

Benchmarking Differentially Private Tabular Data Synthesis

Kai Chen, Xiaochen Li, Chen Gong, Ryan McKenna[†], Tianhao Wang

University of Virginia, [†] Google Research

ABSTRACT

Differentially private (DP) tabular data synthesis generates artificial data that preserves the statistical properties of private data while safeguarding individual privacy. The emergence of diverse algorithms in recent years has introduced challenges in practical applications, such as inconsistent data processing methods, lack of in-depth algorithm analysis, and incomplete comparisons due to overlapping development timelines. These factors create significant obstacles to selecting appropriate algorithms.

In this paper, we address these challenges by proposing a benchmark for evaluating tabular data synthesis methods. We present a unified evaluation framework that integrates data preprocessing, feature selection, and synthesis modules, facilitating fair and comprehensive comparisons. Our evaluation reveals that a significant utility-efficiency trade-off exists among current state-of-the-art methods. Some statistical methods are superior in synthesis utility, but their efficiency is not as good as most machine learning-based methods. Furthermore, we conduct an in-depth analysis of each module with experimental validation, offering theoretical insights into the strengths and limitations of different strategies.

1 INTRODUCTION

Private tabular data synthesis generates artificial data that preserves the statistical properties of real data while protecting individual privacy. This critical problem has a broad range of applications in practice, extending from healthcare [29, 49] to governmental planning [4, 11] and beyond. For instance, in healthcare, there is a need to share patient data for medical treatment while preserving privacy. Similarly, in government, data analysts must analyze sensitive personal attributes, such as gender and disability, while ensuring confidentiality. Addressing this challenge has thus received much attention and led to a growing research area.

Differentiated privacy (DP) has become the gold standard for protecting privacy. DP ensures that the inclusion or exclusion of any single data point does not significantly affect the outcome, thereby protecting each individual data point within a dataset. A substantial body of research has been proposed to address the data synthesis problem with DP. Based on their working principles, they can be broadly classified into two categories: statistical methods and machine learning methods. Statistical methods [23, 35, 38, 55, 62, 63, 67, 67, 69] compress data information through statistical properties, such as low-dimensional data distributions, to achieve data generation. On the other hand, machine learning methods leverage deep learning frameworks designed for data generation, such as generative network [20, 22, 27, 36, 59, 64] and diffusion model [30, 33, 47].

In addition to these efforts to address this problem, many works also focus on providing comprehensive benchmarks. For instance, Du et al. [15] and Tao et al. [57] try to evaluate current methods under the same setting. But their works do not involve recent

algorithms and thus lack completeness. Some other benchmark works [16, 25, 65] focus more on making algorithm analysis and comparison, lacking necessary empirical validation.

In summary, even though several studies have been conducted in this field, we still face several challenges: (1) *Lack of unified evaluation settings*. Beyond algorithmic strategies, evaluation settings, such as data preprocessing, play a significant role in determining algorithm performance. However, these settings are often considered trivial and thus are frequently overlooked in many methods, which potentially leads to unfair comparisons between methods. (2) *Lack of systematic and in-depth analysis*. Current works often focus on proposing new methods, only providing limited intuition about algorithm analysis. Consequently, certain aspects of in-depth analysis, such as analysis for individual algorithm modules, remain under-explored. For example, questions such as how to select marginals more accurately are insufficiently addressed. However, many generative algorithms rely entirely on the selection of marginals, making this result particularly important. (3) *Lack of comprehensive comparison*. Due to various reasons, such as concurrent development or relatively recent introduction, comparisons between recently proposed methods and existing works remain incomplete, particularly between some representative methods like Private-GSD [35] and AIM [39]. Furthermore, current comparisons largely focus on the overall utility of algorithms, with little attention given to experiments analyzing specific working modules.

In light of the above challenges, we believe proposing a new benchmark for evaluating tabular data synthesis is necessary. The contributions of our benchmark work are as follows.

