arXiv:2212.00484v3 [cs.LG] 23 Apr 2024

Differentially-Private Data Synthetisation for Efficient Re-Identification Risk Control

Tânia Carvalho[1*], Nuno Moniz[2], Luís Antunes[1,3], Nitesh Chawla[2]

[1]Faculty of Sciences, University of Porto, Portugal.
[2]Lucy Family Institute for Data & Society, University of Notre Dame, Indiana, USA.
[3]TekPrivacy, Porto, Portugal.

*Corresponding author(s). E-mail(s): tania.carvalho@fc.up.pt;
Contributing authors: nuno.moniz@nd.edu; lfa@fc.up.pt;
nchawla@nd.edu;

Abstract

Protecting user data privacy can be achieved via many methods, from statistical transformations to generative models. However, all of them have critical drawbacks. For example, creating a transformed data set using traditional techniques is highly time-consuming. Also, recent deep learning-based solutions require significant computational resources in addition to long training phases, and differentially private-based solutions may undermine data utility. In this paper, we propose ϵ-PrivateSMOTE, a technique designed for safeguarding against re-identification and linkage attacks, particularly addressing cases with a high re-identification risk. Our proposal combines synthetic data generation via noise-induced interpolation with differential privacy principles to obfuscate high-risk cases. We demonstrate how ϵ-PrivateSMOTE is capable of achieving competitive results in privacy risk and better predictive performance when compared to multiple traditional and state-of-the-art privacy-preservation methods, including generative adversarial networks, variational autoencoders, and differential privacy baselines. We also show how our method improves time requirements by at least a factor of 9 and is a resource-efficient solution that ensures high performance without specialised hardware.

Keywords: Data Privacy, Synthetic Data, Deep Learning, Differential Privacy and Resource-Efficiency.

1 Introduction

Most traditional data de-identification tools use Privacy-Preserving Techniques (PPT) such as generalization (recoding values into broader categories) or suppression (replacing actual values with null values) to preserve the truthfulness of the data [1]. Recent attention in data privacy has focused on data synthesis wherein deep learning models generate artificial data that reflect the statistical properties and structure of the original data set [2]. Although synthetic data generation is often perceived as a means to protect personal information fully, many questions have arisen on memorisation of deep learning models, especially outliers, which typically represent sensitive data points [3]. To address these privacy concerns, Differential Privacy (DP) [4] has been introduced in the synthetic data generation [5].

These approaches are designed to protect an individual's privacy by preventing intruder attacks such as re-identification, also known as singling out (assigning a particular identity to an isolated record) or linkage attacks (linking at least two records concerning the same individual) [6]. However, a significant drawback is that the transformation of all instances may negatively affect data utility. As intruders typically target instances with the highest risk of exposing sensitive information [7], transforming all instances may inadvertently compromise the overall usefulness and effectiveness of the data. Most importantly, current privacy-preserving approaches still exhibit significant computational and time overheads [8].

In this paper, we propose ϵ-PrivateSMOTE, a new strategy for privacy-preserving tabular data sharing using synthetic data generation. Our focus lies on addressing the re-identification risk by strategically replacing high-risk cases with synthetic and similar data points. To our knowledge, interpolation methods have not yet been thoroughly explored in favour of privacy. In this context, we encounter exciting opportunities to explore existing strategies. Our contribution involves leveraging SMOTE-inspired heuristics [9], a sophisticated interpolation technique based on nearest neighbours algorithms, that can operate effectively even in scenarios with limited data quantities [10]. Also, it enables targeted enhancements of privacy assurances by focusing on the highest-risk cases. Our approach takes advantage of the stochastic indistinguishability of DP, which aims to protect against several privacy attacks (e.g., differencing, linkage, and reconstruction attacks) [11] using a privacy budget parameterized by ϵ. The incorporation of the Laplace mechanism, a well-established differentially private mechanism [12], ensures statistically indistinguishable analyses on similar datasets.

While current privacy-preserving solutions for tabular data sharing were designed to modify all instances, ϵ-PrivateSMOTE distinguishes for its ability to strategically enhance privacy assurances in a targeted way offering a promising avenue for secure data sharing and simultaneously maximising data utility. The significance and relevance of our proposal can be summarised as follows:

- ϵ-PrivateSMOTE is easy to implement and configure when compared to traditional privacy solutions using PPT (data granularity reduction);
- ϵ-PrivateSMOTE outperforms deep learning- and differentially private-based solutions in predictive performance;

- ϵ-PrivateSMOTE produces more new data variants with fewer resources than deep learning- and differentially private-based approaches, both in time and computation;
- Contrary to any other privacy-preserving solution tested, ϵ-PrivateSMOTE offers a notable balance between privacy and predictive performance.

The remainder of the paper is organized as follows. Section 2 overviews related work on strategies for privacy-preserving data sharing and sampling strategies. We present ϵ-PrivateSMOTE approach in Section 3. An extensive experimental evaluation is described in Section 4. A discussion of the results obtained is introduced in Section 5, and conclusions are provided in Section 6.

2 Related Work

In the de-identification process, it is fundamental to make assumptions about the intruder's background knowledge since it is unknown what information the person may possess. Such assumptions are based on the selection of quasi-identifiers (QI): attributes that, when combined, can lead to identification, e.g., date of birth, gender, and profession. To limit the disclosure risk of personal information, it is imperative to apply appropriate transformations on such QI [1, 13].

The transformed data set is then evaluated for its re-identification risk. Two standard measures are k-anonymity and record linkage. K-anonymity [14] indicates how many k occurrences occur in the data set w.r.t a specific combination of QI values. An intruder can single out a record when $k = 1$. On the other hand, record linkage (or linkability) [15] aims to measure the ability of re-identification by linking two records using similarity functions. Privacy is also a focus during the machine learning training phase; Membership Inference (MI) risk aims to identify whether a given data point is included in the training set [16].

2.1 Strategies for De-identification

De-identified data is achieved via distinct approaches. We analyse the three most common: traditional, deep learning- and differentially private-based solutions.

Traditional techniques uses PPT to achieve a certain level of privacy. These may be validated using various tools, such as k-anonymity [14] or Differential Privacy (DP) [4]. To achieve k-anonymity, each individual represented in a data set cannot be distinguished from at least $k-1$ other individuals. Data is typically reduced, as is its granularity, which is likely to impact the predictive performance of machine learning models [17]. On the other hand, to achieve DP, the statistical results of the data set are not supposed to be influenced by an individual's contribution. Although DP has gained prominence recently, it is often applied to queries or during the training process in machine learning applications. Nevertheless, a few approaches have emerged for releasing differentially private data sets; for instance, (k, ϵ)-anonymity [18] uses DP on a given attribute. However, their proposal must first guarantee a certain level of k-anonymity and then apply the noise to each group w.r.t QI (equivalence classes). Moreover, this technique is applied to a single attribute. Recently, Muralidhar et al. [19] showed that a fixed ϵ does not guarantee a certain level of confidentiality or

utility on differentially-private data created using noise-added covariances approach nor using added noise to the cumulative distribution function.

