privacy for interactive queries, k -anonymity methods and synthetic data generation for data publishing, or secure federated analysis for multiparty computation without sharing the data itself. Here, we provide an overview based on central criteria describing a context for privacy-preserving data handling, which allows informed decisions in view of the many alternatives. Besides guiding the practice, this categorization of concepts and methods is destined as a step towards a comprehensive ontology for anonymization. We emphasize throughout the paper that there is no panacea and that context matters.

Keywords Anonymization · Privacy-preserving computation · Federated learning · Synthetic data

1 Introduction

The handling of personal and sensitive data is ubiquitous, both in research and in industry. Even though there are various methodological developments regarding the protection of sensitive data, the practice often lacks an overview as well as a guide of how to use methods and tools available for data protection. Those few existing overviews are limited and therefore scarcely used [18,77]. Here, we provide an overview based on central criteria describing a context for privacy-preserving data handling, which allows informed decisions in view of the many alternatives.

Several questions arise for a data provider, when anonymizing data or restricting access to data is necessary or desired: What types of access do data users prefer and why? What

information/data are needed in detail to answer certain research questions? What methods are available to meet data protection requirements and user needs? Is the theoretically best method feasible in terms of the manpower and the cost of implementing and maintaining it? To motivate and exemplify our discussion of methods that address these questions from an application-specific point of view, three general use cases are described, to which we will come back later.

1.1 Motivating use cases

a) A researcher in social sciences wants to use data from an education registry to predict student success based on certain characteristics of the educational institute. Her requirement is that all cell values and marginal distributions correspond to the truth. For this purpose, student exam data in the registry are aggregated before it can be accessed. This is a case we later abbreviate as *cna* (and also *cpa*, if perturbation would be allowed).

b) For analysing the differences of income distribution based on individual-level attributes, a researcher from an official statistics department requires access to scientific-use or open-data from the European Statistics on Income and Living Conditions. To avoid the cumbersome procedure of gaining permission to such data with a limited scope only, she decided

✉ Matthias Templ
matthias.templ@zhaw.ch

Murat Sariyar
murat.sariyar@bfh.ch

¹ Zurich University of Applied Sciences, Institute of Data Analysis and Process Design, Rosenstrasse 3, 8401 Winterthur, Switzerland

² Bern University of Applied Sciences, Engineering and Information Technology Institute of Medical Informatics, Höheweg 80, 2502 Biel, Switzerland

to use synthesized data, which are openly available under the General Public Licence 2.0 and used in many scientific articles for methodological development [2]. This is a case we later abbreviate as *cpo*.

c) A health insurance company wants to share individual-level data with universities and hospitals for analysis purposes. To achieve this goal, it produces scientific-use files. For example, Wei et al. [86] analysed the degree of regional variation and effects of health insurance-related factors with anonymized patients claim data from the largest insurance provider in Switzerland. This is a case that we will call *cpo* later (and *cpa* & *cna* if aggregated data are sufficient for the research questions as well as *fno* & *fpo* if more health insurance companies want to provide data sticking to the data locality principle).

1.2 Perspectives on approaching anonymization

As both computer science and statistics are concerned with data processing, they produced several methods for protecting sensitive data prior and during data analysis. Even though there is a considerable overlap between the methods in both areas, the assumptions are usually diverging. This can be exemplified by the difference between statistical and machine learning approaches: whereas it is usual and often necessary to assume a proper probability distribution for the former field, training examples that are somehow representative for a domain are sufficient for the latter one. With respect to anonymization, differential privacy represents an approach from the computer science domain, proposed by Dwork [26]. It is an approach that does not assume any attack scenario (besides composition attacks based on the number of queries issued), as the principle is to noise the data set in such a manner that the result of the query does not depend on whether an individual record is in the data set or not. However, this is not a free lunch and has usually a price in terms of data utility, and the question of efficiency might arise.

From a traditional and more statistical perspective, a procedure for protecting sensitive data should be based on a disclosure scenario, dealing with risks and utility at the same time. A disclosure scenario depends on the instinctive motivation of an attacker, his or her type of prior knowledge, and the data available to match with. Typically, one defines key variables (quasi-identifiers) by considering external matchable data sources, e.g. public-available registers, data from social networks or mobile phones, commercial data sets on individuals, etc. The disclosure risk is the probability that an attacker can learn something new about an individual, which depends on the number and nature of quasi-identifiers in the data set, the required prediction quality, and the prior knowledge of the attacker. To define the level of accepted risk, several additional parameters must be taken into account, e.g. the IT-security, the circle of users, the distribution form of

data (from open-data to highly secure environments, including federated ones), regulations for specific subject-matter areas, and the sensitivity of information.

1.3 Disclosure risk scenarios

The common *identity disclosure* strategy of an attacker is based on some sorts of record linkage [42], for which k -anonymity is still the basic countermeasure principle, guaranteeing that the probability of singling-out an individual is at least $\frac{1}{k}$. Cases of inferring new information about an individual without necessarily identifying her are named *attribute disclosure*. This is possible due to the lack of diversity of sensitive variables, which is typically measured by or variants of l -diversity [50] or t -closeness [49]. To give a striking toy example: In a small published data set, **all** five men that are living in the town of Winterthur and are within the age range of 45-50 do have mental disorder. In this case, the sensitive information (the disorder) is known for **each** individual person having these characteristics. In opposition to the record linkage scenario, an attacker might have additional private knowledge about an individual. For example, the intruder could belong to the circle of acquaintances of a statistical unit or has obtained information by chance. The calculation of the disclosure risk must be customized to these different scenarios.

When discussing different disclosure risk scenarios and related methods, a trade-off between utility and privacy is often assumed. It might be useful to suspend this assumption for some use cases in view of recent research. For example, some artificial intelligence methods promise to enhance utility by including privacy-enhancing mechanisms as a form of regulation [1]. A generalized form of the trade-off assumption is that preserving privacy requires changes either to the original data or to the way it is analysed, and the impact in terms of utility should be assessed quantitatively as well as qualitatively, e.g. with respect to the accessibility of data.

1.4 Outline

Section 2 first reviews the literature on ontologies and categorizations for anonymization and data access methods. After that, we justify our categorization and the criteria that are used for it, which were already used via their abbreviations to label our motivating use cases. In Sect. 3, central anonymization methods are described under its most appropriate category. Classical methods as well as recent developments are covered there. Section 4 summarizes our proposed categorization and discusses the new ontology in light of other research.

2 The need for categorizing anonymization approaches

In particular, the diversity of fields that develop and use methods for anonymizing data is not conducive for gaining an overall perspective, which makes it difficult for data providers to choose the right approach. To give an example, consider use case (b) in Sect. 1.1. In principle, a high number of different methods could be applied here, but use-case specific there is only one appropriate solution. An ontology or a categorization of methods and concepts helps to find such a solution. Cunha et al. [18] formulated an ontology of methods that distinguish between structured (categorical and numerical data), semi-structured (e.g. graph data, XML), and unstructured data (e.g. textual data, multimedia data) as well as between offline and real-time anonymization. For example, they recommend to use methods for achieving k -anonymity [69] for categorical structured data. Templ [77] has also attempted to define an ontology of methods, but only for microdata perturbation in official statistics. Another example is given by Matsunaga et al. [53], who present several techniques and algorithms for implementing data perturbation (e.g. generalization, suppression, randomization and pseudo-anonymization) and describe a strategy on how and when to use them. However, issues such as benefit-risk trade-offs, different ways of sharing data, constraints defined by applications and methods were not discussed, and hence, many concepts of anonymization could not be covered.