Proposing a Unified Framework for Evaluation. We first propose a generalized framework and align all methods within this framework to ensure fair and objective comparisons. The framework consists of a *data preprocessing module*, a *feature selection module*, and a *data synthesis module*. Notably, this is the first framework to explicitly consider the impact of preprocessing on algorithm comparisons. Moreover, we move forward on the selection and synthesis modules by categorizing them according to their working principle, providing a new perspective to understand them better.

Providing Rigorous Analysis for Current Methods. Given our unified framework, we conduct an in-depth analysis of different modules. We consider different preprocessing methods and their drawbacks and advantages. For the feature selection module, we divide current methods into adaptive methods and non-adaptive methods based on their working principle and formally prove the superiority of introducing scale penalty term and adaptive selection strategy. Finally, we investigate the efficiencies of current synthesis methods and discuss their potential limitations.

Conducting a Comprehensive Comparison. We include current state-of-the-art methods [7, 22, 36, 39, 62, 69] under both statistical and machine learning methods, and newly proposed methods [30, 35] that have not been thoroughly explored. Moreover, our

evaluation is more fine-grained and helps us understand how each module of the algorithms functions independently.

We observed some important experimental findings, i.e., (1) A significant trade-off between utility and efficiency exists among current methods. Two statistical methods, AIM [39] and PrivMRF [7], outperform in utility but show worse time efficiency. Machine learning-based methods, even though relatively inferior in utility, are highly efficient. (2) Preprocessing is crucial for improving algorithm efficiency without significant synthesis errors. It is also algorithm-dependent, with different techniques better suited to specific synthesis methods. (3) All current synthesis modules exhibit limitations in different ways, such as low efficiency or inferior utility. Our source code is available on GitHub.¹

2 PROBLEM FORMULATION

Differential privacy (DP) has become the *de facto* standard for data privacy. It allows aggregated statistical information to be extracted while limiting the disclosure of information about individuals. More formally, the definition of DP is given by:

DEFINITION 1 (DIFFERENTIAL PRIVACY). *An algorithm \mathcal{A} satisfies (ϵ, δ) -differential privacy ((ϵ, δ) -DP) if and only if for any two neighboring datasets D and D' and any $T \subseteq \text{Range}(\mathbf{A})$, we have*

$$\Pr [\mathcal{A}(D) \in T] \leq e^\epsilon \Pr [\mathcal{A}(D') \in T] + \delta.$$

Here, we say two datasets are neighboring ($D \simeq D'$) when they differ on one tuple/sample. To achieve DP in different scenarios, many mechanisms have been employed, such as Gaussian mechanism [41], exponential mechanism [36, 39] and DP-SGD [42]. We provide a detailed introduction to them in the appendix. In our work, we use Rényi DP as a tight composition tool, defined as:

DEFINITION 2 (RÉNYI DP [41]). *We say that an algorithm \mathcal{A} satisfies (α, ϵ) -Rényi DP ((α, ϵ) -RDP) if and only if for any two neighboring datasets D and D'*

$$D_\alpha(\mathcal{A}(D) || \mathcal{A}(D')) \leq \epsilon,$$

where $D_\alpha(Y || N) = \frac{1}{\alpha-1} \ln \mathbb{E}_{x \sim N} \left[\frac{Y(x)}{N(x)} \right]^\alpha$.

RDP has composition and post-processing properties [41], which makes it a suitable choice for complex algorithm design. Moreover, an (α, ϵ) -RDP guarantee can easily be converted to a (ϵ', δ) -DP guarantee via Theorem 1 [41].

THEOREM 1. *If f is an (α, ϵ) -RDP mechanism, then it also satisfy $(\epsilon + \frac{\log 1/\delta}{\alpha-1}, \delta)$ -DP for any $0 < \delta < 1$.*

One promising application of DP is for tabular data synthesis, wherein an artificial tabular dataset is generated that mirrors the statistical characteristics of the original dataset without compromising individual privacy. More formally, assuming that we have a dataset D composed of n records $\{x_1, \dots, x_n\}$, and each record has d attributes $\{A_1, \dots, A_d\}$, we want to generate a dataset D_S similar to D . The two datasets are considered similar based on some similarity metric, such as the ℓ_1 distance or the performance under a downstream task (e.g., answering a range query or training a classification task). We give more concrete metrics in Section 8.