Deep Learning-based solutions such as CTGAN [20] emerged as an alternative to traditional techniques. A comparison of generative methods was performed with particular attention to privacy concerns. This comparison shows that medical data generated by GANs does not model outliers well and that such models are also vulnerable to MI attacks [21].

Differentially Private-based solutions incorporate DP into the process of synthetic data generation for high privacy guarantees. For instance, DPGAN [22], DP-CGAN [5] and PATE-GAN [23] are differentially private GAN models that add noise during the training process of the discriminator. DP has also been added to the tree-based machine learning models to generate synthetic data [24]. Besides, PrivBayes [25] uses DP in a Bayesian Network focusing on marginal distributions. Recent examples on marginals include PrivSyn [26] and PrivMRF [27].

2.2 Sampling Methods

Although privacy-preserving data sharing focuses on data transformations or generative solutions, we suggest that sampling methods, commonly used in imbalanced learning tasks, could be a valuable tool to tackle this challenge.

Popular data pre-processing (i.e., sampling) methods in imbalanced classification tasks include random under- (RUS) and over-sampling (ROS). The former randomly selects cases of the majority class for removal while the latter works by replicating cases of the minority class. The most popular is SMOTE (Synthetic Minority Over-sampling Technique) [9], an over-sampling strategy that focuses on the synthetic generation of rare cases by interpolating existing ones, based on the nearest neighbour algorithm. BorderlineSMOTE [28] is a variation of SMOTE that focuses on oversampling minority examples near class borderlines. In addition to these approaches, several existing sampling methods have been reviewed to address class imbalance problems [29–31].

2.3 Summary

The main challenge in the de-identification process using PPT is to achieve an appropriate level of privacy and its utility, as it involves several steps until the balance between the two metrics is found. Furthermore, a de-identification level is stipulated, e.g., in the case of k-anonymity, an algorithm ensures that the data respects this k level. However, if the algorithm cannot use more generalisation, much information is lost as suppression is applied. In addition, data points that initially would not have a high risk of re-identification are altered, affecting the utility. On the other hand, adding noise through DP can potentially result in data with an increased level of uncertainty, which can be difficult for data analysts to work with and can lead to inaccurate machine learning models.

Finally, a significant drawback of generative models using deep learning solutions is that statistical properties may not be well captured with smaller data sets. Moreover, several implementations of such tools have been shown to violate privacy guarantees, exposing many records to inference attacks and failing to retain data utility [32].

Additionally, a major challenge of deep learning-based solutions is their complexity, as they are computationally costly in memory and require significant execution time [8].

Despite increased efforts to improve de-identification strategies, the above limitations motivate us to explore efficient privacy alternatives. An underinvested area is the intersection with sampling methods used in imbalanced learning tasks, such as SMOTE [9]. In contrast to deep learning-based solutions, the heuristics of interpolation methods are simple, allowing synthetic data to be generated with less effort. We propose to explore such an approach for data privacy.

3 Privacy-Preserving Data Synthetisation

In this section, we present a novel approach for privacy-preservation leveraging the indistinguishability of DP and interpolation methods for synthetic data generation, focusing on cases with a high re-identification risk.

Consider a data set $T = \{t_1, ..., t_n\}$, where t_i corresponds to a tuple of attribute values for an individual's record. Let $V = \{v_1, ..., v_m\}$ be the set of m attributes. A set of quasi-identifiers (QI) consists of attributes values that could be known to the intruder for a given individual where $QI \in V$ and is denoted as $QI = \{v_{j_1}, ..., v_{j_k}\}$. For example, $QI = \{v_{j_1}, v_{j_2}, v_{j_3}\}$ might represent an individual's gender, age, and zip code. A tuple of selected QI is denoted as $t_i^{QI} = \{t_{i,v_j} : v_j \in QI\}$.

Definition 1 (**Re-identification**). *Assume an intruder knows that an individual with the tuple t_r^{QI} is in T. If no other tuple shares identical information, then t_r^{QI} is singled out and the intruder can uniquely identify the individual.*

Additionally, an intruder could infer that there is exactly one person in T who has $t_u^{QI}\langle$gender: male, age: 75, zip: 46711$\rangle$. While singling out this particular individual may not constitute re-identification, the ability to isolate an individual may lead to other privacy attacks [15]. Therefore, it is essential to increase safeguarding measures. For example, if another t_p^{QI} shares the same set of information as t_r^{QI}, there is a 50% probability of re-identification. This reinforces the importance of carefully examining other tuples that may share similar QI. According to the k-anonymity definition, to ensure the protection of cases with a 50% risk of re-identification, the protected version of T must be at least 3-anonymous. This implies that each group w.r.t. QI contains at least three indistinguishable individuals, reducing the probability of re-identification to 33% for each individual. Such criteria can be adjusted by the k parameter according to the use case under analysis. In our study, we focus on protecting single outs and cases with a 50% risk of re-identification. The k-anonymity basis is used to detect and mitigate the risk associated with these highest-risk instances.

Definition 2 (**Highest-risk selection**). *Given T and $T[QI]$, T' corresponds to the highest-risk cases table, where each sequence of values in a tuple t_{i,v_j} occurs at a maximum two times in $T[QI]$.*

As such, instead of generating synthetic instances via interpolation of cases from the minority class, we focus on the highest-risk instances. For such instances, we generate numerical synthetic cases based on randomly weighted interpolation of nearest neighbours, akin to the process in SMOTE [9], without any class discrimination, i.e., focusing solely on case proximity. For the interpolation of nominal attributes, we employ a random selection process where a category is chosen from the neighbours of each instance. However, in cases where only one category is available for the nearest neighbours, the selection is made based on the categories observed in the entire sample to prevent the new value from being equal to the original.

At this point, we still face a major drawback: the inability to deal with lower data variability. In some cases, the nearest neighbour chosen for synthesis may be equal to the instance to be synthesised (e.g. *Age* attribute), generating the same value as the original, which is critical for privacy. To overcome such a drawback, we propose to enforce privacy by forcing a difference between the two values using the standard deviation (*std*), i.e. multiplying the random uniform value [0,1] by *std* or $-std$.