Our proposal aims to support the practical decision-making with respect to the appropriate anonymization method and to broaden the readers' perspectives by discussing a wide range of concepts. Thereby, the needs of the user (data receivers) are brought into focus. Concretely, we present several privacy-preserving data processing approaches with respect to the combination of three central criteria: federated versus centralized data distribution (f or c), providing perturbed or non-perturbed data (p or n), and high-level (aggregated) versus original-level granularity (a or o), leading to 8 possible scenarios, of which fpa and fna are omitted as rather uncommon. We are aware of further criteria such as encrypted versus clear text or batch versus interactive processing, and we will refer to them in due place, e.g. encryption under the heading of federated processing and interactive processing in the context of query servers.

Our justification for the three criteria is as follows: (i) The distinction between federated (f) and centralized (c) approaches is, especially motivated by the developments in computer science with respect to secure multiparty computation [15] and the data locality principle in big data applications [71]. (ii) The difference perturbative (p) and non-perturbative (n) covers the distinction between changing data (e.g. via random variation, noising or swapping)

and suppressing data, locally or globally. Although both approaches are usually combined in practice, this differentiation is widely used for a general decision on the risk-utility trade-off. For example, if the risk is the most important concern or if it is assumed that no anonymization method can achieve a satisfying balance between risk and utility, non-perturbing methods are preferred. In most cases, a balance between risk and utility is desired, requiring the usage of perturbing methods. (iii) Granularity of the data is mainly related to the differentiation between micro (o) and tabular (a) data, which can either be decided or is mandatory, e.g. for producing census tables. Overall, our three criteria cover central aspects that allows to discuss all anonymization methods under general scenarios, which should make it easier for data providers to start and find the appropriate method.

In addition to these general categories, some other aspects of anonymization will be considered as well. First, the difference between offline and online anonymization is relevant for certain applications. Online anonymization is immediately (*on-the-fly*) applied to the results of a query. If a method is applied prior to any query made, we have *offline* anonymization. Another issue is additivity, which should be considered when hierarchies or strata are present in a data set, for example, in cross-tabulated tables, where the total must be equal to the \sum of its subtotals and the subtotals must be equal to the \sum of the individual values. To give a concrete example, the number of 45-50 aged men in Winterthur must sum up to the total of all 45-50 aged men in Winterthur, while the sum of all age classes must sum up to the total number of men living in Winterthur. A further aspect is related to the question whether the results produced are consistent for repeated accesses of information. For the same query one expects the same results, but this does not hold for many perturbative methods such as differential privacy. Finally, we will also consider the utility-risk trade-off of anonymization methods.

3 A new categorization of data anonymization and privacy-preserving access

Combinations of our three central criteria lead to scenarios, to which we assign paradigmatic descriptions in the subsection titles. The methods described in the following under these scenarios are either already frequently applied in practice or are promising new developments which first have to fully prove themselves in real-world applications. For each method described, we include at least one referenced real-world example to emphasize the relevance of it.

3.1 Query servers and related methods for providing aggregated data (cna and cpa)

On a central query server, users do not have access to original data but can issue certain queries to receive aggregated information, which might be additionally obfuscated in order to prevent the leak of sensitive information. The most important methods to achieve privacy for queries in aggregated forms are discussed in the following.

3.1.1 Dashboards as a general use case for query servers

Data dashboards are very popular nowadays to present aggregated information to the public or for internal purposes. One example is the Swisscom Mobility Insights Dashboard¹, where aggregated movement patterns of people are visualized daily. A database is queried internally as soon as the user makes a query via point-and-click. Only those results are reported that are considered not to violate privacy, and others are suppressed *on-the-fly* based on a threshold on the number of entities in a category. Such a threshold rule is often referred to in the literature as a form of a k -anonymity model, which is not quite correct. It is rather a threshold-based primary cell suppression rule. The output is suppressed if the number of entities is below a certain threshold. Here, additional considerations are necessary, since confidential values of suppressed cells can be inferred by comparing and subtracting values as soon as subtotals and totals are reported.

3.1.2 Differential privacy for query servers (cpa)

One prominent method for data obfuscation in this context is differential privacy [26]. With differential privacy, random noise is added to the data set's sensitive attributes, or to a prediction, or to cell values (like counts) of an aggregate. The amount of noise is primarily dependent on how much a single entry or row within a data set will impact the value of the query function. Using differential privacy beyond the setting of interactive queries is criticized in the literature [24]. We will come back to the application of differential privacy in other settings in Sect. 3.4.

If the result of a query is a contingency table, the amount of noise in a differential privacy setting is related to (1) the number of observations in the original data file contributing to the cells of the table and (2) the privacy budget of a user (i.e. number of allowed queries relating to the same type of results). Differential privacy models have two major disadvantages. First, the results of queries are non-consistent and non-additive. The totals do not necessarily have to add up to the sum of the individual values (non-additivity) and one receives

different results for identical queries (non-consistency). Second, differential privacy frequently lead to low utility [20]. For example, Uber has developed an open source framework² to facilitate the integration of differential privacy into existing databases, e.g. allowing to query the average distances of trips without disclosing individual-specific information [47]. The application was criticized for low data utility, see e.g. [48].

While differential privacy methods dynamically add noise to the result of a query, the following methods change the result of each allowed query in advance. This ensures that the level of disclosure risk is known in advance and that the changes to the results are only applied once.

3.1.3 The ABS cell key method for query servers (cpa)

The Australian Bureau of Statistics (ABS) *cellKey* method adds noise to frequency and contingency tables using predefined look-up tables consisting of random number values, which allows producing consistent results from queries on dynamically generated tables [83]. Before the tables are generated, a fix numerical code is assigned to every record. All codes of records falling in a cell are summed, and this sum, the cell key, is used for selecting the random number from the look-up table. The resulting tables are non-additive, as the cells are perturbed independently. Therefore, the marginal sums of the altered values are usually different from the sum of the inner cells contributing to this (partial) sum [58]. An optional additional step may try to achieve additivity by adjusting the noise. However, doing so, the tables are no longer consistent. A software implementation of the *cellKey* is provided by [56].

For example, the German Census 2022³ and the higher education statistics 2021 in Germany [28] are already anonymized with the *cellKey* method [56]. Any result of a query is perturbed with a predefined noise. Even if this method can be used in an interactive environment, it is primarily a static method that is applied prior to any access to a data set.