¹https://github.com/KaiChen9909/tab_bench

3 EXISTING WORK

This section examines existing DP tabular synthesis methods and identifies shortcomings in current benchmarks. These limitations inspire the development of our new benchmark.

3.1 Existing Algorithms

Broadly, existing DP tabular data synthesis methods can be divided into two categories, which are statistical methods and machine learning methods.

Statistical Methods. The exploration of statistical methods started earlier. Shortly after the development of DP, researchers began investigating the data generation problem under the DP. MWEM [23] and DualQuery [18] both release data by repeatedly improving an approximated distribution using Multiplicative Weight approach [24]. Another notable line of research involves Bayesian networks, such as PrivBayes [67] and BSG [5]. In addition, Li et al. [32] try to utilize Copula functions for data generation.

In 2018 and 2020, NIST [45, 46] hosted challenges about DP data synthesis. Among the competing algorithms, PrivBayes, MST [38], and DPSyn [34] exhibited the best performance. All of these methods attempt to privately identify and answer highly correlated low-dimensional marginals. Their main differences lie in their methodology for selecting these low-dimensional marginals and in how they represent the data distribution from these noisy marginals. For example, MST synthesizes data using probabilistic graphical models (PGMs) [40], PrivBayes uses a Bayesian model, and PrivSyn iteratively updates an initialized dataset using an algorithm they call GUM.

After these NIST challenges, more advanced statistical methods were proposed. Zhang et al. proposed PrivSyn [69] by organizing and refining DPSyn. Cai et al. introduced PrivMRF [7], and McKenna et al. proposed AIM [39], both of which dynamically select low-dimensional marginals and employ PGMs for synthesis. Moreover, methods such as FEM [63] and RAP/RAP++ [3, 62] treat synthesis as an optimization problem, utilizing FTPL [28, 54, 56] and relaxed projection [44], respectively, to refine initialized datasets using adaptively selected marginals. More recently, Liu et al. [35] proposed Private-GSD, which can apply genetic algorithms to adjust datasets based on any selected marginals iteratively.

Machine Learning Methods. In addition to statistical methods, machine learning models have been widely explored for DP tabular data synthesis. In NIST 2018, there was an effort to utilize GANs to generate data. However, this method, called DP-GAN [64], did not demonstrate a good performance. Further attempts on generative adversarial networks (GANs), such as DP-GAN [64], DP-WGAN [51], DP-CGAN [59], PATE-GAN [27], and DP-CTGAN [17] also demonstrated limited performance before. Therefore, we omit the comparison of this class of methods.

More recent machine learning approaches, including GEM [36] and DP-MERF [22], represent generative network-based advancements. GEM combines generative networks with adaptive marginal selection mechanisms, while DP-MERF employs random Fourier feature loss to train generative networks. Besides, TabDDPM [30] leverages diffusion models' representational power to fit target data directly. While TabDDPM was not originally designed for DP, it achieves state-of-the-art performance among non-DP methods and

Table 1: Summary of benchmarked algorithms. We match different algorithms within our framework and summarize/categorize their working pipelines. Here, ‘unspecified’ means that related information is not mentioned or does not have a deterministic setting in the original paper. The original baselines column lists some other well-received baseline algorithms.

Category	Method	Data Preprocessing	Feature Selection	Data Synthesis	Original Baselines
Statistical Method	RAP [3]	Unspecified	Adaptive	Relaxed Projection	FEM [63], HDMM [37]
	PrivSyn [69]	Categorical Preprocessing	Non-adaptive	GUM	PrivBayes [67], DualQuery [18], PGM [40]
	PrivMRF [7]	Unspecified	Adaptive	PGM	PrivBayes, BSG [5], DP-WGAN [51], DP-Copula [32]
	RAP++ [62]	Unspecified	Adaptive	Relaxed Projection	RAP, DP-MERF, DP-CTGAN [17], PGM
	AIM [39]	Unspecified	Adaptive	PGM	PrivMRF, RAP, MST [38], MWEM [23]
Private-GSD [35]	Unspecified	Unspecified	Unspecified	Genetic Algorithm	GEM, RAP++, PGM [40]
Machine Learning Method	GEM [36]	Unspecified	Adaptive	Generative Network	RAP, MWEM, DualQuery
	DP-MERF [22]	Unspecified	Non-adaptive	Generative Network	DP-GAN [64], DP-CGAN [59]
	TabDDPM [30]	Unspecified	-	Diffusion Model	-