Nevertheless, we enhance privacy guarantees through DP by replacing the standard uniform random selection with a Laplace-based mechanism that introduces noise from the Laplace distribution. A randomised algorithm M satisfies ϵ-differential privacy [4] if for all data sets D_1 and D_2 differing on at most one element, which differ in a single individual, and for all subsets of possible outcomes S ($S \subseteq Range(M)$),

$$Pr[M(D_1) \in S] \leq e^{\epsilon} \times Pr[M(D_2) \in S] \tag{1}$$

Generally, uniform noise is (0,δ)-differential private [33], where $\epsilon = 0$ means that no information is revealed about any individual's data, also known as "pure" differential privacy; in opposite, data utility is destroyed. The *epsilon* parameter indicates the amount of added noise, often referred to as *privacy budget* or *cost* and δ represents the probability of violating the privacy cost [4]. Therefore, we strengthen privacy protection by leveraging the Laplace mechanism, which is differentially private [12]. In our context, this mechanism is described as follows.

Definition 3 (**Laplace mechanism**). *For each value ϑ in a tuple t_{i,v_j} of $T'[QI]$, the Laplace mechanism A is given by $A = \vartheta + Laplace(1/\epsilon)$.*

We call this proposed solution ϵ-PrivateSMOTE. For clarity and further reproducibility, we provide a detailed description of the execution in Algorithm 1. The Algorithm is organised into distinct steps, which are outlined below.

1. nearest neighbours are computed over the entire sample (line 10);
2. k highest-risk cases are determined based on the selected QI (line 11);
3. *knn* nearest neighbours are found for each highest-risk instance (line 13);
4. DP is incorporated instead of using uniform noise as a privacy protection mechanism (lines 18-19);
5. for stronger privacy guarantees, the synthesis of a new case is enforced when a specific example is equal to the chosen nearest neighbour (line 19 #2) and,

6. nominal attributes are interpolated by randomising the selection based on the uniqueness of a case's neighbours. If there is no randomness in the former case, the uniqueness of the entire sample is considered (lines 20-21).

Algorithm 1 ϵ-PrivateSMOTE

Input: Table T with set of m attributes and n instances, amount of new cases N, number of nearest neighbours knn, amount of noise ϵ, group size k and QI set.
Output: $N * n$ synthetic samples

1: **function** HighestRisk(T, k, QI)
2: eq = compute the occurrences of each group in T w.r.t QI
3: T' = filter rows where eq is less than k
4: **return** T'
5: **end function**
6: **procedure** ϵ-PrivateSMOTE(T, N, knn, ϵ, k, QI)
7: y = class examples $(y_1, ..., y_n)$ in T
8: synthetic_samples = []
9: ζ = one-hot-encoding and standardisation of T
10: nn = fitted nearest neighbours using ζ
11: T' = HighestRisk(T, k, QI)
12: **for** $i \leftarrow 1$ to n instances **do**
13: $nnarray$ = knn nearest neighbours from nn if $i \in$ indexes of T'
14: **while** $N \neq 0$ (populate N times) **do**
15: **for** $attr \leftarrow 1$ to m **do**
16: a = instance in position i
17: b = random instance from $nnaray$
18: **if** $T[attr]$is numerical **then**
19: new_instance = $a + Laplace(0, 1/\epsilon) * (b - a)$ **if** $a \neq b$ #1
else $a + Laplace(0, 1/\epsilon) * rand[std, -std]$ #2
20: **else if** $T[attr]$is nominal **then**
21: new_instance = random category from $T[nnarray]$ unique values
if there are two or more categories,
otherwise from $T[attr]$ unique values
22: **end if**
23: add new_instance to synthetic_samples
24: **end for**
25: add the intact target value $y[i]$ to synthetic_samples
26: **end while**
27: **end for**
28: **end procedure**

3.1 Threat model

Although synthetic data is at the forefront of data de-identification methods, unresolved privacy concerns still require attention and mitigation. Related literature on differentially private-based solutions essentially focuses on presenting privacy assurances concerning different levels of DP [22, 23, 26] and often lacks comprehensive discussions on associated threat models. In addition, deep learning- and differentially private-based models are generally evaluated against MI attacks [34]. However, this type of attack is conventionally applied to learning models [34], specifically during the training phase, making MI attacks less applicable to traditional techniques and

ϵ-PrivateSMOTE. On the one hand, in a k-anonymised data set, each record is deliberately made to resemble multiple other records, making it challenging to determine the presence of a particular individual's data point; on the other hand, as the core of our proposal lies in the heuristics of the interpolation method, the augmentation and DP noise makes it challenging to determine whether a data point corresponds to a specific individual in the original data set. Therefore, membership attacks are not as informative for assessing the privacy of data sets using these approaches.

Linkability [35] is a common privacy attack that involves a concerted effort by an unauthorised entity to establish a correlation between two records [7]. The intruder attempts to link and potentially expose sensitive information in the targeted records, with outliers being particularly vulnerable. The target records of the attack are a collection of N_A original records randomly drawn from the original data set. Suppose an intruder has some knowledge of the targets and access to two data tables, a synthetic data table (T_A) and another obtained from public information (T_B), both of which contain these targets. The intruder's goal is to correctly match records between T_A and T_B given a set of QI common to both data sets.

Definition 4 (**Linkability [35]**). *For each instance in T_A, an intruder finds the k closest synthetic instance in T_B. The resulting indices are $I^A = (I_i^A, ..., I_{N_A}^A)$, where each I_i^A is the set of k indexes of the synthetic records that are nearest neighbours of the i^{th} target in the subspace of feature set.*

For each of the N_A targets, it is checked whether both nearest neighbours sets share the same synthetic data record. In a positive case, an attacker can associate previous unconnected information about a target individual that is present in the original data set. Each successfully established link is scored as a success, where the outcome (O) is determined by:

$$O_i(I_i^A, I_i^B) = \begin{cases} 1, & \text{if } I_i^A \cap I_i^B \neq 0 \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

4 Experimental Evaluation

In this section, we provide a thorough experimental evaluation. Our goal is to answer the following research questions. How does ϵ-PrivateSMOTE compare to competing methods in terms of predictive performance and privacy risk, namely linkability (**RQ1**), how do optimisation paths (prioritisation of a single vector) affect the results of both vectors (**RQ2**), and how does ϵ-PrivateSMOTE behave in terms of time/computational resources complexity (**RQ3**)? In the following, we present our methodology and experimental evaluation methods. Then, we briefly describe the data we used. Finally, we present the results of the experiments.

4.1 Methodology and Methods

We divide our methodology into three phases: *i)* data transformation, *ii)* data utility concerning general distributions and predictive performance evaluation, and *iii)* data

privacy. Regarding data transformation (i), we apply five approaches to original data sets. For each set, we generate several variants employing the following methodologies.

