
EMPIRICAL EVALUATION OF STRUCTURED SYNTHETIC DATA PRIVACY METRICS: NOVEL EXPERIMENTAL FRAMEWORK

Milton Nicolás Plasencia Palacios^{1,2,*}, Alexander Boudewijn¹, Sebastiano Saccani¹, Andrea Filippo Ferraris³, Diana Sofronieva¹, Giuseppe D’Acquisto⁴, Filiberto Brozzetti⁴, Daniele Panfilo¹, and Luca Bortolussi²

¹Aindo SpA, Trieste, Italy

²University of Trieste (UniTs)

³LAST-JD, Alma AI, Alma Mater Studiorum, University of Bologna & DIKE research group, PREC department, Vrije Universiteit Brussel

⁴Luiss University, Rome, Italy

*Corresponding author, plasencia.milton@gmail.com

April 6, 2026

ABSTRACT

Synthetic data generation is gaining traction as a privacy enhancing technology (PET). When properly generated, synthetic data preserve the analytic utility of real data while avoiding the retention of information that would allow the identification of specific individuals. However, the concept of data privacy remains elusive, making it challenging for practitioners to evaluate and benchmark the degree of privacy protection offered by synthetic data. In this paper, we propose a framework to empirically assess the efficacy of tabular synthetic data privacy quantification methods through controlled, deliberate risk insertion. To demonstrate this framework, we survey existing approaches to synthetic data privacy quantification and the related legal theory. We then apply the framework to the main privacy quantification methods with no-box threat models on publicly available datasets.

Keywords synthetic data · privacy · privacy quantification · attacks on privacy · privacy measurements · machine learning and privacy

1 Introduction

Synthetic data generation is increasingly recognized as a powerful privacy-enhancing technology (PET). By preserving the properties of real-world data, synthetic data allow organizations to mitigate privacy risks in research, innovation, and the development of artificial intelligence (AI) systems. This has led to a growing interest in the PET in high-impact sectors such as healthcare [1, 2] and finance [3].

However, there is no universal standard for quantifying privacy risks, leading to fragmented approaches in practice. Privacy is an interdisciplinary concept that spans legal, technical, and ethical dimensions. Practitioners often struggle to justify their choice of privacy metrics or methods, given the lack of clear guidance and the various tools available. Moreover, the practical challenges of meeting regulatory compliance further complicate this task and can limit innovation in data-driven fields [4]. This highlights the urgent need for comprehensive analyses and standardized frameworks for privacy assessment, as identified in prior works [5, 6, 7, 8, 9].

To address these challenges, this paper introduces a benchmarking framework for privacy quantification methods in the context of synthetic data generation. This framework is centered around deliberate and controlled insertion of risk to synthetic datasets in both an idealized manner (directly leaking proportions of real records, based on [10]), and in manners modeling how privacy risks occur in practice (generator overfitting and privacy budgets). To demonstrate the framework, we first survey existing methods for quantifying synthetic data privacy. We then describe our benchmarking

framework and apply it in an experimental analysis of the surveyed privacy quantification approaches. Throughout, we relate our framework, as well as privacy metrics, to relevant legal standards and theory. Our findings indicate that most metrics consistently reflect controlled risks in both theoretical and practical scenarios. Measured risk is typically proportional to inserted risk in idealized settings, though data characteristics, such as a high proportion of categorical attributes, can amplify privacy risks. By combining legal insights, a comprehensive classification of privacy quantification methods under a no-box threat model, and a robust benchmarking framework, our work contributes to the standardization of privacy assessment for synthetic data. This supports researchers and practitioners in advancing synthetic data as a viable PET in real-world applications.

2 Background

2.1 Synthetic data

Contemporary usage of synthetic data has become increasingly significant in healthcare [11, 12], finance, and research, owing to its notable privacy-enhancing advantages and capacity to foster data-driven innovation. Synthetic data is often hailed as a promising privacy-enhancing technology [13, 14, 15, 16]. However, apprehensions about regulatory sanctions frequently deter market actors from fully adopting synthetic data as its status within the regulatory framework remains ambiguous [4, 9], where standardized assessment methodologies and clear legal criteria are lacking [17].

Following Jordon et al. [18], we formalize the concept of synthetic data through Definition 2.1.

Definition 2.1 (Synthetic data, [18]) *Synthetic data is data that has been generated using a purpose-built mathematical model or algorithm (the **generator**), with the aim of solving a (set of) data science task(s).*

Let D denote a tabular database describing data subjects through attribute set $\mathcal{A}(D)$. Rows $d \in D$ are $|\mathcal{A}(D)|$ -tuples with a value $v(d, a)$ for each attribute $a \in \mathcal{A}(D)$. For $A = \{a_1, a_2, \dots, a_m\} \subseteq \mathcal{A}(D)$, we denote by $v(d, A)$ the tuple $(v(d, a_1), v(d, a_2), \dots, v(d, a_m))$. We denote by $\mathcal{D}$ the space of all possible rows over $\mathcal{A}(D)$.

In the remainder, we restrict our analysis to *Synthetic data generated from original real-world data* (hereafter “synthetic data”). We denote by $\hat{D} \sim G(D)$ that synthetic dataset $\hat{D}$ was obtained from a stochastic generator G trained on an original dataset D . This restriction rules out data that is generated through modeling processes that do not infer generation conditions directly from real-world data, such as simulation. Our form of synthetic data has become widespread in recent years, in part due to the success of machine learning-based generation methods [19, 20, 7, 6, 18].

2.2 Legal Background

Synthesizing data from personal data inherently involves processing personal data, as the synthetic output is derived from an original dataset containing personal data. This principle parallels other anonymization processes. Consequently, such processing must fully comply with the relevant regulatory framework (e.g., the GDPR). It is a separate question whether the synthetic data itself can be considered personal data or not. Depending on how effectively privacy risks have been mitigated during its generation, synthetic data may fall under either the personal or non-personal data category.

2.2.1 Legal Definition of Personal Data.

GDPR Article 4(1), defines personal data as “any information relating to an identified or identifiable natural person (‘data subject’).” This definition comprises three essential elements: (i) “any information,” indicating the broad scope of data types that may be considered personal; (ii) “identified or identifiable natural person,” which includes both direct and indirect methods of identification; and (iii) “relating to,” which establishes the connection between the data and the individual [21, 22]. The term “relating to” is particularly significant, as it can be interpreted in multiple ways—such as “in content” (where data explicitly describes an individual), “in purpose” (where data is processed to infer information about an individual), or “in result” (where the use of data indirectly affects an individual) [21, 23].

2.2.2 Anonymization and the law

Anonymity has long been understood as a decisive factor in determining the applicability of data protection law, including under the GDPR. According to the GDPR, an anonymization process should ensure that, once effective techniques have been applied, the data no longer carries the risk of identifying a specific individual and is thus exempt from GDPR’s scope. Recital 26 states that data should be considered anonymized when it has been “rendered in such a manner that the data subject is not or no longer identifiable.” Consequently, synthetic data should be regarded as anonymous only when it precludes any form of direct, indirect, or auxiliary identification. Yet, it would be

misleading to interpret this as establishing a strict binary between personal and non-personal data. Rather, the notion of identifiability under EU law invites a more nuanced, context-aware reading—one that accounts for the risks and means of re-identification as they arise in specific technological and organizational settings.

2.2.3 Reasonableness & proportionality principles

The principle of proportionality plays a central role in balancing the benefits of data processing with privacy protection. The possibility of re-identification—i.e., transforming anonymized data back into personal data—must be evaluated according to its likelihood within a given dataset. Recital 26 clarifies that determining whether an individual is identifiable requires considering all methods “reasonably likely” to be employed:

*To determine whether a natural person is identifiable, account should be taken of all the means **reasonably** likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly. To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the **available technology at the time of the processing and technological developments**.*

In other words, the robustness of the anonymization process is gauged by the “means reasonably to be used” test [21], and the risk of re-identification must be evaluated on the basis of what an adversary could realistically achieve, not in hypothetical or extreme scenarios. If identifying an individual requires unreasonable effort, time, or resources, the data can be considered anonymous and is therefore exempt from data protection rules. The feasibility of re-identification must always be assessed in relation to existing technologies and anticipated future developments.

2.3 Anonymization benchmarks

2.3.1 WP29 typified risks

As previously noted, data is considered anonymized when the data subject is “not or no longer identifiable,” meaning that re-identification is rendered impossible. However, because this determination is inherently ex-post, assessing anonymization can be challenging. To facilitate these evaluations, the Article 29 Data Protection Working Party (WP29) established in 2014 three key risk benchmarks for effective anonymization: Singling Out, Linkability, and Inference [22].

- **Singling Out** corresponds to the possibility to isolate some or all records which identify an individual in the dataset;
- **Linkability** is the ability to link, at least, two records concerning the same data subject or a group of data subjects (either in the same database or in two different databases). If an attacker can establish (e.g. by means of correlation analysis) that two records are assigned to a same group of individuals but cannot single out individuals in this group, the technique provides resistance against “singling out” but not against linkability;
- **Inference** is the possibility to deduce, with significant probability, the value of an attribute from the values of a set of other attributes (exploiting statistical correlations within the data). In an attribute inference attack (AIA), if a dataset D trains a generator $G(D)$ and synthetic data $\hat{D} \sim G(D)$ is accessible alongside auxiliary information, an adversary may combine these to infer unknown target attributes.

These legal risk benchmarks inform the adversarial threat models employed in our analysis. By mapping each WP29 risk to specific adversary objectives and capabilities, the threat model configurations developed in the following sections are designed to reflect what may be considered “reasonably likely” within the meaning of Recital 26 GDPR. In this way, our framework aligns technical risk assessment with the contextual and proportional nature of legal anonymization requirements.

Legally speaking, a robust anonymization procedure — that is, any data processing which produces output that, at a given time and within a specific context, can no longer be considered personal data — must effectively prevent these three typified re-identification risks. Building on these risk factors, this research proposes standardized frameworks for assessing synthetic data as a privacy-enhancing technology (PET) under the GDPR, thereby guiding data controllers, researchers, and policymakers.

2.3.2 Alternative definitions of privacy risks

The WP29 risks do not always align with technical classifications of typified privacy risks (or re-identification risks). Both Hu et al. [24] and Raab et al. [8] classify risks into two categories: 1. Identity disclosure: the ability to identify an

individual in a dataset from a set of known characteristics; 2. Attribute disclosure: the ability to find out a previously unknown attribute value associated with a record in the original data from a set of known characteristics. These definitions align with those of “singling out” and “inference” under the WP29 to an extent. However, they differ from the WP29 definition in two key ways: 1. Raab et al.’s [8] definitions focus on *what can be inferred* from an information release: either an individual’s identity, or a value of one of their attributes. For this reason, linkability is not considered a separate risk; 2. The role of auxiliary information (“known characteristics”) is directly incorporated into Raab et al.’s [8] definitions of identity and attribute disclosure. This provides an idea of how an adversary would exploit a risk in an attack. The required *auxiliary information* could be interpreted as linked data, technically making any form of attack an exploitation of linkability under WP29.

Besides WP29 and the definitions by Hu et al. [24] and Raab et al. [8], attacks compromising confidentiality from the field of adversarial machine learning [25] are also commonly applied to synthetic data generators. In the context of synthetic data, the main risks (not due to deliberate system misuse) are: 1. Membership inference attacks (MIAs): an attack in which the adversary can deduce whether a given record was in the training data of a given machine learning model, based on this model and/or its output. In the context of synthetic data, a MIA is successful if, for some record $d \in \mathcal{D}$, the attacker can infer whether or not $d \in D$ through access to $G(D)$ and/or $\hat{D} \sim G(D)$; 2. Model stealing or shadow modeling: the adversary can reproduce the input-output relationship of a machine learning model (such as a synthetic data generator).