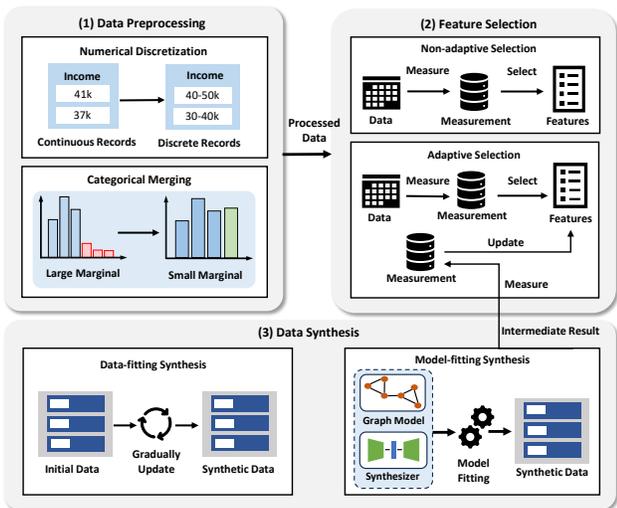


Figure 1: The proposed unified framework. The dataset is first preprocessed and then represented by some selected features. Finally, using the selected features, the synthesis algorithm generates data as the output of the workflow.

can be adapted for DP synthesis using DP-SGD [15]. Consequently, we include TabDDPM in our analysis.

Our work focuses on recently proposed methods that have not been well compared previously and those representing the current state-of-the-art, as summarized in Table 1. Here, we must emphasize that specifically, we don’t consider some LLM-based methods [48, 60]. That is because LLMs are trained on extensive public datasets, which will introduce evaluation bias and affect our further comparison of algorithm modules.

3.2 Existing Benchmark Works

In addition to the proposed algorithms, there are several benchmark studies [15, 16, 25, 57, 65] that investigate the problem of DP data synthesis. However, previous works all have some weaknesses, summarized as follows.

- **Lack of unified comparison setting.** Du et al. [15] have made an experimental evaluation of current works but ignore the importance of a unified setting (e.g., preprocessing), which may significantly influence the comparison fairness.

- **Not include recent advanced works.** Du et al. [15] and Tao et al. [57] focus on utility evaluation, but they both lack investigation for some advanced methods, such as RAP++ and Private-GSD. Fan et al.’s work [16] only focuses on GAN-based methods.
- **Lack of deep analysis.** Yang et al. [65] provide a survey of many methods and further investigate distributed data synthesis. However, they do not delve deep into these algorithms’ working principles. Du et al.’s work and Tao et al.’s work also have a similar weakness in lacking deep algorithm analysis.
- **Lack of comprehensive experiments.** Hu et al. [25] provide analysis for a wide range of DP synthesis algorithms, not only about tabular data synthesis but also trajectory data and graph data. This work, though trying to make an in-depth analysis, lacks comprehensive experiments.

Except for these works, there are other “benchmark-like” works on different aspects of this problem. For instance, Ganey et al. [19] discuss the importance of discretization in tabular data synthesis. Moreover, Stadler et al. [52] focus more on the quantitative evaluation of privacy gain. These works provide in-depth research from different perspectives but lack a more comprehensive viewpoint.

Realizing the weaknesses of current research, our work aims to address these issues by proposing a standardized algorithmic framework (Section 4), providing rigorous analysis (Section 5, Section 6 and Section 7), and conducting detailed experiments (Section 9).

4 FRAMEWORK OVERVIEW

Some previous works [25, 36, 38] have identified that current methods have several common working patterns and proposed unifying algorithmic frameworks, focusing primarily on feature selection and dataset synthesis. However, when dealing with a complex dataset, preprocessing it into a more manageable form should also be a necessary part of the algorithm. As presented Figure 1, our framework extends these approaches by incorporating a data preprocessing module. This section gives an overview of our framework and presents how modules work together to form a complete and cohesive data synthesis pipeline.