Traditional Techniques are applied according to the characteristics of the QI, i.e., the type of attribute values. We assume that any attribute can be a QI as we do not know what information an intruder might possess. Consequently, we collect five different QI sets, each consisting of a random selection of 40% of the attributes, that represent potential background knowledge for an intruder. For each set, we create new transformed variants. We use the ARX tool [36], a benchmark solution. We apply generalisation with specific hierarchies using specific intervals. We also select the outlier suppression option to achieve the desired privacy level. From the solution space, we select the three best solutions w.r.t utility (information loss) as provided by the tool.

Sampling methods are applied to all attributes and instances. We use RUS, SMOTE, and BorderlineSMOTE as implemented in the Python library *imblearn* [37]. Regarding parameterization, we use $ratio \in \{0.25, 0.5, 0.75, 1\}$ for RUS, and $ratio \in \{0.5, 0.75, 1\}$ and $knn \in \{1, 3, 5\}$ for SMOTE and Borderline.

Deep Learning-based solutions are also applied to all attributes and instances. We use the *SDV* library [38], experiment with Copula GAN, TVAE, and CTGAN and the parameters: $epochs \in \{100, 200\}$ and $batch_size \in \{50, 100\}$.

Differentially Private-based solutions are still in an improvement phase[1]. We use *synthcity* [39] that provides a benchmark for synthetic data generators. Our experimentation involved the use of DPGAN and PATE-GAN with the same $epochs \in \{100, 200\}$, $batch_size \in \{50, 100\}$ and $\epsilon \in \{0.1, 0.5, 1.0, 5.0\}$.

ϵ-PrivateSMOTE is applied to the highest-risk cases given the set of QI (the same as for traditional techniques). We use $N \in \{1, 2, 3\}$ (number of times a case is repeated) and $knn \in \{1, 3, 5\}$. Also, we need to specify the amount of added noise. We experiment with $\epsilon \in \{0.1, 0.5, 1.0, 5.0, 10.0\}$. Previous solutions do not include $\epsilon = 10.0$ as it fails to generate several data variants.

All privacy-preserving data variants are evaluated for data utility (ii). For general utility, we use SDMetrics from SDV [38]. Specifically, we use Range Coverage, Boundary Adherence, Statistical Similarity and Correlation to check the extent to which each data variant differs from the original data in terms of general distributions.

We also evaluate the predictive performance using the following methodology. We apply four classification algorithms (Random Forest [40], XGBoost [41], Logistic Regression [42] and Neural Network [43]) using *Scikit-learn* [44], to test all the privacy strategies. Final models for each algorithm are chosen based on a 2*5-fold cross-validation estimation of the evaluation scores for models based on a grid search method. We also obtain performance results on test data – 20% of the original data set. Table 1 summarises this information.

The effectiveness of predictive performance will be evaluated using AUC (Area Under the ROC Curve) [45]. We also apply statistical tests using Bayes Sign Test [46] to evaluate the significance of our experimental results. We use the percentage difference between each pair of solutions as $\frac{R_a - R_b}{R_b} * 100$ where R_a is the solution under

[1]We attempted to use PrivBayes; however, we encountered challenges in generating data variants. The program exhibited long execution times, using only the CPU.

Algorithm	Parameters
Random Forest	$n_estimators \in \{100, 250, 500\}$ $max_depth \in \{4, 7, 10\}$
Boosting	$n_estimators \in \{100, 250, 500\}$ $max_depth \in \{4, 7, 10\}$ $learning_rate \in \{0.1, 0.01\}$
Logistic Regression	$C \in \{0.001, 1, 10000\}$ $max_iter \in \{10e^5, 10e^6\}$
Neural Network	$hidden_layer_sizes \in \{[n_feat], [n_feat/2], [n_feat * 2/3],$ $[n_feat, n_feat/2], [n_feat, n_feat * 2/3], [n_feat/2, n_feat * 2/3],$ $[n_feat, n_feat/2, n_feat * 2/3]\}$ $alpha \in \{0.005, 0.001, 0.0001\}$ $max_iter \in \{10000, 100000\}$

Table 1 Learning algorithms considered in the experimental evaluation and respective hyperparameter grid. Each line of parameter *hidden layer sizes* (Neural Networks) represents an increasing number of layers (1–3).

comparison and R_b is the baseline solution. In this context, ROPE (Region of Practical Equivalence) [47] is used to specify an interval ([-1%, 1%]) where the percentage difference is considered equivalent to the null value. Thus, *1)* if the percentage difference of the specified metric between solutions a and b (baseline) is greater than 1%, solution a outperforms b (win); *2)* if the percentage difference is within the specified range, they are of practical equivalence (draw); and *3)* if the percentage difference is less than -1%, the baseline outperforms solution a (lose).

Each resulting data variant is also evaluated regarding privacy risk (*iii*). We evaluate privacy through linkability risk by comparing each variant against the original data set. The comparison space includes only the specified set of QI [2]. We use *Anonymeter* library [15] for linkability assessment. For this evaluation, it is required a control data set which we use the test data set (20% of original data), and k that indicates the search space for nearest neighbours; we use $k = 10$.

For conciseness, our experimental methodology is illustrated in Figure 1.

4.2 Data

To provide an extensive experimental study, we collected 15 diverse classification data sets from the public and open data repository OpenML [48]. These datasets span various domains, ensuring different representations of real-world scenarios. Common characteristics include: *i)* binary class target, *ii)* number of instances between 1.043 and 15.545, and *iii)* number of features greater than five and less than 103. The average number of instances of the retrieved data sets is 4.916. Also, there is an average of 28 numerical and two nominal attributes. For transparency and reproducibility, all datasets used in our experimental evaluation are publicly available [3].

4.3 Results

The first set of results from the experimental evaluation focuses on an overall assessment of data utility and privacy risk for ϵ-PrivateSMOTE. Concerning predictive

[2] For comparison purposes, we use the same set of QI for sampling and deep learning- and differentially private-based solutions.