Stadler et al. [26] provide formalizations to model the WP29 attacks as specific forms of MIAs each. Both adversarial ML approaches tend to require considerable generator access, as introduced in Section 2.4.

2.4 Threat models

We assume that the adversary aims to obtain sensitive information about individuals whose data is contained in an original dataset used to train a synthetic data generator. Below, we describe adversary resources that can be exploited in privacy attacks when using synthetic data as a PET. These are the tools and information available to the adversary in conducting attacks. These resources can be subdivided into three categories: 1. the synthetic dataset; 2. generator access; and 3. auxiliary information.

Under **no-box generator access**, the adversary exploits access to synthetic data without leveraging information about the model used to construct it. They may leverage auxiliary information sources, such as knowledge of particular individuals in the training data, including if this can be derived from the synthetic data’s performance in downstream tasks; the source of the training data; modeling decisions and circumstances during generator development and training; the structure of the (synthetic) data (e.g. if attribute names are indicative of specific circumstances). This form of auxiliary information is formalized in Definition 2.2.

Definition 2.2 (Auxiliary information) *Let D be a dataset with attributes $\mathcal{A}(D)$. An adversary has **auxiliary information** if there exists a subset $A \subseteq \mathcal{A}(D)$ of attributes and a subset $D' \subseteq D$, such that the adversary knows the values $v(d, A)$ for all $d \in D'$. We refer to the attributes in set A as **quasi-identifiers**. We denote auxiliary information consisting of quasi-identifier set A and records $D' \subseteq D$ by $\text{Aux}(A, D')$.*

Black-box generator access is available when a synthetic data service provider offers a training algorithm for a generator to the end-user as a black box. This means the end-user can: 1. Use the generator to produce an unbounded number of synthetic datasets; 2. Train the generator on any available input dataset matching the algorithms capabilities. However, the end-user does not have direct access to the model’s internal routines and (hyper)parameters.

White-box generator access is available when the synthetic data provider provides not only access to the input-output relationship of their training algorithm (as under black-box access), but also to the full implementation of the algorithm. This could for instance be the scenario where a synthetic data provider makes their training algorithm open-source. It provides adversaries with full information on how the generator G (a typically stochastic function) maps real data D to synthetic data $\hat{D} \sim G(D)$. The conditions to carry out white-box or even black-box attacks are typically not met in practical scenarios [27].

3 Relation to prior research

Several works present **literature reviews and surveys** on synthetic data quality assessment, focusing on or including a classification of privacy evaluation methods. Kabaachi et al. [28] study utility and privacy quantification in a review of 92 papers. They find that, while 85% (78/92) of papers include synthetic data for privacy use cases, only 42% of papers (39/92) use a privacy quantification method. Of the papers conducting privacy quantification, 38% (15/39) rely solely

on either differential privacy or data masking for additional privacy protection. The remaining 62% (24/39) employ at least one of a diverse set of privacy quantification methods.

Hu et al. [24], Osorio et al. [29], and Boudewijn et al. [9] identify common privacy risks, classifying them into several categories. These include classifications based on potential risk factors (e.g. identity and attribute disclosure [24]), potential weaknesses of generated synthetic data (e.g. the use of skewness or background knowledge [29]), and properties of data privacy quantification methods (e.g. mathematical properties, statistical indicators, attack simulations [9]).

Experimental benchmarks for privacy quantification in synthetic data include Hernandez et al. [30] and Yan et al. [31], who propose a number of synthetic data fidelity, utility, and privacy metrics and provide experimental benchmarks. Both works compute multiple metrics on multiple synthetic datasets, proposing the use of the metrics as benchmarks. However, these works primarily benchmark synthetic data as a PET, rather than the privacy quantification methods themselves. By contrast, our aim is to provide empirical benchmarks for privacy quantification methods. This could lead to a more nuanced understanding of their efficacy and sensitivity to parameters and data properties.

Two notable recent works involve **stakeholders or expert panels** in building consensus around synthetic data privacy quantification. Brunelli et al. [32] present a synthetic data privacy and utility toolbox (“*SURE*”, for “*Synthetic Data: Utility, Regulatory compliance, and Ethical privacy*”) involving both statistical privacy indicators and attack simulations. The SURE-project directly invokes feedback of synthetic data end-users, namely data scientists and data protection officers (DPOs). A detailed feedback report of these users is still forthcoming. Pilgram et al. [27] adopt a similar approach to consensus building through expert feedback. Instead of direct stakeholders and end-users, they gather feedback through an expert panel consisting of long-term editorial board members of one journal and one conference on data privacy. They find that attack simulations are regarded positively, as they directly model a concrete risk (identity or attribute disclosure). The expert opinions in these studies are based on custom use of a specific toolbox and descriptions of metrics. Our work, by contrast, provides experimental results under several risk models, offering empirical insight into metrics’ responses to controlled inserted risks.

Legal scholars describe synthetic data as a promising privacy-enhancing technology (PET), particularly in contexts where access to granular data is necessary but privacy concerns are paramount [7]. However, the application of the General Data Protection Regulation (GDPR) to synthetic data lacks clarity [23]. The legal community appears aware of the need for quantitative privacy assessment, welcoming interdisciplinary research into standardization of privacy assessment in synthetic datasets, as the lack of consensus may hinder both innovation and regulatory compliance [6]. Our work contributes to this research gap by showcasing the merits and weaknesses of prevailing privacy quantification methods theoretically and empirically.

4 Synthetic Data Privacy Quantification Taxonomy

In this section, we propose a taxonomy of approaches to the quantification of the degree of privacy protection of synthetic datasets and their generators. The main distinction in our taxonomy centers around the practical approach to measuring privacy, categorized into the three approaches: 1. privacy properties; 2. statistical privacy indicators; and 3. attack simulations.

Figure 1 provides an overview of the approaches to synthetic data privacy quantification. In the remainder of this section, we detail the elements in the taxonomy.

4.1 Differential Privacy and other Privacy Properties

4.1.1 Differential privacy

Differential privacy (DP, see [36]), as formalized in Definition 4.1, is gaining traction as a privacy mechanism in synthetic data generation.

Definition 4.1 (Differential Privacy) *A randomized algorithm $\mathcal{M}$ is (ϵ, δ) -differentially private $((\epsilon, \delta)$ -DP) if for any set S of possible outputs of $\mathcal{M}$, we have:*

$$\mathbb{P}[\mathcal{M}(D) \in S] \leq e^\epsilon \cdot \mathbb{P}[\mathcal{M}(D') \in S] + \delta \quad (1)$$

for all databases D, D' such that $D' = D \setminus \{d\}$ for some d . We refer to $(\epsilon, 0)$ -DP simply as ϵ -DP.

Intuitively put, DP limits the influence of any individual record d in a given input dataset D on the DP generator $G(D)$ through *privacy budget* ϵ . This makes it difficult for adversaries to target specific records in their attacks. The privacy budget is specified by the user, allowing them to directly control the degree of privacy introduced. This, along with mathematical guarantees (see, e.g. DP’s immunity to post-processing [36]) make DP a widespread privacy framework in practice.




Figure 1: Overview of synthetic data privacy quantification methods

A plethora of methods was introduced to obtain DP generators [37, 38, 39, 40, 41, 42]. DP is a property of the generator, added a priori. This is in contrast to metrics and attack simulations that are evaluated a posteriori, once the generator is fully trained and produces synthetic data. We therefore include DP in risk modeling, choosing a range of privacy budgets to control the amount of risk deliberately inserted for benchmarking (see Section 5). DP further allows datasets to be shared securely by responding to specific queries rather than releasing entire datasets, while enabling auditing of queries and protecting privacy through further anonymization techniques [22].

4.1.2 k -Anonymity and related concepts

While DP is a property of information release systems, k -anonymity and related properties pertain to anonymized datasets themselves. The underlying notion of k -anonymity is that *uniqueness is a core risk factor*. For instance, while “place of residence” is not a direct identifier, publicizing a highly unique location can still put an individual at risk. A traditional approach to deal with such risks is *generalization*. In the example, “place of residence” could be generalized to “region”. This offers more privacy protection, as it makes it harder to single out individuals with unique locations. The concept of k -anonymity formalizes this form of protection. Formally, a dataset D is k -anonymous if any combination of quasi-identifier values present in D occurs at least k times. Generalization until k -anonymity is achieved divides the dataset up into equivalence classes with respect to the quasi-identifiers, each with a cardinality of at least k .

By itself, k -anonymity does not prevent AIAs [22]. Suppose in a k -anonymous set, city is generalized to country and age is generalized to age groups. However, 99% of individuals in Italy aged 50-55 have a high income. Then an adversary cannot re-identify the record of a specific individual from Trieste, aged 52. However, they can infer that the individual is highly likely to have a high income. This problem of parameter inference is addressed by l -diversity, requiring that in each equivalence class, the sensitive attributes take on at least l distinct values. The primary objective is to minimize equivalence classes with low attribute diversity, ensuring that attackers with prior knowledge about a specific individual face substantial uncertainty [22]. While conceived as a privacy protection property, the concept of l -diversity can itself be exploited in AIAs. This is discussed in Section 4.3.3.

4.2 Statistical Privacy Indicators

Statistical privacy indicators typically measure the occurrence of potentially hazardous synthetic records. In most cases, this corresponds to identifying if any synthetic records are more *similar to specific original records* than statistically reasonable. In this paper, we use the most common indicator categories, namely: 1. identical match share (IMS, also known as replicated uniques [43], repU [8], and unique exact match [41]); 2. distance-to-the-closest-record (DCR); and 3. k -nearest neighbor (K -NN)-based privacy metrics. Other statistical indicators for privacy assessment in synthetic data for instance include the maximum mean discrepancy (MMD) to test whether overfitting took place [44, 45], but are rare in practice.

4.2.1 Identical match share

IMS metrics simply document how many of the synthetic records are also in the real dataset [41, 14, 46]. In our experiments, we express this as a proportion of the number of rows in the real dataset, thus: $\text{IMS}(D, \hat{D}) := \frac{|D \cap \hat{D}|}{|\hat{D}|}$.

4.2.2 Distance to the closest record

DCR metrics are computed by comparing two quantities: the *Synthetic to Real distance* (SRD) and the *real to real distance* (RRD). Let $\text{Dist} : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ be a distance metric. For a given synthetic record $\hat{d} \in \hat{D}$, the distance to its closest real record, as formalized in equation (2).

$$\text{SRD}(\hat{d}) := \min_{d \in D} \text{Dist}(\hat{d}, d) \quad (2)$$

The RRD is typically computed using a *holdout set*. Partition D into two sets D_1, D_2 . For a given record $d_1 \in D_1$, the RRD is then the distance to its closest record in D_2 . This is formalized in equation (3).

$$\text{RRD}(d_1) := \min_{d_2 \in D_2} \text{Dist}(d_1, d_2) \quad (3)$$

Suppose, for some synthetic record $\hat{d} \in \hat{D}$, we have a statistically unusually small SRD. This means that $\hat{d}$ is very close to some real record $d \in D$ in a manner that may not just be due to stochasticity. $\hat{d}$ may therefore leak sensitive information. DCR metrics assess whether such records exist by comparing the SRD and RRD distributions. In our experiments, we will use a percentile-based comparison, as is common in practice [14, 47, 46]. In particular, we compute the DCR as in equation (4).

$$\text{DCR}(\hat{D}, D) := \frac{\left| \left\{ \hat{d} \in \hat{D} : \text{SRD}(\hat{d}) < \text{RRD}_\alpha \right\} \right|}{\frac{\alpha}{100} \cdot |D_1|} \quad (4)$$

where α is a percentage (we use 2% throughout) and RRD_α is the α^{th} percentile of the RRD distribution. Equation (4) is the ratio of SRD values below the α^{th} RRD percentile and the number of RRD values that percentile (naturally $\alpha\%$ of D_1).