Preprocessing. As observed in our evaluation, many datasets contain attributes with large domain sizes (e.g., exceeding 10^5 in Loan dataset [2]), which pose significant challenges for data synthesis. For statistical methods, large attribute domains lead to expansive, low-dimensional marginals, which will introduce excessive DP noise during synthesis and degrade performance. For machine

learning-based methods, a large domain requires a synthesizer with a substantial model size to learn the relationship between different attributes, leading to slower training and higher resource demands. Additionally, larger model sizes require a higher scale of DP noise under the same privacy budget compared to smaller synthesizers, degrading synthetic performance [21].

Therefore, preprocessing is important and non-negligible in algorithm workflows, which can reduce the dimensionality of marginals by effectively merging those with similar characteristics and compressing the domains of the attributes. However, as shown in Table 1, preprocessing is often overlooked by prior works [3, 7, 22, 30, 35, 36, 39]. This omission in the algorithm pipeline may affect the fairness of comparisons between methods.

Feature Selection. Even after effective preprocessing, the domain of the whole dataset increases exponentially with the number of attributes, making methods that rely on a histogram representation computationally infeasible. A practical approach is to utilize some representative local data features, such as *low-dimensional marginals*, to approximate the full joint data distribution [7, 35, 39, 69]. Therefore, the second step in the framework is to measure the representativeness of the features and select necessary features. To further analyze how to better select features in Section 6, we will categorize existing strategies into two categories:

- **Non-adaptive Feature Selection.** A straightforward approach to this step involves performing one-shot feature selection. Some methods, such as DP-MERF and Private-GSD, predefine a fixed set of features without selection. While others, like PrivSyn, measure and select features in one step.
- **Adaptive Feature Selection.** Adaptive methods incorporate iterative calibration to refine the selection process. Typically, these methods begin with an initial feature selection, and then iteratively update the selected features based on intermediate feedback from previous selection steps.

Data Synthesis. The third step in the framework is to synthesize data that aligns well with the features selected in the feature selection step. Current methods apply a wide range of algorithms to achieve this synthesis step. Broadly, there are two approaches: fitting the dataset or fitting the model.

PrivSyn manually adjusts data records to match the marginals, while Private-GSD achieves this by genetic algorithm [53]. RAP and RAP++ use the relaxed projection mechanism to optimize the records of an initialized dataset. These methods are constructed by adjusting records. Some other methods fit models for data generation. PGM [40] is a classical graphical model for tabular data synthesis, which utilizes a tree-like model to represent the distribution of data. Machine learning models such as generative networks and diffusion models have also been applied here.

It is also notable that this division is not mutually exclusive. We can regard an initialized dataset as a model where each data record is a model parameter, so adjusting data can also be regarded as fitting a model. Our characterization is primarily for explaining algorithm intuition.

5 PREPROCESSING

As mentioned in Section 3, when synthesizing complex datasets, we may encounter attributes with high cardinality. For example, an

Algorithm 1: DP Rare Category Merge

Input: dataset D , merge threshold parameter θ , unique value threshold β , DP parameter ρ_2

Output: preprocessed dataset D

- 1 $V_c \leftarrow$ categorical variables with domain size $\geq \beta$;
- 2 $\rho' \leftarrow \rho_2/|V_c|$;
- 3 **for** $j \in V_c$ **do**
- 4 $b \leftarrow$ 1-way marginal of attribute A_j ;
- 5 $\hat{b} = b + \mathcal{N}\left(0, \frac{1}{2\rho'}\right)$;
- 6 **for** $i = 1 : |\hat{b}|$ **do**
- 7 $\theta' \leftarrow \max\left\{\theta \cdot \sum \hat{b}, 3\sigma\right\}$; \triangleright Merging threshold
- 8 **if** $\hat{b}[i] < \theta'$ **then**
- 9 replace $\hat{b}[i]$ with the rare encoding value;
- 10 **end**
- 11 **end**
- 12 **end**
- 13 **return** D

address or income attribute can have thousands of distinct values. Conducting statistics on such attributes is infeasible due to the high memory requirements and long execution times. Therefore, a preprocessing step is necessary for algorithm comparison. We surveyed some previous works [19, 38, 69] and summarized the data types requiring preprocessing into two types: categorical and numerical. Before introducing this section, we need to clarify a basic assumption: domain information is considered public knowledge. This assumption is reasonable in many cases. For example, the domains of personal attributes in census data are well documented on the IPUMS website [26].