3.1.4 Controlled tabular adjustment (cpa)

Controlled tabular adjustment (CTA) [35] replaces sensitive cell values by their closest safe values, while preserving additivity through small adjustments to the other cells. Like secondary cell suppression, it requires complex mathematical optimization approaches based on mixed integer linear programming, which makes it difficult to apply the method in

¹ <https://mip.swisscom.ch>, accessed 11.08.2021

² <https://github.com/uber-archive/sql-differential-privacy>, status: deprecated. Accessed 11.08.2021

³ <https://www.zensus2022.de/DE/Zensusdatenbank/Geheimhaltung.html>, accessed 11.08.2021

practice. While CTA was applied to test data sets [13,14,33], up to our knowledge no real-world application is described in the literature, probably because the related algorithms are not freely available and due to the fact many statistical offices seems to prefer suppression methods over interval publication methods for tabular output.

3.1.5 Secondary cell suppression (cna)

Secondary cell suppression suppresses additional non-risky cells, as this is often necessary to protect the risky cells. It has two steps. In a first step, those results obtained from records with low frequency (other rules can be applied as well [42]) are suppressed. This is similar to the application of the k -anonymity principle, explained in Sect. 3.3. The resulting table can frequently be compromised just with basic arithmetic's, which requires that—in a second step—additional information is suppressed (hence *secondary cell suppression*). The mathematical formulation of this (NP-hard) problem can be found, e.g. in [30], while related software is described in [43,57].

Suppression of cells in tabular data is widely used in practice, even if this fact is not reflected in the literature. For one exception to this matter of fact, see [16], which evaluated secondary cell suppression on employment data. Secondary cell suppression is typically applied on (hierarchical and/or multi-dimensional) tabular data available on websites of statistical agencies that allow to query these data. For example, the statCube statistical static query server from Statistics Austria allows accessing dozens of data sets in aggregated and partly suppressed form⁴.

3.2 Remote execution or access to the original data (cno)

Remote execution—if done seriously—involves a lot of work on the part of the data provider [41], not only for maintaining software on the server, where the code of the users is executed. The structure of the data must be made available to the external researcher so that they can develop the code to be applied on the data. After the code is developed, it is executed by the data holder (central) on the original data. The results are checked by the data holder and only made available to the researcher if they do not violate anonymity criteria. In addition to that, usually a loop starts: the researcher checks the results, and most probably adjusts his code, which is then again executed by the data holder, and so on. As this iterative process is labour-intensive, remote execution is not often used in practice [10,27].

⁴ https://www.statistik.at/web_de/services/statcube/index.html, accessed 11.08.2021

An example of a remote execution system is the German Research Data Centre, which allows access to dozens of data sets through remote execution. For example, the Diagnosis-Related Groups Statistic (DRG) survey is an annual complete survey of all stationary hospital cases in Germany. The data user has no direct access to this data set, but can send code to be executed. Another example is the remote evaluation system of the Institute for Employment Research (IAB). Researchers can send code to be executed by IAB staff, and the results are checked by the staff before they are sent to the researchers (see [29] for details).

In contrast to remote execution, remote access to the original data under continuous data usage control measures can be granted. One example is the UK Data Service Secure Lab⁵ that provides secure access to a number of non-anonymized data. An example of a data set that is made available is the geospatial data from the UK Labour Force Survey. Another example is the remote access to data from official statistics and social insurance institutions in Germany [87]. A variation to the remote-access scenario is a secure lab, where the network connection is replaced by physical travel to the data holder, who provides a computer with strict security measures in place (no USB interface, no internet, etc.).

3.3 Anonymization for sharing public-use files (cpo)

As soon as one wants to share confidential non-aggregated data, typically a flat file with variables as columns and observations on individuals as rows is used in order to apply anonymization methods (described below) that are designed for structured individual-level data. An example is the anonymization of data by the World Bank⁶. They provide detailed data on a large scale and advertises anonymization with traditional methods [6]. Also, dozens of national statistical agencies including the US Census Bureau, Statistics Denmark, Statistics Netherlands, and Statistics Austrian, as well as companies and institutions in other areas [17] share anonymized data on individual level.

The traditional road to anonymization for data publishing is based on three steps: (1) quantification of the current re-identification risk of each individual person in the data set, (2) anonymization/data perturbation, and (3) measuring the resulting data utility. The first step is based on the disclosure scenario, and usually k -anonymity [68] or uniqueness on subsets [51] is used for population-related data, i.e. the data set is related to whole populations. Such methods are implemented in well-known software tools [66,80]. For survey samples, the disclosure risk should be quantified by taking the sample and

⁵ <https://www.ukdataservice.ac.uk/use-data/secure-lab.aspx>, accessed 11.08.2021

⁶ <https://microdata.worldbank.org/index.php/home>, accessed 11.08.2021

(estimated) population frequency counts into account [32], i.e. the fact that persons of a population may or may not be included in a sample, as well as other peculiarities, such as hierarchical structures (e.g. persons in households or companies). Related methods (e.g. [32]) are implemented in the software tools μ -Argus [44] and sdcMicro [80], the latter being both free and open source.

For the second step, dozens of methods are described [77] and available via software implementations [44,66,80]. Typically, anonymization is a highly iterative and exploratory process with a lot of fine-tuning in the selection of methods and their parameter settings in order to reduce disclosure risk to an acceptable level while providing high data utility. Combination of (global or local) re-coding and suppression are the most common methods. The latter inserts missing values into the (micro-)data set to replace specific values of individual variables. Note that local suppression (find an optimal suppression pattern) is a multivariate problem that is NP-hard and cannot be optimally solved in a reasonable time. If there are many categorical key variables (e.g. more than 6), re-coding and local suppression may not sufficiently reduce the risk of re-identification, or it may lead to a significant loss of information. In this case, swapping using post-randomization (PRAM, [36]) may be a more efficient alternative. PRAM is a probabilistic method and swaps the categories of observations for a categorical variable based on a predefined transition matrix that specifies the probabilities of transitions from one category to another. Depending on the data structure, methods to add (correlated) noise [81], microaggregation [23], and other kinds of swapping methods [19,36,45,62] might be applied as well.

After the data have been anonymized, it is important to assess the information loss and the data quality. This is done, for example, by comparing analysis results on the original and the anonymized data, e.g. comparing contingency tables, output from regression models, distributions, point and variance estimates, etc. Templ [76] argues that rather data- and use-case-dependent measures should be used instead of general purpose measures (such as means and correlations).

3.3.1 Differential privacy for sharing non-aggregated, static data

Differential privacy was originally designed for queries to a database. Each query consumes the privacy budget, and the lower the budget, the more noise is added. If the privacy budget is used up, queries are halted to prevent information reconstruction. In the case of a static application to individual data, however, the number of queries is unlimited, while the Laplace noise ϵ can only be fixed once. It is a challenge to choose ϵ in a way that ensures sufficient anonymization and does not heavily attenuate the data utility. In fact, differential privacy turns out to be not a serious competitor to traditional

anonymization for data publishing. Many articles warned of its misuse in other forms that it is initially designed for [24,60,89].

The static application of differential privacy to anonymize non-aggregated data was used, for example, by the US Census Bureau for its census. Some problems were the following: (i) The occupancy status differs from the population figures (non-additivity); (ii) in some housing units all individuals are reported as occupied, even when children live the household [89]. Further, mortality rates can be distorted substantially, sometimes by more than 100% [38]. Facebook also used differential privacy, which led to poorer data utility than the initial version of data released in 2018 using data aggregation [24].