4.2.3 k -Nearest neighbors metrics

Indicators using k -nearest neighbors (k -NN) generalize DCR-based metrics by comparing neighborhoods of synthetic and real records. We define the synthetic to real k -neighborhood of a synthetic record $\hat{d} \in \hat{D}$, denoted by $N_{\text{SRD}}^k(\hat{d})$ as the k nearest records to $\hat{d}$ in $\hat{D}$ with respect to the SRD. We define the real-to-real k -neighborhood $N_{\text{RRD}}^k(d)$ analogously

using the RRD. We then apply equation (4) with the means of the neighborhoods $N_{\text{SRD}}^k(\hat{d})$ and $N_{\text{RRD}}^k(d)$ in place of the SRD and RRD.

4.3 Attack Simulation

We include attack simulations tailored to the WP29 risks. All three of these risks can be implemented with different attack mechanisms. Below, we provide an overview of possible risk mechanisms and our implementations.

4.3.1 Membership Inference and Singling Out

Identity disclosure is most closely modeled by membership inference attacks (MIAs) [10, 35, 33]. In adversarial ML, a MIA is an attack in which the adversary can determine if a given individual’s record was in the training dataset of a given model [25]. In the context of synthetic data with no-box generator access, MIAs based on uniqueness [10, 48] and on distances [2, 49] are the most common.

Uniqueness-based MIAs detect outliers in the synthetic data and propose these as potentially real original records. The underlying assumption is that outliers in the original data may be more susceptible to MIAs. By definition, they are located in a sparse region of the dataset, due to which the generator may locally overfit.

Anonymeter’s singling out attack [10] looks for unique records with respect to a subset of attributes in the synthetic dataset in a brute-force manner. Anonymeter models attacks as guesses. A singling out attack is essentially a guess of the form “a record d with parameter values $(v(d, a_1), v(d, a_2), \dots, v(d, a_n))$ is a unique record in the original dataset”, in which the combination of values is informed by the synthetic data.

Let $\mathbf{g} := \{g_1, g_2, \dots, g_{N_A}\}$ be a set of $N_A \in \mathbb{N}$ guesses and $\mathbf{o} = \{o_1, o_2, \dots, o_{N_A}\}$ their respective outcomes, where o_i , $i = 1, 2, \dots, N_A$ is defined through equations (5).

$$o_i := \begin{cases} 1, & \text{if guess } g_i \text{ was correct} \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Guess effectiveness is interpreted as a Bernoulli trial with success probability $\hat{r}$. Estimators r, δ_r can then be computed such that $\hat{r} \in r \pm \delta_r$ at confidence level α via the Wilson Score Interval.

We note that several novel approaches further explore the degree to which *isolation* serves as an attack vector. Meeus et al. [35] use a form of outlier detection to identify potential targets for black-box MIAs. D’Acquisto et al. [50] rigorously study the degree to which information release systems increase an adversary’s abilities to isolate and subsequently identify real individual’s information in a novel probabilistic manner.

Domias [33] is a MIA approach grounded in **overfitting detection**. Unlike previous methods (e.g., [51, 52, 53]) which rely solely on synthetic data to infer membership, Domias explicitly targets overfitting by incorporating information from a holdout set of real records. The original dataset D is partitioned into a training set D_{train} , a control set D_{control} , and a reference dataset $D_{\text{reference}}$. A generative model $G(\cdot)$ is trained on D_{train} to produce a synthetic dataset $\hat{D} = G(D_{\text{train}})$. The core premise of this attack is that the generator overfits to outliers within the training set; consequently, these points exhibit a significantly higher density in the learned synthetic distribution than in the true underlying distribution. To quantify this, the method estimates the densities of the synthetic data and $D_{\text{reference}}$. The attack evaluates test points—drawn from D_{train} (specifically targeting outliers) and D_{control} —using Eq. 6:

$$A(x^*) = f\left(\frac{p_{\text{synth}}(x^*)}{p_{\text{reference}}(x^*)}\right) \quad (6)$$

where p_{synth} represents the learned distribution of the synthetic data, $p_{\text{reference}}$ is the learned distribution of the reference dataset, x^* is the target point, and $f : \mathbb{R} \rightarrow [0, 1]$ is a monotonically increasing function. If x^* is an outlier from the training set, the numerator dominates the denominator. Conversely, if x^* belongs to D_{control} , the numerator and denominator yield approximately equal values.

Black-box MIAs generally operate as follows. For a given target record $d_t \in \mathcal{D}$, the adversary aims to predict whether d_t was in the training set D of synthetic dataset $\hat{D} \sim G(D)$. To do so, the adversary exploits their black-box access to the generator’s training algorithm. This typically involves the following steps [26, 54, 42, 35, 33, 34]:

1. The adversary takes k datasets $D_1, \dots, D_k$ with $D_i \subset \mathcal{D} \setminus \{d_t\}$, $i = 1, \dots, k$;
2. For each such dataset D_i , the adversary trains a generator $G(D_i)$ and produces an output dataset $\hat{D}_i \sim G(D_i)$. The adversary stores datasets $\hat{D}_i$ along with the label “no target”;

3. Next, the adversary creates the k datasets $D'_i := D_i \cup d_t$ for $i = 1, \dots, k$, obtained by adding the target record to the chosen datasets;
4. The adversary now trains generators $G(D'_i)$, producing synthetic datasets $\hat{D}'_i \sim G(D'_i)$. The adversary stores the sets $\hat{D}'_i$ along with the label “target”;
5. The adversary now has a labeled collection of synthetic datasets. They leverage this collection to train a model $\mathcal{M}$ that predicts whether the target record d_t was in the dataset used to train a generator G , based on an output synthetic dataset of G .

Step (5) usually involves feature extraction techniques to represent the collection of synthetic datasets as a table or other format of use to easily train $\mathcal{M}$ (typically a machine learning classifier).

Given the high computational intensity of black-box attacks, methods that optimize target selection are highly preferable. Both vulnerable record discovery [35] and strategies targeting overfitted regions [33, 34] are favored specifically for their ability to significantly reduce this computational burden."

4.3.2 Linkability

To the best of our knowledge, Anonymeter’s linkability attack [10] is the only direct no-box implementation of linkability attacks. In this attack, the generator $G(D)$ is trained on original dataset D , including all of its attributes $\mathcal{A}(D)$, giving dataset $\hat{D} \sim G(D)$. Next, attribute set $\mathcal{A}(D)$ is partitioned into two (or more) mutually exclusive attribute sets A_1, A_2 , such that $A_1 \cup A_2 = \mathcal{A}(D)$. The adversary’s objective is now to link the values $v(d, A_1)$ and $v(d, A_2)$ together, giving the overall original record d , for $d \in D$.

To do so, the adversary uses the synthetic dataset $\hat{D}$. In particular, they find the k nearest synthetic neighbors (with respect to the SRD) to the partial records $\{v(d, A_1) : d \in D\}$. Let $\mathbf{1}_{A_1}(d)$ be the set of indices of these k nearest synthetic neighbors of $d \in D$. Then, they repeat this process for partial records $\{v(d, A_2) : d \in D\}$, giving index sets $\mathbf{1}_{A_2}(d)$ of k nearest synthetic neighbors of $d \in D$ with respect to attributes A_2 . A linkability attack is now considered a success if $\mathbf{1}_{A_1}(d) \cap \mathbf{1}_{A_2}(d) \neq \emptyset$. In other words: a linkability attack against a record d is successful if there is a k^{th} synthetic nearest neighbor $\hat{d}$ of d with respect to only attributes A_1 that is also a k^{th} synthetic nearest neighbor of d with respect to only attributes A_2 . SRD calculations are based on the Gower distance. The same guess-based approach to attack efficacy is used as in the uniqueness-based MIA. Throughout, we use the same parameters as used in [10].

4.3.3 Attribute inference

Like identify disclosure attacks, attribute inference attacks (AIAs) can rely on several attack mechanisms. These again include uniqueness and distances, but also ML models trained to identify values of targeted parameters.

Uniqueness-based AIAs seek to infer target t attribute values by isolating unique (equivalence classes of) records in the synthetic dataset $\hat{D}$ with respect to auxiliary information $\text{Aux}(A, D')$. The Generalized Targeted Correct Attribution Probability (GTCAP, [55]) does so by finding l -diverse equivalence classes in $\hat{D}$ with respect to t with a low l -value. Suppose a dataset is anonymized through generalization until k -anonymity is achieved for non-target attributes. Then, the set has several equivalence classes with at least k records each, for which all non-target attributes are the same (or in the same radius for continuous attributes). But if the l -diversity is low in a given equivalence class, then all these records will have similar values of target attribute t .

The l -diversity with low l of equivalence classes with respect to a specific sensitive attribute makes the class susceptible to a parameter inference attack. GTCAP is an approach to probabilistically quantify this susceptibility. The GTCAP is a generalization of the TCAP [56], defined for synthetic record $\hat{d} \in \hat{D}$ by equation (7), where the square brackets are Iverson brackets and $A \subseteq \mathcal{A}(D)$ is a set of quasi-identifiers in the auxiliary information of the adversary.

$$\text{TCAP}(D, \hat{d}) := \frac{\sum_{d \in D} \left[v(d, A) = v(\hat{d}, A), v(d, t) = v(\hat{d}, t) \right]}{\sum_{d \in D} \left[v(d, A) = v(\hat{d}, A) \right]} \quad (7)$$

The GTCAP generalizes the TCAP for continuous attributes, by replacing equality with membership in spheres of given radius around $v(\hat{d}, A)$ and $v(\hat{d}, t)$.

In **distance-based AIAs** (also known as majority-voting attack; local neighborhood attack, see, e.g., [10, 54, 2, 49]), adversaries exploit similarities between synthetic records and records from their auxiliary information (see

Definition 2.2). Suppose an adversary has access to: 1. a synthetic dataset $\hat{D}$; and 2. auxiliary information $\text{Aux}(A, D')$. In a distance-based attack, the adversary simply estimates the value $v(d', t)$ of record $d' \in D'$ and target attribute $t \in \mathcal{A}(D) \setminus A$ by assigning it the value of the nearest synthetic record, i.e.: $v(\hat{d}^*, t)$ for $\hat{d}^* = \arg \min_{\hat{d} \in \hat{D}} \text{Dist}(\hat{d}, d')$. In some cases, this is generalized to an aggregate of the k nearest synthetic neighbors of d' . However, experiments typically show that $k = 1$ leads to the best attack performance [2, 57].

ML-based AIAs (also known as “inference-on-synthetic” [54]) operate in a similar manner to distance-based AIAs, but instead of relying on distances, they rely on ML models. The adversary again has access to synthetic dataset $\hat{D}$ and auxiliary information $\text{Aux}(A, D')$. They want to use $\hat{D}$ to estimate the values of a target attribute $t \in \mathcal{A}(D) \setminus A$ for the records in D' . To do so, they train an ML model M , such that $M(v(\hat{d}, A)) \approx v(\hat{d}, t)$ for $\hat{d} \in \hat{D}$. They then apply M to the auxiliary information to estimate the target attributes of real individuals’ records $v(d, t)$. Model M is a classifier in case t is categorical and a regression model in case t is numeric.

Black-box AIAs can be conducted as a sequence of black-box MIAs [26]. Suppose that $d \in D$ is a target record, for which auxiliary information (in the sense of Definition 2.2) is available for attributes $A \subseteq \mathcal{A}(D)$. Suppose now that t is a categorical target attribute with range T . An adversary can then conduct $|T|$ MIAs, one for each possible value of t . If d was in the training data, then the MIA with d ’s value of t will be successful.

5 Empirical Assessment Framework for Tabular Synthetic Data Privacy Metrics

To empirically assess privacy quantification methods’ efficacy, we introduce the concept of a “*risk model*,” i.e. *a manner to deliberately introduce privacy risks in synthetic datasets with the purpose of assessing how well privacy quantification methods identify these inserted risks*. Below, we introduce three risk models studied in our experiments, motivating their use and detailing their specific parametrization.