5.1 Categorical Attributes Preprocessing

Some prior works [38, 69] have proposed a 3σ merging strategy to preprocess categorical variables, where categories with counts below 3σ (where σ represents the DP noise standard deviation in the counting process) are combined.

By controlling each category’s frequency to be large enough (larger than 3σ), this approach helps mitigate the impact of noise. However, this method has limitations: when we have a large privacy budget, 3σ could be a small value. If we continue using 3σ as the threshold for combining, we may be able to achieve high accuracy, but we cannot reduce the attributes’ domain size to ensure algorithm efficiency.

In response to these limitations, we improve the 3σ merging method, as shown in Algorithm 1, by applying a dual merging threshold $\max\{3\sigma, \hat{n}\theta\}$. Here \hat{n} is the privately measured number of records in the dataset and θ is the threshold parameter. By introducing a fixed threshold, we can avoid the case when σ is too small to reduce the attribute’s domain complexity.

5.2 Numerical Attributes Preprocessing

Continuous numerical variables often exhibit dense distributions within specific intervals, with numerous unique values. Different

Algorithm 2: PrivTree Binning

Input: dataset D , unique value threshold β , divide parameter θ , privacy parameter ρ_1

Output: preprocessed dataset D

```
1 Set  $T = \emptyset$ 
2  $V_n \leftarrow$  numerical variables with domain size  $\geq \beta$ 
3  $\beta_0 \leftarrow 2$ 
4  $\lambda' \leftarrow \frac{2\beta_0-1}{\beta_0-1} \cdot \sqrt{\frac{|V_n|}{2\rho_1}}$ 
5  $\delta' \leftarrow \lambda' \cdot \ln \beta_0$ 
6 for  $j \in V_n$  do
7    $\mathcal{T} \leftarrow \text{PrivTree}(D[j], \lambda', \delta', \theta)$ 
8   Append  $\mathcal{T}$  to  $T$ 
9 end
10 Apply  $T$  to discretize dataset  $D$ 
11 return  $D$ 
```

from categorical attributes, these unique values have numerical correlation, making value combining (like what we do for categorical attributes) impossible. Thus, discretization is a proper preprocessing method for numerical attributes. Here, we outline two potential discretization approaches:

Uniform Binning. Some previous works [12, 19] use uniform binning for continuous data preprocessing. It partitions an attribute’s domain into equal-length intervals, relying only on the attribute’s domain range and a predefined number of bins. Formally, this method can be expressed as Uniform Bin(x) = $\lfloor \frac{x-x_\ell}{h} \rfloor$, where x_ℓ is the lower bound of the attribute’s domain, and h is the length of the uniform interval determined by the bin number.

Uniform binning is advantageous because it requires no detailed data information, avoiding the need for additional privacy budget allocation. However, it has significant drawbacks, particularly for attributes with uneven distributions. For example, when data is highly concentrated around specific values, uniform binning can lead to inefficient binning, as it may allocate unnecessary bins to sparsely populated areas.

PrivTree. A limitation of uniform binning is its reliance on a predetermined number of bins, which introduces concerns about hyperparameter selection. To address this, PrivTree decomposition [57, 68] can be used as a self-adaptive discretization method.

PrivTree employs a tree structure to iteratively divide the domain of an attribute, with splits continuing until intervals contain only a small number of records. However, since PrivTree utilizes sensitive data for domain division, it requires a fraction of the privacy budget to guarantee differential privacy. We use PrivTree on multiple attributes whose domain size is larger than a threshold, with algorithm details in Algorithm 2. We provide the proof of DP guarantee of this algorithm in the appendix.