3.3.2 Synthetic data

Another way of perturbing the data is by generating synthetic data that exhibit the same characteristics as the original data. Methods for creating synthetic data stem from the machine learning (ML), AI and statistical modelling domains. They learn parameters of models on the original data and use them to generate artificial data. Synthetic data are associated with low disclosure risk [54,75,78] and advantages in the simplicity of the process - once synthetic data are generated. However, it is not always possible to generate “synthetic data is as-good-as-real” and data utility is often lower than using traditional statistical disclosure control techniques [73]. Moreover, data can be associated with complex cluster structures, logical rules, missing values, outliers, multiple relationships and complex sampling schemes, rendering it highly difficult to synthesize such data [82]. However, synthetic data might still be useful for certain general analyses, training and education, open-data for method development, and for remote execution (see Sect. 3.2). On the other hand, perfect synthetic data may compromise privacy.

Two examples of synthetic data that are used heavily in literature are the synthetic data of the European Statistics on Income Living Conditions (EU-SILC), produced by [2] and the synthetic version of the Structure of Earnings Survey [79]. Software for generating such synthetic data is, for example, provided by synthpop [63] for data without any of the mentioned complicated structures and by simPop [82] for more complex data sets. Recently, also methods from artificial neural networks are used to synthesize data, for example, using generative additive networks (GAN) [59,73].

3.4 Federated privacy-preserving analysis and computation (fpo and fno)

Using new techniques for data analysis that stick to the data locality principle promises to unlock data that was previously deemed too sensitive and complex for being anonymized

and published. One prominent example is the retrospective analysis of genetic data across multiple sites in a federated privacy-preserving computing environment [15]. The predominant scenario is fno, as data locality usually comes with the reassurance that only highly aggregated end results will be published. We included fpo (and fpa) as well here, because there are use cases, where the data provider wants to ensure that internal privacy breaches are not likely and that not all sorts of results are passed on unperturbed (see the discussion in Sect. 3.1 and below). Central concepts for privacy-preserving computation used in this context are

- Differential privacy [12,26], where noise is added to the retrieved information and the analysis results, making it impossible to reverse engineer the individual inputs.
- Secure multiparty computation (SMC) [15,46] where the data analysis is distributed across multiple parties so that no single party can see the full set of the inputs.
- Homomorphic encryption (HE) [9,72] where the data are encrypted before it is passed on so that it can be analysed but not converted back into the original information.
- Federated learning and analysis [88] where parties share insights from the analysis of their data without sharing the data itself.

Secure multiparty computation and homomorphic encryption allow computations on the original but encrypted data sets that are dispersed across different sites. Both are the backbone for many federated analysis applications in banking and fintech, insurance sector, healthcare and medicine, retailing industry, as well as in recommendation systems [7,88]. As described in Sect. 3.2, it is especially useful for developing analysis tools and producing results in a collaborative manner. For ensuring semantic and syntactic interoperability among those sites, meta-data and data entry examples should be provided. The two most well-known drawbacks of such systems are architectural complexity and performance issues [7].

One prominent tool for federated analysis in medicine is provided by the R package DSOpal [52], the DataSHIELD implementation for Opal, a core data warehouse application that provides all the necessary tools to import, transform and describe data. DSOpal has been used for various projects. Two of them were conducted by the international research network *InterConnect* with respect to diabetes and obesity. They set up a federated database infrastructure, allowing a secure analysis of harmonized datasets across participating studies from around the world, without sharing individual-level data [22]. A similar privacy computation approach was used for diabetes predictions in a study conducted by the main hospital in Shanghai and further 16 branches across China [37].

Recently, differential privacy has also been proposed for a federated form of applying machine learning methods. Differential privacy is used for guiding the aggregation of results in each step of the data processing by obfuscating the (intermediate) results, as they might leak sensitive information (rare case of fpa). The main drawback for utilizing differential privacy in this context is the lack of reliability of the results. In particular, updates of models on the same data require increasing amounts of noise for protecting sensitive information [24,86]. In addition to that: if data from one client consists of multiple entries related to one person, the concept of differential privacy may no longer work. Moreover, [40] showed—based on the work of [70]—that GAN models can predict a victim’s data very well and thus undermine obfuscation generated by differential privacy. Differential privacy for decentralized data sets has been mainly applied to image data [55].

There are suggestions to change the perspective on the drawbacks of differential privacy. Instead of mourning about the reduced utility, adding noise could be regarded as some kind of regulation, even when privacy is not an issue. One example is PATE (private aggregation of teacher ensembles), which coordinates the development of ML models on different data sets by ensuring that the models find general patterns instead of over-fitting and thereby disclosing sensitive information, e.g. that a certain person has cancer [64]. Models trained on non-overlapping training sets are assessed to have learned general patterns, if they produce the same results on test data. Otherwise, they rely too heavily on the training data, in which case PATE adds noise in such a manner that the parameter estimates stays roughly the same if a single training sample is omitted or changed. Thereby, the ML model does not focus on specifics, which is similar to the use of dropouts in deep neural networks [84].

PATE and similar methods of federated learning are very susceptible to so-called Byzantine [8] or poisoning attacks [74]. One single, non-colluding malicious peer is sufficient to impact the federated learning to the extent that a globally false model results. Homomorphic encryption may help here to prevent such security attacks, but have other issues [25].

3.5 Summarizing overview as a guideline for the practice

Our proposal for an ontology on concepts and methods for privacy-preserving data access, analysis and publishing is summarized in Table 1. It covers 11 general methods in 4 categories.

We provide pro’s, con’s and implications (touching risk-utility aspects). General rules for their application are also given with references to use cases. In addition to that, we indicate whether the approach can be used interactively (online versus offline (before data release)), whether one can rely on

receiving the same data with the same access request (consistency), and whether the analysis can rely on additivity in resulting contingency tables of categorical data as these issues are very central for users. The reason to aggregate *cpa* & *cna* as well as *fpo* & *fno* lies in recent proposals of hybrid approaches, most often relying on the principles of differential privacy. Category 1 covers methods for anonymizing tabular data as well as for producing highly aggregated results. Category 2 is the one that is typically associated with anonymization: providing perturbed versions of microdata (individual-level data) for scientific and public use. Category 3 subsumes approaches for centralized access to microdata without perturbation. It relies heavily on data usage control and output checking and does not produce publicly available data. Category 4 covers federated approaches such as secure multiparty computation and is especially appropriate for multi-centered data analysis on a large amount of data. Such an overview helps in decision-making related to the appropriate approach for anonymization and related approaches. It is important to extend such an ontology to cover other types of data, such as unstructured text, signal and image data.

4 Discussion

Every data provider should have justified reasons for deciding whether and which form of privacy-preserving processing of the data is suitable for her particular data usage scenarios. For example, it is not the case that synthetic data is the panacea for all situations, nor any other concept. For some purposes, federated data analysis and interactive queries are useful, but for many scenarios outside the big-data context, the implementation and maintenance costs are too high. In addition to that, not all sorts of computations can be decomposed in a fashion that is necessary for federated analysis. In particular, the requirements of reproducible research may favour traditional methods with their focus on releasing anonymized data to a research community with certain utility assurances.