5.1 Risk models

5.1.1 Leaky risk model

By leaky risk model, we refer to the following experimental setup. The real dataset D is divided into three sets: training (D_t), control (D_c), and release (D_r). The “synthetic” dataset $\hat{D}$ is then obtained by taking D_r and adding a fraction $f_t := \frac{|\hat{D} \cap D_t|}{|D_t|}$ of records from the training set. The leaky risk model allows us to control the direct impact of leaked private information on the privacy quantification method. By inserting risks in controlled unit steps, we can observe whether risk metrics respond to inserted risk linearly. Ultimately leaking the entire original dataset ($f_t = 1$) allows us to observe whether the risk measures reach their theoretical maximum when all original records are leaked.

5.1.2 Overfitting risk model

The leaky risk model provides insight into the response of the assessment frameworks to controlled, deliberately inserted privacy leaks. However, this is not how privacy risks emerge in practice when using synthetic data as a PET. In practice, generator overfitting is considered a main source of risk [58, 26, 9]. When a generator overfits, it might memorize specific patterns from the real data, rather than representing them stochastically. Thereby, they can put involved real individuals at risk by inferring and reproducing their personal information.

To account for risks due to generator overfitting, we proceed as follows. We measure models’ performance through validation loss (see Appendix F). First, we train the generator optimally, giving optimal loss function value L^* . We define the *overfit ratio* f_o . For a given overfit ratio f_o , we then retrain the generative model, but terminate training prematurely once its loss function satisfies $L = f_o \cdot L^*$. This allows us to evaluate the risks in the generator’s output synthetic data when it is trained to a factor f_o of optimality.

When $f_o = 1$, we have $L = L^*$, so no deliberate overfitting took place. When $f_o = 2$, the generator’s loss function value is twice that of its previously determined optimum.

5.1.3 Differential Privacy

A third approach to deliberate risk insertion is through DP. In particular, we vary the privacy budget ϵ of ϵ -DP. This allows us to assess how the privacy quantification methods to DP as a privacy mechanism. We then evaluate the metrics on their output synthetic datasets. By using large ϵ values (up to 100), we aim to simulate conditions where the differential privacy mechanism is nearly disabled, effectively mimicking an overfitting scenario. This allows us to illustrate how model behavior changes as privacy constraints are relaxed.

5.2 Assessing robustness of privacy quantification methods

5.2.1 Outliers

Privacy quantification methods may be sensitive to their parametrization, or to specific data properties, e.g. the presence of outliers [26, 58, 59]. We therefore propose combining the threat models with systematic outlier removal to assess these vulnerabilities. In particular, we propose repeating the experiments after removing $x\%$ of outliers from the original datasets, with x ranging from 0% up to 10%. Generators should then be retrained on the outlier-removed original sets, after which the experiments are repeated. In our experiments, we let x take values of 1%, 2%, 5%, and 10%, showcasing the relation between the number of removed outliers and privacy. We conducted these experiments for overfit ratios of both $f_o = 1$ and $f_o = 1.6$ to compound the effects of outlier removal and deliberate overfitting.

5.2.2 Privacy quantification methods’ parametrization

The quantification methods may also be sensitive to specific parameters. In our experiments, we account for this through the following additional experiments: 1. we assess k -NN-based privacy indicators with different values of k ; 2. for the GTCAP, we repeat the experiments with several radii; Some of the obtained findings seem general in nature and are corroborated by prior work, such as k -NN metrics being optimal with $k = 1$ [49, 2]. For GTCAP, practitioners should make context-sensitive decisions regarding the importance of parameters.

5.3 Accounting for specific versus general inference

5.3.1 On “Relating to” in the GDPR

Understanding what it means for data to “relate to” an individual is central to determining whether data counts as personal under the GDPR and whether synthetic data qualifies as anonymized. The legal literature, including Cesar [23], draws a threefold distinction: data may relate to an individual *in content*, *in purpose*, or *in result*. Each of these has different implications for assessing privacy risks.

“Relating to in content” refers to data that contains specific information about an identifiable individual—for example, “Jan Jansen is poor.” This form is the most direct and is widely recognized as personal data under the GDPR.

“Relating to in purpose” applies when data, though not directly identifying, is with the intention of affecting, evaluating, or making decisions about an individual. For instance, consider synthetic data used to generate a behavioral profile for targeted advertising. Even if the underlying data is anonymized, if the synthetic model is used with the purpose of tailoring content to “a likely diabetic user in postcode 1234”—and this targeting impacts Jan Jansen—it may still count as relating to him “in purpose.” Here, the relevant behavior being influenced is not necessarily the data subject’s own, but rather the behavior of those interacting with or making decisions about the data subject.

“Relating to in result” refers to outcomes that affect individuals based on insights derived from data. For example, a generator trained on real health data may not contain Jan Jansen’s records, but if its predictions lead to changes in insurance policy pricing that impact him specifically, it arguably “relates to” him in result.

These nuances matter because only “in content” data necessarily involves information *specific to an identifiable person*. “In purpose” and “in result” cases often deal with general, population-level insights. Treating these general inferences as personal data risks overextending the GDPR’s scope, and may conflate statistical accuracy with privacy breaches. A synthetic dataset that closely matches population-level statistics is not inherently privacy-violating. However, if a generative model overfits and memorizes information that is highly specific and unique to an individual, then the synthetic output may include traces of “in content” data, amounting to a genuine privacy leak.

To operationalize this distinction, we adopt and build upon baseline approaches from Giomi et al. [10] and Sundaram et al. [60], which allow us to differentiate between information deducible from general patterns and information that can only result from direct memorization.

5.3.2 Anonymizer control set baseline

Giomi et al. [10] propose the use of control sets in attack simulations. This works as follows. Subdivide the original dataset D into a training set D_{train} and a control set D_{control} . Train the generator on D_{train} only. Evaluate the efficacy of the attacks on D_{train} , giving an effectiveness metric value e_{train} . Next, apply the the same attack to the D_{control} , giving effectiveness score e_{control} . This score indicates how well the attack performs against data that was not to train $G(D_{\text{train}})$. As such, efficacy is merely due to the *general* utility of D , not to memorization of $G(D_{\text{train}})$.

Next, compute the proportion E to which the efficacy of the attack on D_{train} is a consequence of information gained through $G(D_{\text{train}})$'s possible memorization of information in D_{train} . This is computed using equation (8), where e^* is the worst-case behavior of e_{train} , that is: the value of e_{train} in case all attacks on D_{train} are successful.

$$E := \frac{e_{\text{train}} - e_{\text{control}}}{e^* - e_{\text{control}}} \quad (8)$$

Attack effectiveness on the control set (e_{control}) is not due to memorization of specific information, but rather due to the degree to which the information of the synthetic data generalizes to its entire underlying population. The use of the baseline thus accounts for the distinction between specific training data and general population information.

5.3.3 Canary records baseline for AIAs

Houssieau et al. [54] and Sundaram et al. [60] note that the success of an AIA may be due to general correlations between attributes in the data (i.e. general patterns), rather than specific learned information about individual records. This is referred to as the **“base-rate problem”**. An individual’s wealth may, for instance, correlate strongly with their education level, age, and place of residence. An AIA can then succeed by exploiting this *general*, population-level correlation, rather than by exploiting specific information memorized or leaked about the individual (*“in content”*) by the synthetic dataset or generator.

Sundaram et al. [60] propose accounting for the degree to which AIAs’ efficacy is due to general correlations rather than memorization as follows. Let D be a real dataset and let t be a target attribute. For target record d , let d' be defined through equation (9).

$$v(d', a) = \begin{cases} v(d, a), & \text{if } a \neq t \\ r, & \text{if } a = t \end{cases} \quad (9)$$

where r is a random value sampled from the uniform distribution over the range of t . Note that this transformation is applied deliberately to the target attribute of each target record. General correlations between t and remaining attributes then have no impact on d' . Let D' be the dataset obtained by taking D and replacing d with d' . Suppose we use $\hat{D}' \sim G(D')$ in an AIA with target record d' and target attribute t . If this AIA is successful, this cannot be due to general correlations, due to the randomization. Thus, if the AIA is still successful, this indicates that the generator may leak *specific* information relating to the individual corresponding to record d *“in content”*. In our experiments, we conduct additional experiments with canary record baselines by randomly selecting 100 canary records per dataset and averaging the success rates of AIAs against these canaries.

6 Experiments

6.1 Synthetic data generation

The experiment with synthetic datasets was conducted using the following methods: RealTabFormer [61], PATEGAN [39], Synthpop [62], AIM [63]. These models were selected to ensure a diverse representation of synthetic data generation approaches. They span a range of methodologies, including classical statistical techniques (Synthpop), deep generative adversarial networks (PATEGAN), and transformer-based architectures (RealTabFormer). In particular, this set includes both differentially private (PATEGAN, AIM) and non-differentially private (Synthpop, RealTabFormer) models, allowing for a comparative evaluation of privacy metrics across fundamentally distinct privacy-preserving paradigms. By selecting models that differ in architecture, training strategies, and privacy mechanisms, our evaluation aims to provide insights that are not limited to a single class of generative models but rather reflect broader trends and trade-offs in the field of synthetic tabular data generation.

RealTabFormer is a transformer-based framework designed for generating non-relational tabular data using an autoregressive model. Additionally, it can generate relational data by leveraging a sequence-to-sequence (Seq2Seq) architecture. In our experiments, we focus solely on the autoregressive model, specifically employing the GPT-2 architecture.

PATEGAN is a generative adversarial network designed for generating synthetic tabular data that uses the PATE mechanism [64]. This mechanism incorporates differential privacy, a formal privacy guarantee that limits the amount of information about individual records in the training data that can be inferred from the output. It ensures that synthetic data generated respects privacy constraints, making it safer to use for sharing or analysis in sensitive domains like healthcare or finance. Synthpop is an R package that generates synthetic tabular data using a sequential modeling approach. It synthesizes variables one at a time, modeling the conditional distribution of each variable based on the variables already synthesized. It uses statistical modeling models like decision trees or regression.

AIM is a differentially private model that follows the select-measure-generate paradigm. This approach involves three key steps: 1. Selection: Choosing queries, typically low-dimensional marginals, from a collection; 2. Measurement: Adding noise to the selected queries to ensure differential privacy; and 3. Query Refinement: Retaining the queries that best explain the noisy measurements. The AIM framework iteratively performs the selection and measurement steps to progressively enhance the quality of the selected queries.

6.2 Metrics evaluation

Below, we detail the specific implementations of privacy quantification methods outlined in Section 4. We restrict the experiments to the metrics most common in practice, namely no-box models. These prevail in 1. **academia**: see the literature reviews discussed in Section 3; 2. **industry**: Ganey et al. [42] outline metrics of twelve synthetic data providers, while [65, 48, 15] describe metrics of one additional providers each. All rely on no-box metrics; 3. **Governmental organizations** such as the European Union *Joint Research Centre (JRC)* and the United States *National Institute of Standards and Technology* use no-box statistical privacy metrics for synthetic data privacy assessment [14, 41]. This restriction aligns with the reasonableness and proportionality principle, as more advanced threat models typically become computationally intensive. Furthermore, the synthetic data provider has direct control over generator access. More advanced threat models also often require highly specific assumptions. For instance, DOMIAS [33] and MAMA-MIA [34] require that the generator is marginals, conditionals, or copulas-based. Table 1 provides an overview of the metrics compared in the experiments.