[3] https://www.kaggle.com/datasets/up201204722/original

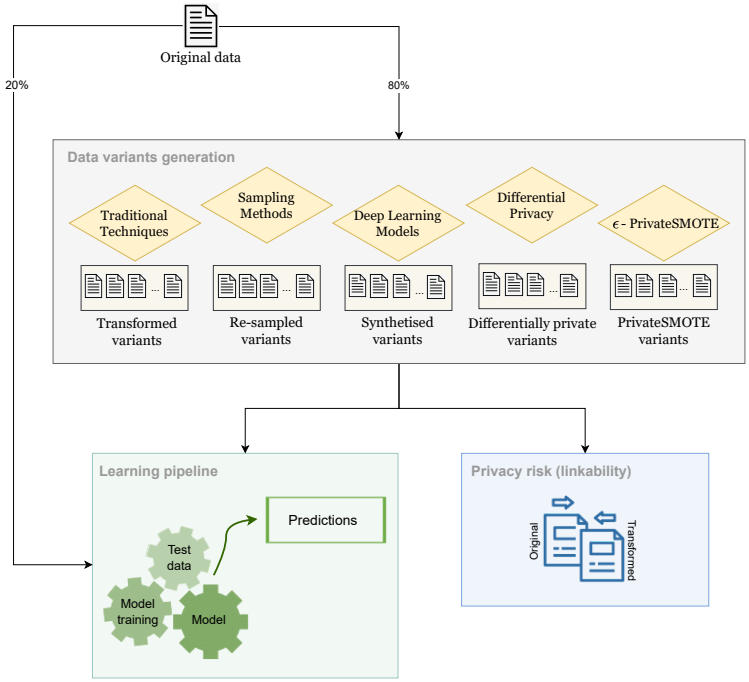

Fig. 1 Methodology of the experimental evaluation.

performance, results refer to the out-of-sample performance of the best models for each data variant, estimated via cross-validation. Then, we calculate the percentage difference between the best-estimated model for each data variant solution and the baseline – best solution for each original data set.

Figure 2 illustrates the correlation between the linkability and the percentage difference of predictive performance for each ϵ across all produced data variants.

Overall, we observe a central tendency around zero for both the percentage difference of predictive performance and linkability. Nevertheless, higher ϵ values can provide solutions with high predictive performance and lower privacy risk. Such a result suggests that ϵ-PrivateSMOTE is able to provide solutions that maximise both vectors, in particular, for $\epsilon \geq 0.5$. Although it is well-known that lower ϵ values, i.e. higher noise, result in higher privacy [4] and consequently may negatively affect the predictive performance, we observe that 0.1-PrivateSMOTE has higher linkability and yet can produce several solutions with equivalent predictive performance. This outcome is potentially ascribed to the larger values at the tails of the distribution, which may lead to an increase in the overall distribution, as we can evidence in Figure 3.

This Figure shows different utility measures concerning all ϵ values for all data variants created with ϵ-PrivateSMOTE. We conclude that lower ϵ values present better Range Coverage meaning that the data variants cover the full range of each numeric column. On the other hand, lower ϵ fails to adhere to the real boundaries. Such conditions result in more spread values at the tails, facilitating linkability since these values

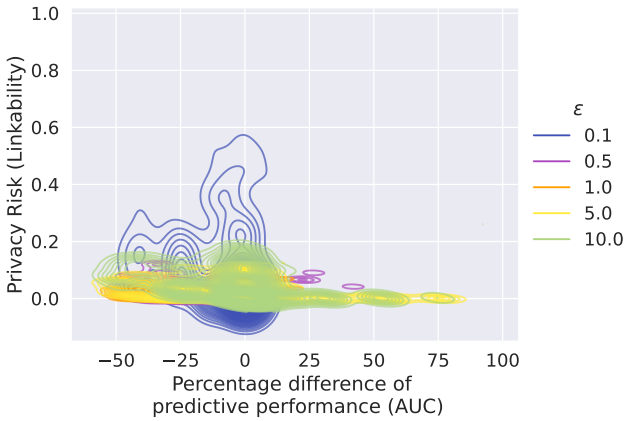

Fig. 2 Relationship between privacy risk and predictive performance of all data variants produced with ϵ-PrivateSMOTE for each ϵ.

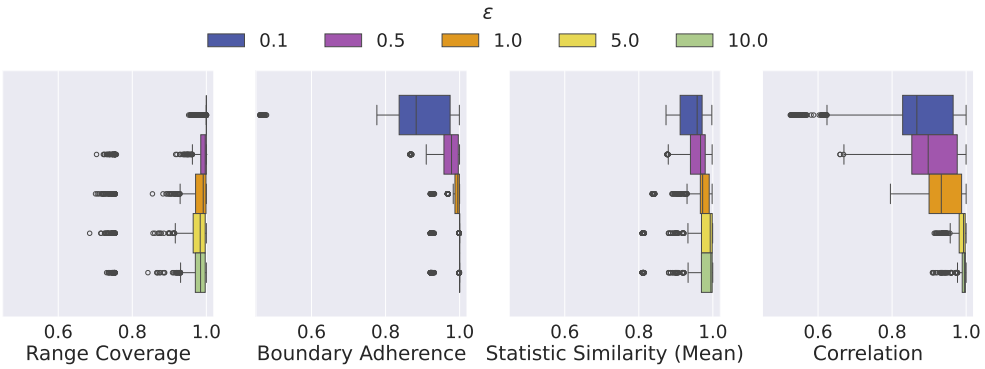

Fig. 3 Data utility measures for each ϵ across all data variants using ϵ-PrivateSMOTE.

are rarer. Also, it provides lower Statistic Similarity and Correlation compared to high ϵ values. In general, we observe better utility for higher ϵ values.

Optimisation Paths

In the subsequent analysis, we present a comparative evaluation of ϵ-PrivateSMOTE against state-of-the-art privacy-preservation approaches, explicitly examining the optimisation paths. These paths focus on prioritising a single vector and presenting the corresponding values of the other vector. Thus, Figure 4 demonstrates the effects of selecting the best solutions in predictive performance and the corresponding values for linkability, and the reverse path. Note that PPT refers to traditional techniques.

Focusing on prioritising predictive performance and each approach individually, PPT present a high predictive performance but also a higher linkability risk. In our experimental setup, we use lower generalization levels, removal of outliers, and removal of features with too many unique values due to suppression. These transformations

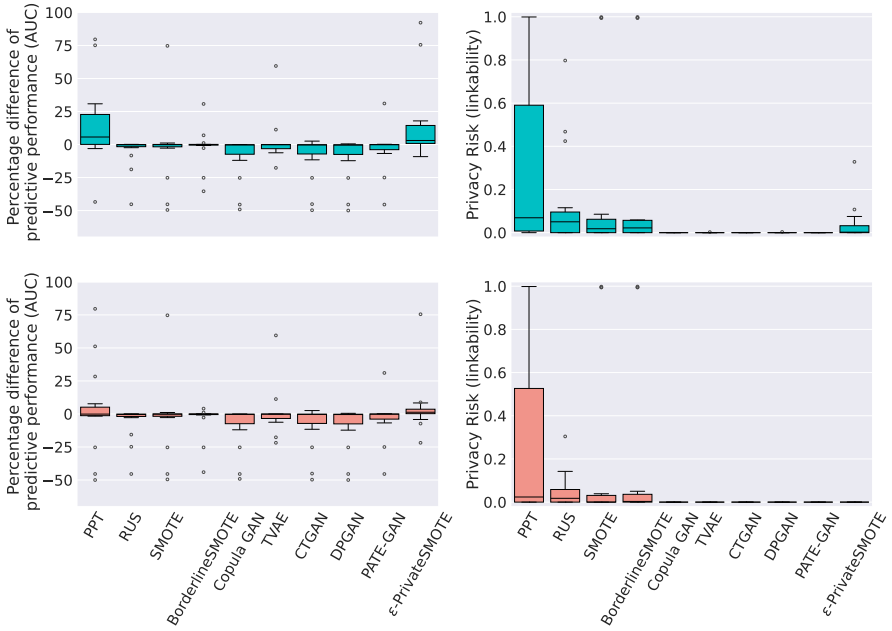

Fig. 4 Best predictive performance results and corresponding privacy risk (blue) and best privacy risk results and corresponding predictive performance results for each transformation technique (red).