Many research questions can be answered by aggregated or noised data that have the same marginals as the raw data. However, if reasoning for individual cases is required (for example, in clinical settings), methods from our *a* (aggregated) and *p* (perturbed) categories should be avoided. Even methods in the *f* (federated) category could be problematic in such cases, as the data quality and differences between the local data repositories cannot be assessed adequately. Almost all methods need extensions to tackle challenges such as high-dimensional data [65], anonymization of event-history data [39], anonymization of trajectory/mobility data [85], or time-varying sensitive features.

Although the original default settings of differential privacy were interactive queries with aggregated results, nowa-

days, it is often used in other settings as well, e.g. for noising the output of the prediction of machine learning models. Training ML methods on original data and noising the results by differential privacy principles is seen to be an alternative to traditional anonymization. However, differential privacy caused a lot of discussion, and various authors criticized this use of the DP methodology [5,31], especially when used in a non-interactive environment [24]. Results modified by means of differential privacy are non-consistent and non-additive, a potential major concern for many users and data holders, and further positive empirical evidence of its usefulness is needed.

We have concentrated our discussion on preventing identity disclosure, other disclosure risks, such as those related to attribute or membership disclosure, are only hinted to in this paper, as the former risk is to be addressed in almost every case associated with the protection of personal data. For the identity disclosure risk, the following three general advises for the practice can be extracted from our discussion: (i) If individual-level data are provided for external user or data must be provided for reproducible research, traditional methods (anonymization) should be used, as they address the utility-risk trade-off by design. (ii) In case of interactive queries, differential privacy is a feasible and intuitive alternative to traditional methods for multi (approx. > 5)-dimensional tables, simply because too much perturbation for controlled tabular adjustment (using the *cellKey* method or secondary cell suppression) would be needed if the number of possible queries are not restricted. (iii) For privacy-preserving federated analysis, new methods such as PATE, or secure multiparty computation via homomorphic encryption are promising for ML predictions, but still have to show their usefulness and feasibility in practice.

Some data providers and users do not accept non-consistent and non-additive results (e.g. produced by applying differential privacy), even if the final analyses might not be affected by it. Data users usually want to explore the data to gain a feeling to interesting research questions. This has to be taken into account in addition to the risk-utility considerations that focus on a limited number of analyses or general purpose metrics. Hence, there are solid and scientific reasons why producing public-use file are the most desired forms in practice. Yes, once the data are out there, no guarantee can be given that de-anonymization is impossible, as the computational and methodological developments cannot be assessed in advance. However, such indeterminable risks should not guide concrete decisions.

Our categorization is destined to be comprehensive and simple at the same time, which generate advantages over existing taxonomies. For example, Cunha et al. [18] would classify our motivating examples (b) and (c) in Sect. 1.1 as cases for applying *k*-anonymity, related principles such as *l*-diversity or differential privacy, not considering possibil-

Table 1 Overview of the methods for privacy-preserving data access, analysis and publishing. Online: whether the approach can be used interactively. Consistency: whether one can rely on receiving the same data with the same access request. Additivity: whether the analysis can rely on additivity in resulting contingency tables of categorical data (either explanatory or target variables). In the right column of the table, references to real-world use cases are provided as well.

Method	on-line	consis-	ten-ty	addit-	ivity	pro's	con's and implications	when to apply
1) cpa and cna						Producing anonymized tabular data and highly aggregated results		
Adding random noise using differential privacy (cpa)	Yes	No	No	No	Simpler than secondary cell suppression and controlled tabular adjustment. Often low disclosure risks.	Cell values are not truth-preserving (e.g. yielding households with only children). The lower the privacy budget, the higher the noise leading to low data utility [67].	Esp. for $p \geq 5$ -dim. tables, and when non-truth-preserv. methods accepted. Use case: [47].	
Controlled tabular adjustment (cpa)	No	Yes	Dep-ends	Replac-ing sensitive cell values by closest safe values and small adjustments made to other cells may result in low risk and relatively high utility when $p < 5$.	Resource-intensive. Changing cell values of a table might not find acceptance from users.	When non-truth-preserving methods do find acceptance and $\approx p < 5$. Use case: [13].		
Deterministic noise by the <i>cellKey</i> method (cpa)	Yes	Yes	No	Implementation in software exists and it avoids some flaws of differential privacy. Well-established method.	Cell values are not truth-preserving.	When non-truth-preserving methods do find acceptance. Use case: [28].		
Secondary cell suppression (cna)	No	Yes	Yes	Standard and popular method. Reported cell values are true.	Some cells are suppressed in addition to risky ones. Often non-feasible for large, high-dim. tables. Can result in too many suppressions. Resource-intensive and high comput. costs.	Especially when noise and rounding are not an option and when $\approx p < 5$. Use case: [16].		
2) cpo					Providing perturbed versions of microdata (individual-level data) for scientific and public use			
Synthetic data	No	Yes	Yes	Very low disclosure risk. Once synthesized, low to zero man power needed. Eliminates the bureaucratic burden associated with gaining access to sensitive data.	For various types of analyses and data structure, the data utility is too low [82]. Consequently, synthetic data can lead to incorrect findings and incorrect decisions.	Data for education, remote execution, open-data, micro-simulation tasks. Use case: [2].		

Table 1 continued

Method	on-line	consistency	additivity	pro's	con's and implications	when to apply
Classical anonymization	No	Yes	Yes	Ones done, no further costs. Detailed (non-aggregated) data can be shared. Often data need only be changed very slightly to greatly reduce disclosure risk.	In case of a high number of quasi-identifiers, data utility can be quite low [77]. Needs experience and expertise to produce anonymized data with high utility.	Open-data or scientific-use files to be shared. When reproducible research is a need. Use case: [6].
Differential privacy	No	Yes	Yes	Once done, no further costs. Detailed (non-aggregated) data can be shared.	Non-plausible values possible, e.g. households with only children [89]. Yield often less utility than classical methods [25]. Utility diminishes when remaining priv. budget is low.	Open-data or scientific-use files to be shared. Use case: [34].
3) cno	Providing centralized access to microdata without perturbation. Relying on data usage control and output checking					
Remote execution	No	Yes	Yes	Very low disclosure risk with the opportunity to analyse the original data. Requires synthetic data for developing and testing the code before applying it. Access to original pseudo-anonymized data.	Highly reduced flexibility. Output checking can hardly be automatized, is costly and needs maintenance [10,11,27,41].	When remote access is not possible. When conformation and not exploration is the focus. Use case [29].
Remote access	Dep-ends	Yes	Yes		Not allowed in some countries. Computations are carried out on a server without admin rights of the user. Needs final careful output checking. Needs maintenance.	If privacy laws permit and staff are available to review the results. Use case: [87]
4) fpo, fno	Federated computation on original microdata, especially for multi-centered data analysis on a large amount of data					
Federated analysis using secure multiparty computation	Yes	Dep-ends	Dep-ends	Yields exact results with low disclosure risks. After the system is established, low running costs.	Not feasible for every sort of analysis. Differential privacy might be necessary for intermediate results, lowering utility [25]. High costs for implementation. Reduced flexibility. Might not be safe from privacy attacks using GANs [40]. Maintenance work when data structures changes.	When large amounts of data are distributed and frequent statistical questions are to be answered with mutual benefits for all data providers. Use case [55].
Federated learning using PATE	Yes	No	No	Non-sharing of data but of trained model parameters or predictions. Training models on data from several peers. Prediction on test data.	For users: black box approach. Noising results reduces data utility. High costs for installment. Maintenance work when data structures changes.	When large amounts of data are distributed and machine learning methods have to be trained, e.g. BERT-fine-tuning with large corpora. Use case: [21,22].