All metrics are evaluated using the risk models from Section 5. For the leaky risk model, we range the leak fraction f_l from zero (no leaks) to one (all training data leaked) in equal steps of 0.2; for the overfitting model, we range the overfit ratio f_o from one (no deliberate overfitting) to two ($L = 2L^*$) in equal steps of 0.2. For the DP risk model, we use the OpenDP Smartnoise AIM [63] and PATEGAN [39] model to train ϵ -DP synthetic data generators with privacy budget ϵ in $\{1.0, 5.0, 10.0, 50.0, 100.0\}$.

6.2.1 Statistical indicators

Implementation of the IMS is straight-forward. For DCR and k -NN, we use 2% percentiles. We evaluate the Euclidean distance in an embedded space obtained through the embedding methods from [46]. To ensure that the measure yields a value of zero when no information is leaked and a value of one when the entire data is leaked, we normalize the DCR between its best case and its worst case. Thus, a privacy score of zero indicates no perceived risks, and a value of one indicates that all synthetic records pose risk to real records.

6.2.2 MIAs

for the uniqueness-based MIA, we use Anonymeter’s singling out attack [10]. We conduct experiments with outliers with respect to single attributes, as well as experiments in with the number of attributes varying between three and twelve, each consisting of at most¹ 2000 attacks ($N_A = 2000$). The final reported risk is the highest risk R . Bootstrapping with $n = 1000$ is used to measure the variance of the methods.

For MIAs with overfitting detection, we adopt the approach introduced by van Breugel et al. [33]. This method involves fitting two instances of a chosen density estimator (such as the Kernel Density Estimator (KDE) or the Block Neural Autoregressive Flow (BNAF) [68]): one on the reference set and another on the synthetic dataset. For each test point, densities are computed using these fitted estimators. The DOMIAS score is then calculated as the ratio between the synthetic density and the reference density. Finally, the area under the ROC curve (AUC-ROC) is computed to evaluate the attack. The resulting AUC values were normalized to restrict the measure to the interval $[0, 1]$, allowing for a standardized interpretation. The results, presented in Appendix A, demonstrate that the presence of outliers in the original data is crucial for this attack’s success. The attack’s mechanism hinges on the synthetic dataset being overly similar to the real one. When this happens, the synthetic probability density function (PDF) essentially memorizes the training data’s distribution. This memorization, in turn, causes the synthetic PDF to assign high density values to the original training records, while assigning only medium values to records from other sets. Our findings perfectly align with this principle. For instance, the leaky and overfitting setups showed an increasing trend in attack success, which is expected as both scenarios produce synthetic data that is too faithful to the training set. On the other side, the attack’s performance degraded in the outlier removal experiment. This confirms our hypothesis, as the very records that contribute most to the identifiable signal (the outliers) had been removed from the learning set, making the training data much harder for the synthetic PDF to distinguish.

¹The number of guesses is dependent on how many predicates single out unique synthetic records. This may be fewer than 2000 if no more vulnerable records can be detected.

Table 1: Synthetic data degree of privacy quantification methods used in this study. Risk: risk measured or controlled; Aux: auxiliary information; MIA: membership inference attack; AIA: attribute inference attack; disc.: disclosure; SO: singling out; Link: linkability; ML: machine learning; IoS: inference-on-synthetic; LN: local neighborhood. Exactly the attribute disclosure attacks require access to auxiliary information. Nb: examples are for reference only: they may implement same attack methods and mechanisms in different manner than our implementations.

Method	Type	Risk	Mechanism	Aux.	Example(s)	Our implementation
Differential privacy	Generator property	Propensity	NA	No	[63, 37]	OpenDP AIM
Differential privacy	Generator property	Propensity	NA	No	[39]	Synthetic PATEGAN [66]
IMS	Statistical indicator	Propensity	NA	No	[43, 41, 14, 46]	Standard
DCR	Statistical indicator	Propensity	NA	No	[14, 46]	Percentiles-based
k -NN	Statistical indicator	Propensity	NA	No	[14, 46]	Percentiles-based
Outlier-based MIA	Attack	SO	uniqueness	No	[10] (SO), [35]	Anonymeter [10] SO
Outlier-based MIA (DOMIAS)	Attack	SO	uniqueness	Yes	[33]	ROC AUC classifier
Linkability attack	Attack	Link	distance	Yes	[10] (Link)	Anonymeter [10] Link
Classifier inference	Attack	AIA	ML	Yes	[54] (IoS)	XGBoost, accuracy
Regression inference	Attack	AIA	ML	Yes		XGBoost, RMSE
Distance-based AIA	Attack	AIA	distance	Yes	[10] (AIA), [54] (LN), [2, 49]	Anonymeter [10] AIA
GTCAP	Attack	AIA	uniqueness	Yes	[55, 56, 67, 57]	See [55]

6.2.3 AIAs

In our experiments, numerical data was normalized with a MinMaxScaler. Hence, we consistently chose of the radii in the interval $(0, 1)$. In particular, we apply the GTCAP radii of 0.1. We use $A = \mathcal{A}(D) \setminus \{t\}$ to test the worst-case scenario in which the adversary has access to all but one attribute of a given record. To assess distance-based AIAs, we use the Anonymeter inference attack [10] as an instance of distance-based AIAs. Like the singling out attack, this is evaluated in the form of an estimator of guess effectiveness. In our experiments, we use the same parameters as in [10]. To conduct ML-based AIAs, we used the default unregularized *Scikit Learn* implementation of eXtreme Gradient Boosting (XGBoost). For the Adult dataset, we used a classifier, evaluated with accuracy. For the Census and Texas datasets, we used a regressor evaluated using the root mean square error (RMSE). For classification models, we evaluate the accuracy. For regression models, we evaluate the root mean square error (RMSE). To bring RMSE values into a range comparable with accuracy metrics, the normalization shown in equation (10) is applied: the original RMSE is divided by the data range, resulting in values within the interval $[0, +\infty]$, where 0 represents perfect model performance and 1 indicates that the prediction error is comparable to the variability of the original data. Values exceeding 1 suggest poor model performance. An additional normalization factor is introduced to shift the range to $[-\infty, 1]$, where 1 corresponds to perfect model performance.

$$\text{NRMSE}(y, \hat{y}) = 1 - \frac{\text{RMSE}(y, \hat{y})}{\text{range}(y)} \quad (10)$$

Recall, this heuristic is designed to align accuracy metrics for categorical data and regression metrics for numerical data within a similar range of values, where a perfect model is consistently represented by a value of 1 across both types of metrics. However, it is important to recognize the underlying differences between the metrics and avoid overinterpreting normalized values as directly equivalent to accuracy.

For Linkability, we follow use the default Anonymeter implementation [10].

6.3 Datasets and experimental set-up

We tested the methodology on three commonly used public datasets, namely: Adult [69], the Texas Inpatient Public Use Data File (“Texas”) [70], and the 1940 Census full enumeration from IPUMS USA (“Census”) [71]. The adult dataset has 48842 rows and fifteen attributes, of which six are numerical (float), and the remaining nine are categorical. The Texas and Census datasets are considerably larger, at respectively 193 and 97 attributes. We limit our experiments to subsets of 28 and 37 attributes for these sets, respectively. We take randomly selected, mutually exclusive subsets of cardinality 25,000 (Census) and 20,000 (Texas) as training and control sets for the Texas and Census data. See Appendix H for more information. All experiments were conducted on a computer equipped with AMD Ryzen Threadripper 2950X 16-Core Processor (CPU) and 128 Gib of RAM.

6.4 Results

Table 2 contains the results with no inserted risk and maximal inserted risk for each of the metrics, dataset, and risk model. Appendix A shows the full results, as well as the correlation matrices of the privacy quantification methods per experiment for all risk models. Note that while each risk measure is normalized, different methods measure different quantities. Comparisons should therefore be based on the *response to inserted risk*, rather than the exact values of the measurements. Table 3 contains the computation times of the metrics per dataset and risk model.

7 Discussion

We note that two of the generative models in our experiments failed to offer proper utility, namely: PATEGAN, and AIM, (see Appendix G). These generators failed to capture the distributions of their input data, and may have low privacy risks due to their lack of realism. We restrict our discussion to the results of the other generators.

7.1 Quantification approaches

Based on the taxonomy studied in Section 4, we identified three main approaches to privacy risk quantification for synthetic data as an anonymization method: formal properties; statistical indicators; and attack simulations. Out of these only attack simulations involve explicitly modeling adversaries’ threat models. This allows for privacy measurement to involve properties of both the synthetic datasets and of their generators. On the other hand, relying on specific assumptions (on model access, available auxiliary information, attack vectors, etc.) may lead to metrics that do not generalize well to other scenarios and future developments, misaligning with the Reasonableness and Proportionality

Table 2: Results of the leaky and overfitting risk models. RTF: RealTabFormer; O: outlier; D: distance; ML: machine learning. By “no risk”, we indicate that no risk was deliberately added, i.e. $f_l = 0$ for the leaky risk model; $f_o = 1$ for the overfitting risk model; for the DP risk model, we equate “no risk” to a privacy budget of $\epsilon = 0$. Similarly, “max risk” refers to $f_l = 1$; $f_o = 2$; and $\epsilon = 100$. We use the asterisk (*) to denote the maximum risk, which exceeds any risk achieved with the previous values of f_l , f_o , or ϵ .

Method	Leaky			Overfit			DP					
							PATEGAN			AIM		
	Adult	Texas	Census	Adult	Texas	Census	Adult	Texas	Census	Adult	Texas	Census
IMS (no risk)	0.0	0.0	0.0	0.0	0.0	0.0010	0.0	0.0	0.0	0.0	0.0	0.0
IMS (max risk)	1.0	1.0	1.0	0.8684	0.0162	0.9428	0.0	0.0	0.0	0.0	0.0	0.0
IMS stdev (max risk)	0.0045	0.0038	0.0034	0.0034	0.0008	0.0017	0.039	0.0	0.0015	0.0	0.0	0.0
DCR (no risk)	0.0011	0.0082	-0.0013	0.0008	-0.0009	0.0046	0.0002	-0.0203	0.0064	-0.0024	-0.0204	-0.0186
DCR (max risk)	1.0	1.0	1.0	0.5325	0.4783	0.5700	0.2933	-0.0204	0.1115	-0.0204	-0.0204	-0.0046
DCR stdev (max risk)	0.0001	0.0	0.0	0.0011	0.0010	0.0010	0.0032	0.0	0.0005	0.0	0.0	0.0005
MIA, O (no risk)	0.0060	0.0400	0.0101	0.0278	0.0310	0.0270	0.0148	0.0153	0.0285	0.0136	0.0213	0.0137
MIA, O (max risk)	0.9990	0.9990	0.9989	0.7620	0.6744	0.7895	0.4836	0.0257	0.1025	0.0101	0.0239	0.0318
MIA, O stdev (max risk)	0.0010	0.0010	0.0011	0.0204	0.0216	0.0298	0.0272	0.0126	0.0296	0.0173	0.0138	0.0282
DOMIAS, O (no risk)	0.0040	-0.0066	0.0052	0.0966	0.0966	0.0120	0.0162	-0.0088	0.0074	-0.0014	-0.0016	0.0
DOMIAS, O (max risk)	0.5368	0.0312	0.3152	0.3542	0.1444	0.2202	0.2154*	-0.0016	0.0136	0.0024*	-0.0012*	0.0018*
Link. (no risk)	0.0015	0.0116	0.0030	0.0004	0.0111	0.0035	0.0015	0.0015	0.0065	0.0	0.0025	0.0035
Link. (max risk)	0.6433	0.9934	0.6336	0.2854	0.3820	0.2874	0.1589	0.0065	0.0521	0.0005	0.0045	0.0020
Link. stdev (max risk)	0.0211	0.0035	0.0213	0.0199	0.0218	0.0203	0.0161	0.0049	0.0103	0.0017	0.0059	0.0046
AIA, ML (no risk)	0.1570	0.1904	0.0052	0.0077	0.3370	0.0060	0.0887	0.6611	0.0	0.0	0.0	0.0
AIA, ML (max risk)	0.4499	0.9749	0.1535	0.3107	0.7193	0.1733	0.2175	0.4377	0.0337	0.0	0.0	0.0
AIA, D (no risk)	0.0835	0.1312	0.2153	0.0665	0.1262	0.2013	0.1086	0.0887	0.2159	0.0176	0.0177	0.0988
AIA, D (max risk)	0.9922	0.9920	0.9579	0.5945	0.5557	0.6393	0.3621	0.1161	0.2857	0.0241	0.0229	0.1720
AIA, D stdev (max risk)	0.0041	0.0080	0.0390	0.0551	0.0599	0.1221	0.0819	0.0887	0.1629	0.0585	0.0963	0.1968
GTCAP (no risk)	0.0019	0.0005	0.0015	0.0089	0.0	0.0012	0.0817	0.0	0.0424	0.0	0.0	0.0001
GTCAP (max risk)	0.9665	1.0	0.9897	0.5094	0.0156	0.8114	0.4513	0.0	0.1516	0.0	0.0	0.0017