allow for obtaining approximate predictive performance values with the original data but could potentiate privacy leaks. Regarding sampling methods, the predictive performance is very similar to the original data. Since most of the original data sets are pretty balanced, this may explain the predictive performance results. Also, these three show a certain degree of linkage risk, mainly because they retain instances in their original form. The scenario is a bit different for deep learning- and differentially private-based solutions. Despite good privacy values, these techniques show a decline in predictive performance. Finally, ϵ-PrivateSMOTE provides a median very similar to PPT with a lower linkability risk, near zero.

Concerning the reverse path, we observe a decrease in linkability risk, especially for ϵ-PrivateSMOTE, which is equivalent to deep learning- and differentially private-based solutions. However, this improvement comes at the cost of predictive performance. This consequence is also demonstrated with PPT.

In a general overview, ϵ-PrivateSMOTE presents a median with practical equivalence to the baseline while maintaining a reduced linkability risk. Since it does not alter privacy-secure instances, this may help explain its performance. Also, as we produce twice/triple new cases (N parameter), this can worsen the linkage ability because it creates more uncertainty about the cases.

5 Discussion

It is well known that the maximisation of data privacy is at the expense of data utility [1]. Thus, it is likely that the higher the privacy risk, the higher the predictive performance. However, ϵ-PrivateSMOTE has proven its effectiveness in providing solutions that guarantee lower losses for predictive performance while maintaining lower linkability (Figure 2). Comparing the efficacy in terms of predictive performance and privacy risk of ϵ-PrivateSMOTE with the state-of-the-art approaches (**RQ1**), the results show that our proposal can produce competitive results with the baseline for both scenarios. ϵ-PrivateSMOTE stands out from the remaining approaches (Figure 4), especially compared to PPT, which provides a competitive result in terms of predictive performance. Although PPT preserve non-QI attributes in their original form, ϵ-PrivateSMOTE maintains a similar predictive performance while protecting all attributes with a lower linkability risk.

Concerning the optimisation paths, we observe better linkability results achieved by prioritising privacy. This outcome suggests that it is beneficial to prioritise privacy first for ϵ-PrivateSMOTE. On the other hand, the privacy improvements for PPT may not compensate for the loss in predictive performance. For the remaining tested solutions, it is practically indifferent which vector is prioritised (**RQ2**).

To provide a deeper understanding of the difference between ϵ-PrivateSMOTE and the remaining approaches, Figure 5 provides a comparison that reports the results of the Bayes Sign Test for predictive performance to infer the statistical significance of the paired differences in the outcome of the approaches. With this aim, we identify the *oracle* model and use it as the baseline regarding privacy prioritisation. The oracle identification relies on the assumption that a strategy is able to consistently estimate the best possible model in out-of-sample performance from the explored hyperparameter configurations. Figure 5 illustrates the best estimated (in cross-validation settings) model for each transformation technique and the *oracle* for each data set.

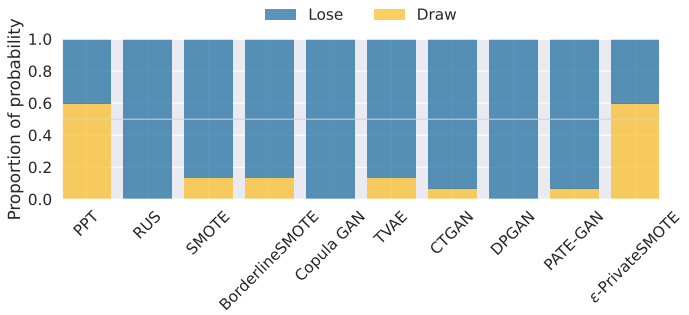

Fig. 5 Comparison between the best-estimated hyperparameter configuration per transformation technique and the oracle configuration. Illustrates the proportion of probability for each candidate solution drawing or losing significantly against the oracle according to the Bayes Sign Test for predictive performance.

Results show that the probability of practical equivalence with the oracle is 60% for both PPT and ϵ-PrivateSMOTE. Even though privacy is prioritised, these strategies demonstrate the ability to ensure favourable predictive performance results. However, in contrast to the results observed with ϵ-PrivateSMOTE, the higher linkability risk previously identified with the use of PPT is noteworthy. The remaining transformation techniques present a probability of practical equivalence to the oracle lower than 15%.

Additionally, we provide in Figure 6 the comparison of all transformation techniques in terms of general data utility for the privacy optimisation path. We transformed the ranges of PPT data variants by selecting the lowest bound to maintain the same data type as the original data.

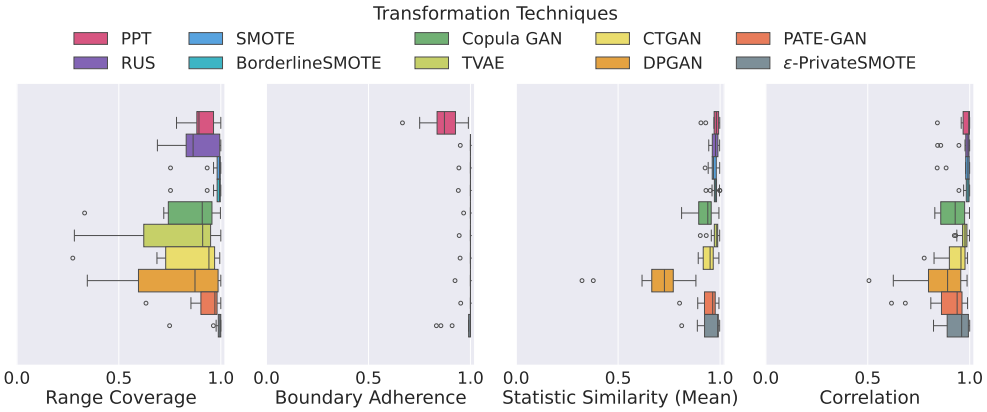

Fig. 6 Data utility measures for each transformation technique across best data variants w.r.t privacy optimisation.