ity of generating synthetic data. In addition to that, their taxonomy gives no advice for the motivating example (a). Templ's categorization [77] would place (b) and (c) in our category cpo as candidates for applying data perturbation methods, and would also give no clues about the appropriate approach for case (a). Our proposal provides a more strongly differentiated classification in terms of central properties, which allows considering viable alternatives for these different cases. Example (a) would be subsumed under cna in our proposal, proposing secondary cell suppression that leads to consistent and additive tables produced in an offline process. If the utility is assessed to be too low, a cpa-approach such as *cellKey* could be suggested, while hinting to the associated problems.

Motivating example (b) would be classified as cpo in our proposal, as the generation of synthetic data can be seen as a data perturbation method that guarantees to retain the general characteristics while perturbing the data in a complex sampling scheme, leading to a low disclosure risk. From an ontological point of view, it might be asked whether the generation of synthetic data should be placed in a separate category, as it does not really perturb the original data. On the one hand, subsuming it under cpo allows considering viable alternatives, for example, if the data user requires having real-world data. On the other hand, a separate category could help to infer characteristics of synthetic data generation that are conducive for developing better alternatives or extensions.

Motivating example (c) would be classified as cpo, relying on classical anonymization techniques. For some research questions, it might be sufficient to use aggregated data, in which case methods under cpa or cna would be proposed. To make such distinctions requires having an overview of potential use cases and their characteristics, which is often not that simple. To facilitate the related decisions, it might be conducive if the data provider lists research questions that are associated with individual-level and aggregated public-use files generated from their individual-level data. Extending example (c) by considering data from more than one health insurance company and a researcher that is interested in overall costs for certain patient groups, fno or fpo would be suggested.

In conclusion, the possibility to swiftly generate suggestions and recommend alternative anonymization approaches for such use cases, shows the main advantage of our proposal over existing taxonomies. It might be useful to extend our proposal by means of relying on an upper-level ontology, such as BFO (basic fundamental ontology, [4]) or SKOS (simple knowledge organization for the web [61]), which are widely used in different domains. However, whether the effort is really worth it is difficult to estimate. For example, [3] developed a SKOS-based ontology for privacy-related issues in the context of state surveillance activities. It represents general privacy requirements and rules, and significant efforts

are necessary to produce a practice-relevant guideline such as our taxonomy.

Funding Open access funding provided by ZHAW Zurich University of Applied Sciences

Research Data Policy and Data Availability Statements Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no conflicts of interest to declare that are relevant to the content of this article.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abadi, M., Erlingsson, U., Goodfellow, I., McMahan, H.B., Mironov, I., Papernot, N., Talwar, K., Zhang, L.: On the protection of private information in machine learning systems: Two recent approaches. In 2017 IEEE 30th Computer Security Foundations Symposium (CSF), pages 1–6, (2017)
2. Alfons, A., Kraft, S., Templ, M., Filzmoser, P.: Simulation of close-to-reality population data for household surveys with application to EU-SILC. *Stat. Methods Appl.* **20**(3), 383–407 (2011)
3. Arguedas, V.F., Izquierdo, E. and Chandramouli, K.: Surveillance ontology for legal, ethical and privacy protection based on SKOS. In 2013 18th International Conference on Digital Signal Processing (DSP), pages 1–5, (2013)
4. Arp, R., Smith, B. and Spear, A.D.: Building ontologies with basic formal ontology. The MIT Press, (2015)
5. Bambauer, J., Muralidhar, K., Sarathy, R.: Fool's gold: An illustrated critique of differential privacy. *Vanderbilt J. Entertain. Technol. Law* **16**(4), 701–755 (2014)
6. Benschop, T., Welch, M.: A practice guide for microdata anonymization. In Joint UNECE/Eurostat work session on statistical data confidentiality, the Hague, Netherlands, (2019)
7. Blake, M., McWaters, J., and Galaski, R.: The next generation of data-sharing in financial services: Using privacy enhancing techniques to unlock new value. *World Economic Forum*, pages 1–35, (2019)
8. Blanchard, P., El Mhamdi, E.M., Guerraoui, R. and Stainer, J.: Machine learning with adversaries: Byzantine tolerant gradient descent. In: Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., and Garnett R., editors, *Advances in*