Table 3: Mean of computation times of the various measurements in seconds for the RealTabFormer and AIM models

Method	Adult	Leaky Texas	Census	Adult	Overfit Texas	Census	Adult	DP Texas	Census
IMS	22.50	66.30	43.41	22.50	66.30	43.41	22.50	66.30	43.41
DCR	6.91	134.72	83.35	6.91	134.72	83.35	6.91	134.72	83.35
k -NN	6.91	134.72	83.35	6.91	134.72	83.35	6.91	134.72	83.35
DOMIAS (O)	2.11	46.45	68.45	2.21	45.71	65.45	2.90	47.86	71.45
MIA (D)	6.91	134.72	83.35	6.91	134.72	83.35	6.91	134.72	83.35
Link.	18.06	103.01	99.26	12.44	30.02	63.50	9.72	28.74	63.53
AIA (ML)	0.35	0.58	0.45	0.36	0.54	0.43	0.036	0.55	0.46
AIA (D)	163.23	1380.70	1603.49	74.98	438.33	1117.48	71.99	445.90	1278.39
GTCAP	417.05	2093.73	5580.93	420.77	2115.10	5488.23	435.16	2238.01	5789.22

Principle (cf. Section 2.2.3). This is particularly true for black-box and white-box attacks, which additionally tend to involve unreasonable assumptions in terms of computational resources.

On the other hand, DP and statistical privacy indicators are “threat model-agnostic”: their risk quantification is independent of possible risk exploitations in specific attacks. However, our results indicate that statistical privacy metrics are strongly correlated with metrics based on specific attacks in both idealized and practical risk scenarios. Furthermore, the statistical metrics showcased more robustness (see the standard deviations in Table 2), and are considerably more computationally efficient than attack-based methods (Table 3). Experiments with k -NN metrics consistently indicate that the DCR (i.e. the case $k = 1$) provides the most stringent risk indication (Appendix B). This is consistent with prior research (see [49, 2]) and may explain why the choice for $k = 1$ is common in distance-based attacks (e.g. in [54, 10]).

Under the **leaky risk model**, most metrics increase linearly with f_l , indicating a proportional response to the amount of direct leaks. Furthermore, with few exceptions, measures M satisfy $M \approx f_l$ throughout, indicating that all of the injected risk is detected. In particular, Table 2 shows that most metrics (nearly) achieve their maximum value when $f_l = 1$, except for the linkability attack, the ML-based AIA and DOMIAS. A possible reason why linkability does not attain its largest value is that its efficacy depends on the partitioning of attributes. If partial records have many duplicates, linking remains difficult even with significant leaks. For the ML-based AIA, the behavior is due to the strong performance of the ML models on the control set. The baseline computation thus reduces the overall value of the metric. The efficacy of the DOMIAS attack hinges on two factors: the density estimator’s sophistication and the statistical separation between member and non-member data. A simple estimator may fail to detect the local density spikes caused by model overfitting, leading to a low DOMIAS score due to poor separability. Conversely, if non-member points are highly similar to the training points, occupying the same high-density regions, the computed densities for both groups overlap significantly. This overlap also results in a low DOMIAS score, which, in this context, is interpreted as the synthetic dataset offering strong privacy protection. Thus, a high DOMIAS score requires both a sensitive estimator and discernible differences between member and non-member densities.

Under the **overfitting risk model**, responses to risk between the various methods again follow a similar pattern, resulting in large correlations. With the maximum amount of risk inserted during the experiments ($f_o = 2$). Thus, even with an overfit generator whose validation loss is about twice its minimum, measured risks are about half as severe as in the more theoretical worst-case of the leaky risk model. Correlations between metrics remain strong, but weaker than under the leaky risk model, particularly on the Texas dataset. The relation between overfitting and detected risk underscores that overfitting leads to risks in synthetic data generation. This can be used pro-actively to foster privacy protection through early stopping techniques during generator training.

Under the **differential privacy risk model**, there was little response to an increase in privacy budget. This lack of measured risks makes measured correlations uninterpretable and is likely due to the aforementioned lack of utility of the DP generators’ output (see Appendix G).

Broader implications for anonymization. Taken together, the experimental results underscore that the anonymization status of synthetic data is an empirical matter. It depends on measurable properties of the generator and its outputs rather than on assumptions about synthetic data as such. Consequently, privacy assessment must be dataset-specific and sensitive to both model behaviour and underlying data characteristics. The findings also show that privacy risks may persist even where utility is low, and that high utility does not necessarily imply unacceptable risks. Ensuring alignment with the GDPR’s identifiability standard therefore requires explicit verification that no metric detects leakage consistent

with information “relating to” identifiable individuals in content, rather than reliance on general statistical similarity alone.

7.2 Risk factors and attack vectors

The taxonomy studied in Section 4 also indicates that uniqueness and similarities between real and synthetic records are core privacy risk factors.

7.2.1 Uniqueness

Unique records in real-world datasets are naturally more susceptible to singling out as defined by WP29 (see Section 2.3). Uniqueness is therefore a risk factor, regardless of the anonymization technique. Also, uniqueness is therefore an attack vector in various attack-based privacy metrics, both for singling out and AIA (see Table 4). Statistical privacy metrics typically do not address uniqueness, while DP operates by limiting impact of individual records, thus reducing uniqueness. Interestingly, only the uniqueness-based MIA showed some response to DP privacy budget on the Texas dataset, though this was not very pronounced.

To document the impact of uniqueness on privacy metrics, we conducted experiments with generators trained on original datasets with removed outlier fractions. We did so with both no and considerable deliberate overfitting ($f_o = 1$ and $f_o = 1.6$, respectively), see Appendix C. For the *Texas dataset*, outlier removal had a significant impact on most attacks, which struggled to succeed after the outliers were removed, especially under considerable overfitting ($f_o = 1.6$). On the *Adult* and *Census* datasets, only the DOMIAS attack showed a significant degradation in attack efficacy after outlier removal. The attacks likely succeeded by exploiting the model’s highly specific generations for unique, easily identifiable outlier data points. Consequently, after these outliers were removed from the training data, the attacks failed, as their most obvious targets were gone.

These results corroborate the theory that outliers pose significant privacy hazards, potentially arising from local overfitting around these unique points [72, 48]. The results indicate that this particularly affects MIAs, which is intuitive as these typically identify unique individuals to target. This finding is also consistent with the MIA literature (see, e.g. [26, 35, 73, 33]). The DCR, by contrast, was able to detect risks robustly even after outlier removal. This is evidenced by the consistently higher DCR value with overfit model than with non-overfit model, regardless of the proportion of outliers removed.

Our results indicate that dataset properties can affect the extent to which outlier removal mitigates privacy risks. The *Texas* dataset is heavily categorical (21 out of 28 attributes), with many attributes having few values. This limited combination space may make outliers more obvious targets for a wider range of attacks. Conversely, the *Adult* and *Census* sets have smaller proportions of categorical attributes (9 out of 15 and 11 out of 37, respectively) and broader value ranges, perhaps making their outliers exploitable only by highly specific methods like DOMIAS.

7.2.2 Similarity, ML, and combined factors

Similarity between synthetic and real records can increase susceptibility to all three WP29 attacks. For singling out, if some unique real record does not have a corresponding highly similar synthetic record, then the uniqueness does not present as an attack vector. This is corroborated by our experimental results: uniqueness-based attacks and distance-based metrics are strongly correlated (Appendix A); and outlier removal reduces risks measured through similarity-based metrics for the Texas dataset (Appendix C). Thus, similarity-based metrics can provide direct insight into uniqueness as a risk factor.

Similarity facilitates linkability, as it enables linking synthetic records to given real records, providing insight into the linkability of the training data. For AIAs, similarity can serve as a direct attack vector even in the absence of uniqueness. However, in AIAs against non-outliers, the base-rate problem may be more pronounced. Our experiments with canary records indicate that this problem is significant for non-outliers and may result in an overestimation of AIA risks (Appendix D), particularly when adhering to the legal perspective that only “relating to in content” poses a privacy breach. This was particularly true for the GTCAP, for which the canary baseline experiment indicated that nearly all of the success could be ascribed to the base-rate problem.

Our results also indicate that distance-based privacy metrics strongly correlate with no-box attacks that leverage ML as an attack vector. This has an intuitive explanation, as classifiers or regressors will assign similar values to similar records. In the black-box case, Meeus et al. [35] show how exploiting the combination of uniqueness and similarity can make MIAs possible that would otherwise not meet the reasonableness principle (Section 2.2.3) due to their computational requirements.

7.2.3 Implications for GDPR-Aligned Risk Assessment

The combined results across uniqueness-, similarity-, and ML-based vectors show that privacy risks in synthetic data depend on structural features of the dataset and on generator learning dynamics. Accordingly, whether synthetic data qualify as anonymized under the GDPR hinges on empirically demonstrating the absence of detectable leakage under threat models aligned with the “means reasonably likely to be used” standard. This reinforces a risk-based, rather than assumption-based, understanding of synthetic data as a privacy-enhancing technology.

8 Conclusion

This paper proposed a framework for assessing the efficacy of tabular synthetic data privacy metrics. To showcase the framework, it first provided a survey of existing privacy metrics, covering mathematical privacy properties; statistical privacy indicators; and simulated attacks. The further classification of the latter into singling out, linkability, and inference attacks links them to relevant legal theory on anonymization.

Our framework evaluates metrics from the taxonomy across three novel *risk models*. This enables the empirical measurement of metrics’ responses to deliberately added risk. The risk models respectively insert risk through: 1. direct leaks; 2. generator overfit; 3. varying privacy budgets for DP generators. Through baseline computations, it covers the discrepancy between two types of information inferred by generative models: specific information about the generator’s training data; and general information about the underlying population.

Experiments with the framework showed that no-box privacy quantification methods correlate strongly, indicating that uniqueness-based and similarity-based risks coincide under no-box risk models. This suggests that the choice of a privacy quantification method might be best informed by robustness and efficiency, which would favor statistical indicators. However, such indicators only explicitly target distances between real and synthetic data points. Future research should study the possibility to reliably, robustly, and efficiently quantify uniqueness as a privacy risk factor for synthetic data at the level of probability distributions, following [74].

Overall, the results indicate that synthetic data can only be regarded as anonymized when empirical evidence confirms that the generator does not reproduce information relating to identifiable individuals in content. The proposed framework contributes to this determination by translating GDPR-relevant concepts—such as singling out, linkability, inference, and the “means reasonably likely to be used” standard—into measurable, dataset-specific tests. This facilitates a clearer alignment between technical evaluation practices and legal anonymization requirements, and underscores the need for multi-metric, context-aware assessment strategies rather than reliance on any single indicator or on assumptions about synthetic data’s inherent privacy properties.