While PPT present high results for Statistic Similarity (Mean) and Correlation, it holds the worst result for Boundary Adherence, which may also explain the highest linkability risk. Concerning deep learning- and differentially private-based solutions, the Boundary Adherence is high but the remaining measures present a higher variability, which is reflected in the predictive performance results (Figure 4). ϵ-PrivateSMOTE also shows some variability for Statistic Similarity (Mean) and Correlation, probably due to the oversampling of the highest-risk cases. In summary, SMOTE, BorderlineSMOTE and ϵ-PrivateSMOTE present a median very close to the maximum for all data utility measures. However, only ϵ-PrivateSMOTE provides a good trade-off between linkability and predictive performance.

Although ϵ-PrivateSMOTE can provide competitive results, the biggest advantage is its simplicity (**RQ3**). We compare deep learning- and differentially private-based solutions with ϵ-PrivateSMOTE in terms of data variant generation. As shown in Table 2, ϵ-PrivateSMOTE can generate a large number of data variants much faster while consuming much fewer resources than the compared solutions. We do not present the computational cost results for the PPT as the parameters were handset. Sampling methods are omitted due to their linkability risk, making them non-competitive.

Technique	Nr. of generated data variants	Elapsed Time (min)	CPU (%)	GPU (%)	RAM (%)	Time \ Variant
Copula GAN	60	207,60	10,43	17,60	37,42	3,46
TVAE	60	125,26	13,09	16,76	37,42	2,08
CTGAN	60	206,94	10,59	41,66	37,44	3,44
DPGAN	240	857,37	8,19	23,07	47,34	3,57
PATE-GAN	240	355,33	14,59	9,96	47,42	1,48
ϵ-PrivateSMOTE	3375	553,28	22,41	0	22,41	0,16

Table 2 Comparison of computational costs between deep learning- and differentially private-based models and ϵ-PrivateSMOTE.

Table 2 provides the number of data variants generated for each approach, considering the number of QI set and the number of parameters of the transformation techniques. For example, PATE-GAN produces 16 data variants for each original data set, resulting in a total of 240 data variants. It is noticeable that despite ϵ-PrivateSMOTE generating more than 14 times the number of variants, it demonstrates efficiency by being approximately 9 times faster than PATE-GAN and using fewer resources. We experimented with GPU utilisation; however, our proposed approach requires more time to generate a data variant on the GPU due to a portion of operations running on the CPU and others on the GPU. The operations between the CPU and GPU introduce additional latency. Therefore, we opted to implement ϵ-PrivateSMOTE on the CPU for improved efficiency. Such results were conducted on a Linux-based system with the following specifications: kernel version 5.15.0-91-generic, 2.85GHz 24-Core AMD EPYC Processor, 256GiB RAM and a GeForce RTX 3090.

To conclude, ϵ-PrivateSMOTE emerges as a robust method that provides similar or better privacy guarantees than traditional and generative methods while optimising predictive performance, and exhibiting similar statistical properties with original data. Most importantly, it significantly reduces computation time and operates with reduced resource requirements when generating new data variants. However, we still face some limitations, especially regarding potential inference-based attacks. A further concern is the obfuscation of outliers. A data set generated with ϵ-PrivateSMOTE may still contain outliers, which can have significant privacy implications. Also, our solution focuses on tabular data and cannot handle other types of data such as images. All of these potential issues will be investigated in future work.

6 Conclusions

In this paper, we propose ϵ-PrivateSMOTE, a new strategy for privacy-preserving data sharing that combines data synthetisation via noise-induced interpolation with Differential Privacy to address the highest-risk cases concerning k-anonymity. Experimental results show that: *i)* ϵ-PrivateSMOTE is capable of providing competitive results concerning predictive performance; *ii)* it provides lower linkability risk, especially when privacy is the priority; and, *iii)* ϵ-PrivateSMOTE proved to be a more resource-efficient solution than deep learning- and differentially private-based approaches, improving time complexity by a factor of 9. Moreover, it performs well on lower dimensional data sets and does not require learning all real data.

Declarations

- Funding: Not applicable.
- Conflicts of interest/Competing interests: I declare that the authors have no competing interests as defined by Springer, or other interests that might be perceived to influence the results and/or discussion reported in this paper.
- Ethics approval: Not applicable.
- Consent to participate: Not applicable.
- Consent for publication: all of the material is owned by the authors and/or no permissions are required.
- Availability of data and material: Data is provided within the manuscript.
- Code availability: The code of the proposed method is available at https://github.com/tmcarvalho/privateSMOTE.
- Authors' contributions: Tânia Carvalho and Nuno Moniz wrote the main manuscript. All authors reviewed the manuscript.

References

[1] Carvalho T, Moniz N, Faria P, Antunes L. Survey on Privacy-Preserving Techniques for Microdata Publication. ACM Comput Surv. 2023 jul;55(14s). https://doi.org/10.1145/3588765.

[2] El Emam K, Mosquera L, Hoptroff R. Practical synthetic data generation: balancing privacy and the broad availability of data. O'Reilly Media; 2020.

[3] Nikolenko SI. Synthetic data for deep learning. arXiv preprint arXiv:190911512. 2019;.

[4] Dwork C. Differential privacy: A survey of results. In: International conference on theory and applications of models of computation. Springer; 2008. p. 1–19.

[5] Torkzadehmahani R, Kairouz P, Paten B. Dp-cgan: Differentially private synthetic data and label generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops; 2019. p. 0–0.

[6] European Commission.: Opinion 05/2014 on Anonymisation Techniques. Accessed Jan 2023. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf.

[7] Carvalho T, Faria P, Antunes L, Moniz N. Fundamental privacy rights in a pandemic state. Plos one. 2021;16(6):e0252169.

[8] Bird T, Kingma FH, Barber D. Reducing the computational cost of deep generative models with binary neural networks. arXiv preprint arXiv:201013476. 2020;.

[9] Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. SMOTE: synthetic minority over-sampling technique. Journal of artificial intelligence research. 2002;16:321–357.

[10] Pradipta GA, Wardoyo R, Musdholifah A, Sanjaya INH. Improving classifiaction performance of fetal umbilical cord using combination of SMOTE method and multiclassifier voting in imbalanced data and small dataset. Int J Intell Eng Syst. 2020;13(5):441–454.

[11] Dwork C, Roth A, et al. The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science. 2014;9(3–4):211–407.

[12] University of Pennsylvania.: Lecture1: Introduction to Differential Privacy and the Laplace Mechanism. Accessed Jun 2023. https://www.cis.upenn.edu/~aaroth/chatgpt_lecture_notes.pdf.