- Neural Information Processing Systems, vol 30. Curran Associates, Inc., (2017)
9. Blatt, M., Gusev, A., Polyakov, Y., Goldwasser, S.: Secure large-scale genome-wide association studies using homomorphic encryption. *Proc Nat Acad Sci* **117**(21), 11608–11613 (2020)
 10. Bond S., Brandt M., and de Wolf P-P.: Guidelines for the checking of output based on microdata research. Technical report, ONS, DeStatis, CBS, 2013. Project No: 262608. Data without Boundaries. WORK PACKAGE 11 (Improved Methodologies for Managing Risks of Access to Detailed OS Data). D11.8 - Final reports of synthetic data CTA, ECTA, cell suppression & Guidelines for output checking
 11. Bond S., Brandt M. , and de Wolf P-P.: Guidelines for output checking. Technical Report European Commission, FP7 - SP4 Capacities, Project number 262608, Data without boundaries, (2016)
 12. Bonomi, L., Jiang, X., Ohno-Machado, L.: Protecting patient privacy in survival analyses. *J. Am. Med. Inform. Assoc.* **27**(3), 366–375 (2019)
 13. Castro, J.: Present and future research on controlled tabular adjustment. In: Joint UNECE/Eurostat work session on statistical data confidentiality, Tarragona, Spain (2011)
 14. Castro, Jordi, González, José A.: A linear optimization-based method for data privacy in statistical tabular data. *Optimiz. Methods. Softw.* **34**(1), 37–61 (2019)
 15. Cho, H., Wu, D.J., Berger, B.: Secure genome-wide association analysis using multiparty computation. *Nat. Biotechnol.* **36**(6), 547–551 (2018)
 16. Cohen, S., Bogong, T.L.: A comparison of data utility between publishing cell estimates as fixed intervals or estimates based upon a noise model versus traditional cell suppression on tabular employment data. Research report of the Bureau of Labor Statistics, Washington, D.C
 17. Crampin, A.C., Dube, A., Mboma, S., Price, A., Chihana, M., Jahn, A., Baschieri, A., Molesworth, A., Mwayeghele, E., Branson, K., Floyd, S., McGrath, N., Fine, P.E.M., French, N., Glynn, J.R., Zaba, B.: Profile: The Karonga health and demographic surveillance system. *Int. J. Epidemiol.* **41**(3), 676–685 (2012)
 18. Cunha, M., Mendes, R., Vilela, J.P.: A survey of privacy-preserving mechanisms for heterogeneous data types. *Comput. Sci. Rev.* **41**, 100403 (2021)
 19. Dalenius T. and Reiss S.P.: Data-swapping: A technique for disclosure control. In: Proceedings of the Section on Survey Research Methods, vol 6, pages 73–85. American Statistical Association, (1982)
 20. Davis J.S.II and Osonde A.O.: Privacy preservation in the age of big data: a survey. RAND Corporation, Santa Monica, CA, (2016)
 21. Devlin, J., Chang, M.W., Lee, K. and Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Vol 1 (Long and Short Papers), pages 4171–4186, Stroudsburg, PA, USA, (2019). Association for Computational Linguistics
 22. Doiron, D., Marcon, Y., Fortier, I., Burton, P., Ferretti, V.: Software Application Profile: Opal and Mica: open-source software solutions for epidemiological data management, harmonization and dissemination. *Int. J. Epidemiol* **46**(5), 1372–1378 (2017)
 23. Domingo-Ferrer, J., Mateo-Sanz, J.M.: Practical data-oriented microaggregation for statistical disclosure control. *IEEE Trans. Knowl Data Eng.* **14**(1), 189–201 (2002)
 24. Domingo-Ferrer, J., Sánchez, D., Blanco-Justicia, A.: The limits of differential privacy (and its misuse in data release and machine learning). *Commun. ACM* **64**(7), 33–35 (2021)
 25. Domingo-Ferrer, J., Blanco-Justicia, A., Manjón, J., Sánchez, D.: Secure and privacy-preserving federated learning via co-utility. *IEEE Internet Things J.* **9**(5), 3988–4000 (2021)
 26. Dwork C.: Differential privacy: A survey of results. In: Proceedings of the 5-th International Conference on Theory and Applications of Models of Computation, TAMC 2008, page 1-19, Berlin, Heidelberg, (2008). Springer-Verlag
 27. Emily G., Greci C., Kotrotsios Y., Parker S., Scott J., Welpton R., Wolters A., and Woods C.: Handbook on Statistical Disclosure Control for Outputs. Technical report, (2019)
 28. Endeled, T.: Die Geheimhaltung mit der Cell-Key-Methode. *WISTA*, **6**, (2019)
 29. FDZ IAB. Datenfernverarbeitung und gastaufenthalte am fdz der ba im iab. Technical report, Forschungsdatenzentrum (FDZ) der Bundesagentur für Arbeit (BA) im Institut für Arbeitsmarkt- und Berufsforschung (IAB), (2021)
 30. Fischetti, M., Salazar-González, J.J.: Complementary cell suppression for statistical disclosure control in tabular data with linear constraints. *J. Am. Stat. Assoc.* **95**, 916–928 (2000)
 31. Francis P.: Dear differential privacy, put up or shut up. Technical report, (2020). MPI-SWS-2020-005
 32. Franconi L. and Poletti S.: Individual risk estimation in μ -Argus: a review. In J. In: Domingo-Ferrer, editor, Privacy in Statistical Databases, Lecture Notes in Computer Science, pages 262–272. Springer, (2004)
 33. García, S.H., Salazar-González, J.J.: Enhanced controlled tabular adjustment. In Joint UNECE/Eurostat work session on statistical data confidentiality, Tarragona, Spain (2011)
 34. Garfinkel S.: Differential privacy and the 2020 us census. MIT Case Studies in Social and Ethical Responsibilities of Computing, (Winter 2022), 1 (2022). <https://mit-serc.pubpub.org/pub/differential-privacy-2020-us-census>
 35. Giessing, S.: Pre-tabular perturbation with controlled tabular adjustment: Some considerations. In: Domingo-Ferrer, J. (ed.) Privacy in Statistical Databases. pp. pp. 48–61. Springer International Publishing, Cham (2014)
 36. Gouweleew, J., Kooiman, P., Willenborg, L., De Wolf, P-P.: Post randomisation for statistical disclosure control: Theory and implementation. *J. Official Statist.* **14**(4), 463–478 (1998)
 37. Guo X., Yao Q., Kwok J., Tu W., Chen Y., Dai W., and Yang Q.: Privacy-Preserving Stacking with Application to Cross-organizational Diabetes Prediction, pages 269–283. Springer International Publishing, Cham, (2020)
 38. Hauer, M.E., Santos-Lozada, A.R.: Differential privacy in the 2020 census will distort covid-19 rates. *Socius*, 7:1–6, 2021. online first
 39. Heldal, J.: Anonymised integrated event history datasets for researchers, pp. 1–7. In Joint UNECE/Eurostat work session on statistical data confidentiality, Tarragona, Spain (2011)
 40. Hitaj B., Ateniese G., and Perez-Cruz F.: Deep models under the GAN: Information leakage from collaborative deep learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17, page 603-618, New York, NY, USA, (2017). Association for Computing Machinery
 41. Hochguertel, T., Weiss, E.: De facto anonymity in results. In Joint UNECE/Eurostat work session on statistical data confidentiality, Tarragona, Spain (2011)
 42. Hundepool A., Domingo-Ferrer J., Franconi L., Giessing S., Nordholt-Schulte E., Spicer V., and de Wolf P-P.: Statistical Disclosure Control. Wiley, (2012)
 43. A. Hundepool, R. Ramaswamy, de Wolf P-P., L. Franconi, S. Giessing, D. Reipsilber, J.J. Salazar, C. Castro, G. Merola, and P. Lowthian. τ -Argus software, version 4.1.7, 2018
 44. Hundepool A., Van deWetering A., Ramaswamy R., Franconi L., Capobianchi A., DeWolf P-P., Domingo-Ferrer J., Torra V., Brand R., and Giessing S.: μ -Argus version 5.1 software and users manual, (2015)