References

- [1] Mauro Giuffrè and Dennis Shung. Harnessing the power of synthetic data in healthcare: innovation, application, and privacy. *npj Digital Medicine*, 6, 10 2023.
- [2] Andre Goncalves, Priyadip Ray, Braden Soper, Jennifer Stevens, Linda Coyle, and Anna Paula Slaes. Generation and evaluation of synthetic patient data. *BMC Medical Research Methodology*, 2020.
- [3] Vamsi K. Potluru, Daniel Borrajo, Andrea Coletta, Niccolò Dalmaso, Yousef El-Laham, Elizabeth Fons, Mohsen Ghassemi, Sriram Gopalakrishnan, Vikesh Gosai, Eleonora Kreačić, Ganapathy Mani, Saheed Obitayo, Deepak Paramanand, Natraj Raman, Mikhail Solonin, Srijan Sood, Svitlana Vyetrenko, Haibei Zhu, Manuela Veloso, and Tucker Balch. Synthetic data applications in finance, 2024.
- [4] Michal S. Gal and Orla Lynskey. Synthetic data: Legal implications of the data-generation revolution. *Iowa Law Review*, 109(1087):1087–1154, 2024.
- [5] Jerome P Reiter. Synthetic data: A look back and a look forward. *Trans. Data Priv.*, 16(1):15–24, 2023.
- [6] Ana Beduschi. Synthetic data protection: Towards a paradigm change in data regulation? *Big Data & Society*, 11(1):20539517241231277, 2024.
- [7] Steven M. Bellovin, K. Dutta, Preetam, and N. Reiter. Privacy and synthetic datasets. *Stanford Technology Law Review*, 2018.
- [8] Gillian M Raab, Beata Nowok, and Chris Dibben. Practical privacy metrics for synthetic data, 2024.
- [9] Alexander Boudewijn, Andrea F. Ferraris, Daniele Panfilò, Vanessa Cocca, Sabrina Zinutti, Karel Schepper, and Carlo Chauvenet. Privacy measurement in tabular synthetic data: State of the art and future research directions. 10 2023.

- [10] Matteo Giomi, Franziska Boenisch, Christoph Wehmeyer, and Borbála Tasnádi. A unified framework for quantifying privacy risk in synthetic data, 2022.
- [11] A. Gonzales, G. Guruswamy, and S. R. Smith. Synthetic data in health care: A narrative review. *PLOS Digital Health*, 2(1):e0000082, 2023.
- [12] M. Giuffrè and D. L. Shung. Harnessing the power of synthetic data in healthcare: Innovation, application, and privacy. *NPJ Digital Medicine*, 6(1):186, 2023.
- [13] James Jordon, Lukasz Szpruch, Florimond Houssiau, Mirko Bottarelli, Giovanni Cherubin, Carsten Maple, Samuel N. Cohen, and Adrian Weller. Synthetic data - what, why and how? 2022.
- [14] J. Hradec, M. Craglia, M. Di Leo, S. De Nigris, N. Ostlaender, and N. Nicholson. *Multipurpose synthetic population for policy applications*. JRC technical report, 2022.
- [15] Morgan Guillaudeux, Olicia Rousseau, Julien Petot, Zineb Bennis, Charles-Axel Dein, Thomas Goronflot, Nicolas Vince, Sophie Limou, Matilde Karakachoff, Matthieu Wargny, and Pierre-Antoine Gourraud. Patient-centric synthetic data generation, no reason to risk reidentification in biomedical data analysis. *NPJ Digital Medicine*, (6(37)), 2023.
- [16] Jean-Francois Rajotte, Robert Bergen, David L. Buckeridge, and Khaled El Emam. Synthetic data as an enabler for machine learning applications in medicine. *iScience*, 25(11):105331, 2022.
- [17] A. Boudewijn and A. F. Ferraris. Legal and regulatory perspectives on synthetic data as an anonymization strategy. *Journal of Personal Data Protection Law*, 1:17–31, 2024.
- [18] James Jordon, Lukasz Szpruch, Florimond Houssiau, Mirko Bottarelli, Giovanni Cherubin, Carsten Maple, Samuel N. Cohen, and Adrian Weller. Synthetic data – what, why and how?, 2022.
- [19] Alvaro Figueira and Bruno Vaz. Survey on synthetic data generation, evaluation methods and gans. *Mathematics*, 10(15), 2022.
- [20] Abdul Jabbar, Xi Li, and Bourahla Omar. A survey on generative adversarial networks: Variants, applications, and training, 2020.
- [21] *Opinion 4/2007 on the concept of personal data*, 2007.
- [22] Article 29 Working Party on the Protection of Individuals with regard to the Processing of Personal Data. Opinion 05/2014 on anonymisation techniques, 2014.
- [23] Cesar A.F. López and Abdullah Elbi. On the legal nature of synthetic data. *Proceedings of the Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- [24] Jingchen Hu and Claire McKay Bowen. Advancing microdata privacy protection: A review of synthetic data, 2023.
- [25] Ram Shankar Siva Kumar, David O'Brien, Kendra Albert, Salomé Viljööen, and Jeffrey Snover. Failure modes in machine learning systems, 2019.
- [26] Theresa Stadler, Bristena Oprisanu, and Carmela Troncoso. Synthetic data – anonymisation groundhog day, 2022.
- [27] Lisa Pilgram, Fida K. Dankar, Jorg Drechsler, Mark Elliot, Josep Domingo-Ferrer, Paul Francis, Murat Kantarcioglu, Linglong Kong, Bradley Malin, Krishnamurthy Muralidhar, Puja Myles, Fabian Prasser, Jean Louis Raisaro, Chao Yan, and Khaled El Emam. A consensus privacy metrics framework for synthetic data, 2025.
- [28] Bayrem Kaabachi, Jérémie Despraz, Thierry Meurers, Karen Otte, Mehmed Halilovic, Fabian Prasser, and Jean Louis Raisaro. Can we trust synthetic data in medicine? a scoping review of privacy and utility metrics, 11 2023.
- [29] Pablo A Osorio-Marulanda, Gorka Epelde, Mikel Hernandez, Imanol Isasa, Nicolas Moreno Reyes, and Andoni Beristain Iraola. Privacy mechanisms and evaluation metrics for synthetic data generation: A systematic review. *IEEE Access*, 2024.
- [30] Mikel Hernandez, Gorka Epelde, Ane Alberdi, Rodrigo Cilla, and Debbie Rankin. Standardised metrics and methods for synthetic tabular data evaluation, 09 2021.
- [31] Chao Yan, Yao Yan, Zhiyu Wan, Ziqi Zhang, Larsson Omberg, Justin Guinney, Sean Mooney, and Bradley Malin. A multifaceted benchmarking of synthetic electronic health record generation models. *Nature Communications*, 13, 12 2022.
- [32] Dario Brunelli, Shalini Kurapati, and Luca Gilli. Sure: A new privacy and utility assessment library for synthetic data. In *2024 IEEE International Conference on Blockchain (Blockchain)*, pages 643–648, 2024.

- [33] Boris van Breugel, Hao Sun, Zhaozhi Qian, and Mihaela van der Schaar. Membership inference attacks against synthetic data through overfitting detection, 2023.
- [34] Steven Golob. *Privacy Vulnerabilities in Marginals-based Synthetic Data*. PhD thesis, 2024.
- [35] Matthieu Meeus, Florent Guepin, Ana-Maria Cretu, and Yves-Alexandre de Montjoye. Achilles’ heels: Vulnerable record identification in synthetic data publishing, 2023.
- [36] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [37] Lucas Rosenblatt, Xiaoyan Liu, Samira Pouyanfar, Eduardo de Leon, Anuj Desai, and Joshua Allen. Differentially private synthetic data: Applied evaluations and enhancements, 2020.
- [38] Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative adversarial network, 2018.
- [39] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*, 2019.
- [40] Robin C. Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective, 2018.
- [41] Christine Task, Karan Bhagat, and Gary Howarth. SDNist v2: Deidentified Data Report Tool, March 2023.
- [42] Georgi Ganey and Emiliano De Cristofaro. On the inadequacy of similarity-based privacy metrics: Reconstruction attacks against "truly anonymous synthetic data". *arXiv preprint arXiv:2312.05114*, 2023.
- [43] Gillian M. Raab. Utility and disclosure risk for differentially private synthetic categorical data. In Josep Domingo-Ferrer and Maryline Laurent, editors, *Privacy in Statistical Databases*, pages 250–265, Cham, 2022. Springer International Publishing.
- [44] Sina Rashidian, Fusheng Wang, Richard Moffitt, Victor Garcia, Anurag Dutt, Wei Chang, Vishwam Pandya, Janos Hajagos, Mary Saltz, and Joel Saltz. Smooth-gan: Towards sharp and smooth synthetic ehr data generation. In Martin Michalowski and Robert Moskovitch, editors, *Artificial Intelligence in Medicine*, pages 37–48, Cham, 2020. Springer International Publishing.
- [45] Cristóbal Esteban, Stephanie L. Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans, 2017.
- [46] Daniele Panfilo, Alexander Boudewijn, Sebastiano Saccani, Andrea Coser, Borut Svava, Carlo Rossi Chauvenet, Ciro Antonio Mami, and Eric Medvet. A deep learning-based pipeline for the generation of synthetic tabular data. *IEEE Access*, pages 1–1, 2023.
- [47] Michael Platzer and Thomas Reutterer. Holdout-based empirical assessment of mixed-type synthetic data. *Frontiers in Big Data*, 2021.
- [48] Sobhan Moazemi, Tim Adams, Hwei Geok NG, Lisa Kühnel, Julian Schneider, Anatol-Fiete Näher, Juliane Fluck, and Holger Fröhlich. Nfdi4health workflow and service for synthetic data generation, assessment and risk management, 2024.
- [49] Edward Choi, Siddharth Biswal, Bradley Malin, Jon Duke, Walter F. Stewart, and Jimeng Sun. Generating multi-label discrete patient records using generative adversarial networks, 2018.
- [50] M. Naldi G. D’Acquisto. *Big Data e Privacy by design. Anonimizzazione, pseudonimizzazione, sicurezza*. Giappichelli, 2017.
- [51] Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. Logan: Membership inference attacks against generative models, 2018.
- [52] Dingfan Chen, Ning Yu, Yang Zhang, and Mario Fritz. Gan-leaks: A taxonomy of membership inference attacks against generative models. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, pages 343–362, 2020.
- [53] Benjamin Hilprecht, Martin Härterich, and Daniel Bernau. Monte carlo and reconstruction membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies*, 2019.
- [54] Florimond Houssiau, James Jordon, Samuel N. Cohen, Owen Daniel, Andrew Elliott, James Geddes, Callum Mole, Camila Rangel-Smith, and Lukasz Szpruch. Tapas: a toolbox for adversarial privacy auditing of synthetic data, 2022.
- [55] Rémy Chapelle and Bruno Falissard. Statistical properties and privacy guarantees of an original distance-based fully synthetic data generation method, 2023.