[13] Domingo-Ferrer J. A survey of inference control methods for privacy-preserving data mining. In: Privacy-preserving data mining. Springer; 2008. p. 53–80.

[14] Samarati P. Protecting respondents identities in microdata release. IEEE transactions on Knowledge and Data Engineering. 2001;13(6):1010–1027.

[15] Giomi M, Boenisch F, Wehmeyer C, Tasnádi B.: Anonymeter. Accessed Jun 2023. https://github.com/statice/anonymeter.

[16] Long Y, Bindschaedler V, Gunter CA. Towards measuring membership privacy. arXiv preprint arXiv:171209136. 2017;.

[17] Brickell J, Shmatikov V. The cost of privacy: destruction of data-mining utility in anonymized data publishing. In: Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining; 2008. p. 70–78.

[18] Holohan N, Antonatos S, Braghin S, Mac Aonghusa P. (k, ϵ)-Anonymity: k-Anonymity with ϵ-Differential Privacy. arXiv preprint arXiv:171001615. 2017;.

[19] Muralidhar K, Domingo-Ferrer J, Martínez S. ϵ-Differential Privacy for Microdata Releases Does Not Guarantee Confidentiality (Let Alone Utility). In: International Conference on Privacy in Statistical Databases. Springer; 2020. p. 21–31.

[20] Xu L, Skoularidou M, Cuesta-Infante A, Veeramachaneni K. Modeling tabular data using conditional gan. Advances in Neural Information Processing Systems. 2019;32.

[21] Yale A, Dash S, Dutta R, Guyon I, Pavao A, Bennett KP. Assessing privacy and quality of synthetic health data. In: Proceedings of the Conference on Artificial Intelligence for Data Discovery and Reuse; 2019. p. 1–4.

[22] Xie L, Lin K, Wang S, Wang F, Zhou J. Differentially private generative adversarial network. arXiv preprint arXiv:180206739. 2018;.

[23] Jordon J, Yoon J, Van Der Schaar M. PATE-GAN: Generating synthetic data with differential privacy guarantees. In: International conference on learning representations; 2018. .

[24] Mahiou S, Xu K, Ganev G. dpart: Differentially Private Autoregressive Tabular, a General Framework for Synthetic Data Generation. arXiv preprint arXiv:220705810. 2022;.

[25] Zhang J, Cormode G, Procopiuc CM, Srivastava D, Xiao X. Privbayes: Private data release via bayesian networks. ACM Transactions on Database Systems (TODS). 2017;42(4):1–41.

[26] Zhang Z, Wang T, Li N, Honorio J, Backes M, He S, et al. PrivSyn: Differentially Private Data Synthesis. In: 30th USENIX Security Symposium (USENIX Security 21). USENIX Association; 2021. p. 929–946. Available from: https://www.usenix.org/conference/usenixsecurity21/presentation/zhang-zhikun.

[27] Cai K, Lei X, Wei J, Xiao X. Data synthesis via differentially private markov random fields. Proc VLDB Endow. 2021 jul;14(11):2190–2202. https://doi.org/10.14778/3476249.3476272.

[28] Han H, Wang WY, Mao BH. Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. In: International conference on intelligent computing. Springer; 2005. p. 878–887.

[29] Tarekegn AN, Giacobini M, Michalak K. A review of methods for imbalanced multi-label classification. Pattern Recognition. 2021;118:107965.

[30] Spelmen VS, Porkodi R. A review on handling imbalanced data. In: 2018 international conference on current trends towards converging technologies (ICCTCT). IEEE; 2018. p. 1–11.

[31] Basha SJ, Madala SR, Vivek K, Kumar ES, Ammannamma T. A Review on Imbalanced Data Classification Techniques. In: 2022 International Conference on Advanced Computing Technologies and Applications (ICACTA); 2022. p. 1–6.

[32] Stadler T, Oprisanu B, Troncoso C. Synthetic data–anonymisation groundhog day. In: 31st USENIX Security Symposium (USENIX Security 22); 2022. p. 1451–1468.

[33] He J, Cai L, Guan X. Differential private noise adding mechanism and its application on consensus algorithm. IEEE Transactions on Signal Processing. 2020;68:4069–4082.

[34] van Breugel B, Sun H, Qian Z, van der Schaar M. Membership inference attacks against synthetic data through overfitting detection. arXiv preprint arXiv:230212580. 2023;.

[35] Giomi M, Boenisch F, Wehmeyer C, Tasnádi B. A Unified Framework for Quantifying Privacy Risk in Synthetic Data. arXiv preprint arXiv:221110459. 2022;.

[36] Prasser F, Eicher J, Spengler H, Bild R, Kuhn KA. Flexible data anonymization using ARX—Current status and challenges ahead. Software: Practice and Experience. 2020;50(7):1277–1304.

[37] Imbalanced-learn developers.: Imbalanced learn. Accessed Jan 2023. https://imbalanced-learn.org/stable/index.html.

[38] Patki N, Wedge R, Veeramachaneni K. The Synthetic Data Vault. In: 2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA); 2016. p. 399–410.

[39] Qian Z, Cebere BC, van der Schaar M.: Synthcity: facilitating innovative use cases of synthetic data in different data modalities. Available from: https://arxiv.org/abs/2301.07573.

[40] Ho TK. The random subspace method for constructing decision forests. IEEE transactions on pattern analysis and machine intelligence. 1998;20(8):832–844.

[41] Chen T, Guestrin C. Xgboost: A scalable tree boosting system. In: Proc. 22nd ACM Int. Conf. SIGKDD; 2016. p. 785–794.

[42] Agresti A. An introduction to categorical data analysis. Wiley; 1996.

[43] Mcculloch W, Pitts W. A Logical Calculus of Ideas Immanent in Nervous Activity. Bulletin of Mathematical Biophysics. 1943;5:127–147.

[44] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine learning in Python. the Journal of machine Learning research. 2011;12:2825–2830.

[45] Weng CG, Poon J. A new evaluation measure for imbalanced datasets. In: Proceedings of the 7th Australasian Data Mining Conference-Volume 87; 2008. p. 27–32.

[46] Benavoli A, Corani G, Demšar J, Zaffalon M. Time for a change: a tutorial for comparing multiple classifiers through Bayesian analysis. The Journal of Machine Learning Research. 2017;18(1):2653–2688.

[47] Kruschke J, Liddell T.: The bayesian new statistics: Two historical trends converge. SSRN Electron. J.

[48] Vanschoren J, van Rijn JN, Bischl B, Torgo L. OpenML: Networked Science in Machine Learning. SIGKDD Explorations. 2013;15(2):49–60. https://doi.org/10.1145/2641190.2641198.