45. Ito S. and Hoshino N.: Data swapping as a more efficient tool to create anonymized census microdata in japan. In J. (eds.) In: Domingo-Ferrer, editor, Privacy in Statistical Databases, Lecture Notes in Computer Science, volume 8744, pages 185–199. Springer, Cham, (2014)
46. Jagadeesh, K.A., Wu, D.J., Birgmeier, J.A., Boneh, D., Bejerano, G.: Deriving genomic diagnoses without revealing patient genomes. *Science* **357**(6352), 692–695 (2017)
47. Johnson, Noah, Near, Joseph P., Song, Dawn: Towards practical differential privacy for sql queries. *Proc. VLDB Endow.* **11**(5), 526–539 (2018)
48. Klucar J.: Uber's differential privacy. probably isn't. <https://github.com/frankmcherry>, Feb (2018)
49. Li N., Li T., and Venkatasubramanian S.: t-closeness: Privacy beyond k-anonymity and l-diversity. In 2007 IEEE 23rd International Conference on Data Engineering, pages 106–115, (2007)
50. Machanavajjhala, A., Kifer, D., Gehrke, J., Venkatasubramanian, M.: l-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data* **1**(1), 3 (2007)
51. Manning, A.M., Haglin, D.J., Keane, J.A.: A recursive search algorithm for statistical disclosure assessment. *Data Min. Knowl. Disc.* **16**(2), 165–196 (2008)
52. Marcon Y.: DSOpal: DataSHIELD Implementation for Opal, (2021). R package version 1.3.0
53. Matsunaga R., Ricarte I., Basso T., and Moraes R.: Towards an ontology-based definition of data anonymization policy for cloud computing and big data. In 2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W), pages 75–82, (2017)
54. McClure, D., Reiter, J.P.: Assessing disclosure risks for synthetic data with arbitrary intruder knowledge. *Stat. J. IAOS* **32**, 109–126 (2016)
55. McMahan H.B., Moore E., Ramage D., Hampson S., and Arcas B.A.: Communication-efficient learning of deep networks from decentralized data, (2017)
56. Meindl, B.: cellKey: Implementing ABS cell-key method for adding noise to frequency and continuous tables, 2020. R package version 0.19.1
57. Meindl, B.: sdcTable: Methods for statistical disclosure control in tabular data, 2020. R package version 0.31
58. Meindl B. and Enderle T.: cellKey-consistent perturbation of statistical tables. In Joint UNECE/Eurostat work session on statistical data confidentiality, the Hague, the Netherlands, (2019)
59. Mendelevitch, O., Lesh, M.: Security and Privacy From a Legal, Ethical, and Technical Perspective, chapter Beyond Differential Privacy: Synthetic Micro-Data Generation with Deep Generative Neural Networks, pages 1–14. 09 (2020)
60. Mervin J.: Can a set of equations keep u.s. census data private. *Science*, (2019)
61. Miles A. and Pérez-Agüera JR.: Skos: Simple knowledge organisation for the web. *Catal Classif Quart* **43**(3-4):69–83, 2007
62. Muralidhar, K., Sarathy, R.: Data shuffling- a new masking approach for numerical data. *Manage. Sci.* **52**(2), 658–670 (2006)
63. Nowok, B., Raab, G.M., Dibben, C.: synthpop: Bespoke creation of synthetic data in R. *J. Stat. Softw.* **74**(11), 1–26 (2016)
64. Papernot N., Song S., Mironov I., Raghunathan A., Talwar K., and Erlingsson U.: Scalable private learning with PATE. *arXiv e-print arXiv:1802.08908*, (2018)
65. Prasser, F., Bild, R., Eicher, J., Spengler, H., Kohlmayer, F., Kuhn, K.A.: Lightning: Utility-driven anonymization of high-dimensional data. *Trans. Data Privacy* **9**(2), 161–185 (2016)
66. Prasser F. and Kohlmayer F.: Putting statistical disclosure control into practice: The ARX data anonymization tool. In *Medical Data Privacy Handbook*, (2015)
67. Ruggles S.: Implications of differential privacy for census bureau data and scientific research. Technical Report 2018-6, Data without boundaries. Task Force on Differential Privacy for Census Data. Institute for Social Research and Data Innovation (ISRDI), University of Minnesota, (2018)
68. Samarati, P.: Protecting respondents identities in microdata release. *IEEE Trans. Knowl. Data Eng.* **13**(6), 1010–1027 (2001)
69. Samarati P. and Sweeney L.: Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Technical Report SRI-CSL-98-04, SRI International, (1998)
70. Shokri R. and Shmatikov V.: Privacy-preserving deep learning. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15, page 1310–1321, New York, NY, USA, (2015). Association for Computing Machinery
71. Silva M.J., Rijo P., and Francisco A.: Evaluating the impact of anonymization on large interaction network datasets. In Proceedings of the First International Workshop on Privacy and Security of Big Data, PSBD '14, page 3–10, New York, NY, USA, (2014). Association for Computing Machinery
72. Sim, J.J., Chan, F.M., Chen, S., Meng Tan, B.H., Mi Aung, K.M.: Achieving GWAS with homomorphic encryption. *BMC Med. Genom.* **13**(7), 90 (2020)
73. Stadler T., Oprisanu B., and Troncoso C.: Synthetic data – anonymisation groundhog day, (2022)
74. Sun, G., Cong, Y., Dong, J., Wang, Q., Liu, J.: Data poisoning attacks on federated machine learning, (2020)
75. Templ, M.: Providing data with high utility and no disclosure risk for the public and researchers: An evaluation by advanced statistical disclosure risk methods. *Austrian J. Stat.* **43**(4), 247–254 (2014)
76. Templ, M.: Quality indicators for statistical disclosure methods: A case study on the structure of earnings survey. *J. Offic. Stat.* **31**(4), 737–761 (2015)
77. Templ, M.: Statistical disclosure control for microdata: methods and applications in R. Springer International Publishing, Cham, Switzerland (2017)
78. Templ M. and Alfons A.: Disclosure risk of synthetic population data with application in the case of EU-SILC. In *Privacy in Statistical Databases., Lecture Notes in Computer Science*, pages 174–186. Springer, (2010)
79. Templ, M., Filzmoser, P.: Simulation and quality of a synthetic close-to-reality employer-employee population. *J. Appl. Stat.* **41**(5), 1053–1072 (2014)
80. Templ, M., Kowarik, A., Meindl, B.: Statistical disclosure control for micro-data using the R package sdcMicro. *J. Stat. Softw.* **67**(4), 1–36 (2015)
81. Templ M. and Meindl B.: Robustification of microdata masking methods and the comparison with existing methods. *Privacy in Statistical Databases. Lecture Notes in Computer Science.* Springer, 5262:177–189, (2008)
82. Templ, M., Meindl, B., Kowarik, A., Dupriez, O.: Simulation of synthetic complex data: The R package simPop. *J. Stat. Softw.* **79**(10), 1–38 (2017)
83. Thompson, G., Broadfoot, S., Elazar, D.: Methodology for the automatic confidentialisation of statistical outputs from remote servers at the Australian Bureau of Statistics. In Joint UNECE/Eurostat work session on statistical data confidentiality, Ottawa, Canada (2013)
84. Wang S., Wang X., Zhao P., Wen W., Kaeli D., Chin, P. and Lin X.: Defensive dropout for hardening deep neural networks under adversarial attacks. In Proceedings of the International Conference on Computer-Aided Design, ICCAD '18, New York, NY, USA, 2018. Association for Computing Machinery
85. Ward, K., Lin, D., Madria, S.: A parallel algorithm for anonymizing large-scale trajectory data. *ACM/IMS Trans. Data Sci.* **1**(1), 1–26 (2020)
86. Wei, K., Li, J., Ding, M., Ma, C., Yang, H.H., Farokhi, F., Jin, S., Quek, T.Q.S., Poor, H.V.: Federated learning with differential

- privacy: Algorithms and performance analysis. *IEEE Trans. Inf. Forensics Secur.* **15**, 3454–3469 (2020)
87. Wirth H., Rockmann U., Müller D., Goebel J., and Mika T.: Remote access to data from official statistics agencies and social security agencies. Technical report, Rat für Sozial- und Wirtschaftsdaten (RatSWD), Berlin, 5(6): 1-41, (2019)
88. Yang Q., Fan L., and Yu H.: Federated Learning. Privacy and Incentive. *Lecture Notes in Computer Science*. Springer Nature Switzerland AG, Cham, Switzerland, (2020)
89. Zamarripa C. and Williams B.: Census Bureau announces traditional redistricting data not recommended for use this decade, (2021)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.