- [56] Mark Elliot. Creating the best risk-utility profile : The synthetic data challenge, 2015.
- [57] Yingrui Chen, Jennifer Taub, and Mark Elliot. The trade-off between information utility and disclosure risk in a ga synthetic data generator. *Joint UNECE/Eurostat Work Session on Statistical Data Confidentiality*, 2019.
- [58] Nicholas Carlini, Matthew Jagielski, Chiyuan Zhang, Nicolas Papernot, Andreas Terzis, and Florian Tramer. The privacy onion effect: Memorization is relative, 2022.
- [59] Hailong Hu and Jun Pang. Membership inference attacks against gans by leveraging over-representation regions. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, CCS '21*, page 2387–2389, New York, NY, USA, 2021. Association for Computing Machinery.
- [60] Meenatchi Sundaram Muthu Selva Annamalai, Andrea Gadotti, and Luc Rocher. A linear reconstruction approach for attribute inference attacks against synthetic data, 2024.
- [61] Aivin V. Solatorio and Olivier Dupriez. Realtabformer: Generating realistic relational and tabular data using transformers, 2023.
- [62] Beata Nowok, Gillian M. Raab, and Chris Dibben. Synthpop: Bespoke creation of synthetic data in r. *Journal of Statistical Software*, 74(11):1–26, 2016.
- [63] Open Differential Privacy (OpenDP). The smartnoise system: Aim, 2024. <https://github.com/opendp/smartnoise-sdk>.
- [64] Nicolas Papernot, Martín Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*, 2016.
- [65] Dario Brunelli, Shalini Kurapati, and Luca Gilli. Sure: A new privacy and utility assessment library for synthetic data. In *2024 IEEE International Conference on Blockchain (Blockchain)*, pages 643–648, 2024.
- [66] Zhaozhi Qian, Bogdan-Constantin Cebere, and Mihaela van der Schaar. Synthcity: facilitating innovative use cases of synthetic data in different data modalities, 2023.
- [67] Jennifer Taub, Mark James Elliot, and Gillian M. Raab. Creating the best risk-utility profile : The synthetic data challenge. 2019.
- [68] Nicola De Cao, Wilker Aziz, and Ivan Titov. Block neural autoregressive flow. In *Uncertainty in artificial intelligence*, pages 1263–1273. PMLR, 2020.
- [69] Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996. DOI: <https://doi.org/10.24432/C5XW20>.
- [70] Center for Health Statistics Texas Department of State Health Services. Texas hospital discharge data public use data., 2005. <https://www.dshs.texas.gov/texas-health-care-information-collection/general-public-information/hospital-discharge-data-public>.
- [71] Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. Ipums usa: Version 8.0 extract of 1940 census for u.s. census bureau disclosure avoidance research. minneapolis, mn: Ipums, 2008. DOI: <https://doi.org/10.18128/D010.V8.0.EXT1940USCB>.
- [72] Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer: Evaluating and testing unintended memorization in neural networks, 2019.
- [73] Zexi Yao, Nataša Krčo, Georgi Ganev, and Yves-Alexandre de Montjoye. The dcr delusion: Measuring the privacy risk of synthetic data. In *European Symposium on Research in Computer Security*, pages 469–487. Springer, 2025.
- [74] Giuseppe D’Acquisto, Aloni Cohen, Maurizio Naldi, and Kobbi Nissim. From isolation to identification. In *Privacy in Statistical Databases: International Conference, PSD 2024, Antibes Juan-Les-Pins, France, September 25–27, 2024, Proceedings*, page 3–17, Berlin, Heidelberg, 2024. Springer-Verlag.
- [75] Milton Nicolás Plasencia Palacios, Sebastiano Sacconi, Gabriele Sgroi, Alexander Boudewijn, and Luca Bortolussi. Contrastive learning-based privacy metrics in tabular synthetic datasets, 2025.

A Complete experimental results




Figure 2: Risk assessment methods evaluated using the leaky risk model




Figure 3: Risk assessment methods evaluated using the overfit risk model - RTF [61]




Figure 4: Risk assessment methods evaluated using the overfit risk model - Synthpop [62]




Figure 5: Risk assessment methods evaluated using the DP risk model - PATEGAN [39]




Figure 6: Risk assessment methods evaluated using the DP risk model - AIM [63]




Figure 7: correlation matrices risk assessment methods using leaky risk model (from left to right: Adult, Texas, Census dataset)




Figure 8: correlation matrices risk assessment methods using overfitting risk model - RTF(from left to right: Adult, Texas, Census dataset)




Figure 9: correlation matrices risk assessment methods using overfitting risk model - Synthpop(from left to right: Adult, Texas, Census dataset)




Figure 10: correlation matrices risk assessment methods using DP risk model - AIM (from left to right: Adult, Texas, Census dataset)




Figure 11: correlation matrices risk assessment methods using DP risk model - PATEGAN (from left to right: Adult, Texas, Census dataset)

B Experiments with k -NN-based indicators

For each dataset and risk model, we experimented with k -NN-based privacy indicators for various values of k . Both the leaky and overfitting risk models show a clear dominance of $k = 1$ over other values.




Figure 12: Results with k -NN-based privacy metric for various k using the leaky risk model




Figure 13: Results with k -NN-based privacy metric for various k using the overfit risk model - RTF [61]








Figure 14: Results with k -NN-based privacy metric for various k using the overfit risk model - Synthpop [62]

C Experiments with outlier removal

We used the local outlier factor (LOF) to embeddings obtained through contrastive learning to identify outliers [75]. Different proportions of outliers were removed to evaluate their impact on the measurements. We used the overfitting risk model, with overfit ratios of 1.0 and 1.6.

For each metric, we applied bootstrap sampling with replacement from the original dataset (1,000 resamples per configuration apart from MIA and GTCAP where we resample 10 times). Metrics were recomputed for each resample, allowing us to estimate confidence intervals. The resulting error bands correspond to 95% confidence intervals. However, we observed that the variance across bootstrap samples is very low, typically on the order of 10^{-3} . As a result, we decided not to plot the confidence intervals, as doing so would add visual clutter without contributing to the interpretability of the figures.

Figure 15: Results with k -NN-based privacy metric for various k using the DP risk model - PATEGAN [39]Figure 16: Results with k -NN-based privacy metric for various k using the DP risk model - AIMFigure 17: IMS, DCR, and MIA with outlier removal in the original dataset prior to generator training, with both no overfitting ($f_o = 1$) and overfitting ($f_o = 1.6$) - RTF [61]Figure 18: GTCAP and ML Inference with outlier removal in the original dataset prior to generator training, with both no overfitting ($f_o = 1$) and overfitting ($f_o = 1.6$) - RTF [61]




Figure 19: Anonymeter’s methods with outlier removal in the original dataset prior to generator training, with both no overfitting ($f_o = 1$) and overfitting ($f_o = 1.6$) - RTF [61]




Figure 20: IMS, DCR, and MIA with outlier removal in the original dataset prior to generator training, with both no overfitting ($f_o = 1$) and overfitting ($f_o = 1.6$) - Synthpop [62]




Figure 21: GTCAP and ML Inference with outlier removal in the original dataset prior to generator training, with both no overfitting ($f_o = 1$) and overfitting ($f_o = 1.6$) - Synthpop [62]




Figure 22: Anonymeter’s methods with outlier removal in the original dataset prior to generator training, with both no overfitting ($f_o = 1$) and overfitting ($f_o = 1.6$) - Synthpop [62]

D Canary Record Baseline

In this section, we compare the canary record baseline [54, 60] with the training set score for evaluating AIAs (see Section 5.3.3) on the *Adult* dataset. We repeated the baseline computations for 100 randomly selected target records. The average AIA success rates using standard inference, machine learning-based inference attack, and GTCAP are presented below.



Figure 23: Comparison between Canary Record Baseline and Risk on the training set using the leaky risk model

The Canary Record Baseline method is indeed better aligned with modern privacy auditing, particularly for synthetic data, where generalization gaps are less indicative of privacy risk. It offers a more realistic and detailed way to measure privacy risks, especially in situations where someone might actively try to extract sensitive information from synthetic data.

E Impact of radii on the GTCAP

The impact of radii on GTCAP was analyzed to understand its effect on the risk measure within the training set, without using the control set as a baseline. In this experiment, a subset of attributes was selected for evaluation.

- Adult dataset: The target variable was ‘income’, while the key variables were ‘workclass’, ‘education’, and ‘marital status’.
- Texas dataset: The target variable was ‘length of stay’, with key variables being ‘illness severity’, ‘pat country’, and ‘sex code’.
- Census dataset: The target variable was ‘incwage’, and the key variables were ‘nchild’, ‘race’, and ‘sex’.

The results indicate that increasing the radii leads to a higher risk. This is because larger radii expand the range of values considered equivalent, thereby increasing the potential overlap and risk.




Figure 24: Results with GTCAP privacy metric for various radii using the overfit risk model

F Validation loss computation

Validation loss in tabular generative models—such as Realtabformer or Synthpop—is commonly used to monitor overfitting during training. These models aim to learn the joint distribution of the data, allowing them to generate




Figure 25: MLE utility scores of synthetic datasets per dataset and generator (utility of training set for reference)

realistic synthetic rows. To assess their generalization ability, we hold out a portion of the dataset as a validation set and compute the loss by measuring how well the model can reconstruct or generate these unseen rows. This involves comparing predicted values against the actual values in the held-out records. If the training loss continues to decrease while the validation loss starts increasing, it indicates that the model is memorizing specific training rows rather than learning a generalizable data distribution. This not only degrades the model’s generation quality on new data but also raises the risk of privacy leakage, as it may inadvertently replicate individual records from the training set. Therefore, tracking validation loss is crucial both for assessing generalization and for ensuring the privacy-preserving nature of synthetic data generation.

G Utility evaluation

To obtain a comprehensive evaluation, we assessed the utility of the generated datasets. In particular, we measured Machine Learning Efficacy (MLE), which evaluates how well a binary classifier can distinguish between real and synthetic data samples. For this, we constructed training and test sets by combining an equal number of real and synthetic records. An XGBoost Classifier was then trained to differentiate between the two. The MLE is then the accuracy of the classifier. An MLE score close to 0.5 suggests that the synthetic data is very similar to the real data and therefore highly useful. On the other hand, a score closer to 1.0 indicates that the classifier can easily tell them apart, which implies low utility. Figure 25 shows the scores per dataset and generator.

In our experiments, only REaLTabFormer and Synthpop produced synthetic data with strong utility scores. By contrast, the differentially private synthesizers scored poorly on utility. This suggests that their high privacy performance may result from generating data that simply lacks the real dataset’s structure.

H Datasets

We use the following datasets: Adult [69], the Texas Inpatient Public Use Data File (“Texas”) [70], and the 1940 Census full enumeration from IPUMS USA (“Census”) [71]. We restrict the Texas and Census dataset to the following attributes:

Table 4: Attribute selection

Texas	Census
ADMIT_WEEKDAY, APR_DRG, APR_MDC, CERT_STATUS, COUNTY, ETHNICITY, FIRST_PAYMENT_SRC, ILLNESS_SEVERITY, LENGTH_OF_STAY, PAT_AGE, PAT_COUNTRY, PAT_STATE, PAT_STATUS, PAT_ZIP, PRINC_DIAG_CODE, PROVIDER_NAME, PUBLIC_HEALTH_REGION, RACE, RISK_MORTALITY, SEX_CODE, SOURCE_OF_ADMISSION, TOTAL_CHARGES, TOTAL_CHARGES_ACCOMM, TOTAL_CHARGES_ANCEIL, TOTAL_NON_COV_CHARGES, TOTAL_NON_COV_CHARGES_ACCOMM, TOTAL_NON_COV_CHARGES_ANCEIL, TYPE_OF_ADMISSION	AGE, AGEMARR, AGEMONTH, BPL, CHBORN, CITIZEN, EDUC, ELDCH, EMPSTAT, FAMSIZE, HIGRADE, HISPAN, HRSWORK1, INCWAGE, IND1950, MARST, MBPL, MOMRULE_HIST, MTONGUE, MBPL, MOMRULE_HIST, MTONGUE, NATIVITY, NCHILD, NSIBS, OCC1950, POPLOC, POPRULE_HIST, RACE, RELATE, RELATED, SEX, SPLOC, SPRULE_HIST, STEPMOM, STEPPOP, SUBFAM, UOCC, WKSWORK1, YNGCH

The resulting datasets used in the experiments are summarized in Table 5.

Table 5: Datasets used in the experiments. #Attributes: number of attributes; #Numeric: number of numeric attributes; #Categorical: number of categorical attributes; #Rows: number of rows

	Adult	Texas	Census
#Attributes	15	28	37
#Numeric	6	6	26
#Categorical	9	21	11
#Rows	48,842	60,000	75,000