6 FEATURE SELECTION

Selecting features determines the performance of many algorithms. Similar to preprocessing methods, it is important for statistical methods, while some machine learning methods can automatically learn the characteristics of the dataset by training models. In this

section, we will briefly introduce the currently proposed selection methods and analyze them.

6.1 Existing Selection Methods

Non-adaptive Feature Selection. Non-adaptive methods perform all feature computations and operations at the beginning of the algorithm based on the characteristics of the marginals. The selected features will then be delivered to the synthesis modules without any further refinement.

Some algorithms directly predefine a set of features instead of paying attention to selecting representative features. DP-MERF [22] leverages random Fourier features to capture correlations among numerical variables, while employing 2-way marginals for categorical variables. In addition, some methods with strong fitting ability can work on all two-way marginals. For example, Liu et al. [35] conduct experiments on all two-way marginals to demonstrate the performance of Private-GSD.

Instead of predefining some features, PrivSyn [69] selects the most highly correlated marginals while respecting the privacy budget in one round. They measure each marginal using metric $\text{InDif}_{i,j} = |M_{i,j} - M_i \times M_j|$, where M denotes the attribute marginal. The selection process involves minimizing the expected error

$$\sum_i (N_i x_i + \text{InDif}_i (1 - x_i)),$$

where $x_i \in \{0, 1\}$ denotes the selection decision and N_i is DP noise.

Adaptive Feature Selection. Different from non-adaptive methods, adaptive methods continuously update feature selection with feedback from the previous selection steps. In each synthesis round, adaptive feature selection conducts many queries on the current estimation, and the features with large errors are selected. These features will be used for the selection of future rounds. This strategy is used by methods like RAP [3], RAP++ [62], GEM [36], PrivMRF [7] and AIM [39]. These methods differ in several key aspects:

- First, in terms of initialization, PrivMRF uses a carefully designed criterion to select a small set of features as the starting point, while AIM initializes with all 1-way marginals. In contrast, RAP, RAP++ and GEM do not specify any initialization.
- Secondly, the selection criteria also differ among these methods. Both AIM and PrivMRF’s marginal selection criteria include a punishment term proportional to the marginal scale, allowing for a balance between noise level and feature representativeness. Other methods measure the features without a penalty term, which is one-sided and potentially introduces more errors.
- Finally, the selection mechanism varies: AIM and GEM utilize the exponential mechanism [8], whereas RAP and RAP++ employ the Gumbel mechanism [3], allowing them to select more than one feature in each round; PrivMRF takes a different approach, directly adding Gaussian noise to the selection criteria for selection and select the largest one.

Compared to the non-adaptive methods, the adaptive methods require multiple rounds of data synthesis or computation during the feature selection. Therefore, the efficiency of non-adaptive methods heavily depends on the efficiency of the data synthesis algorithms. Then, a natural question arises: which selection methods are theoretically better in utility? Or do the non-adaptive methods, which

come at the cost of higher computational complexity, ensure continuous optimization in the correct direction during feature selection? We will discuss this in detail in the next section.

6.2 Analysis for Selection Algorithms

Some previous studies [7, 36, 39] have discussed the important properties of a good selection mechanism. In this subsection, we formally investigate some aspects of it. Since all methods, except for DP-MERF, which predefines features without selection, focus on marginal selection, we only discuss marginal selection here.

Importance of Scale Penalty Term. For any marginal M , we have two choices: choose and privatize it or do not choose it. Let the marginal estimated by available information be M_0 , and the privatized marginal be \hat{M} . We have $\hat{M} = M + \mathcal{N}(0, \sigma)$, where σ is known privacy budget. The expected error can be derived as,

$$\begin{cases} \text{Selection error: } \|M - \hat{M}\|_1 = n\sigma\sqrt{2/\pi} \\ \text{Unselection error: } \|M - M_0\|_1. \end{cases}$$

Therefore, we can compare these two errors, denoted as $\|M - M_0\|_1 - n\sigma\sqrt{2/\pi}$, to determine the selection result, which demonstrates the importance of scale penalty term. This equation is also how some selection criteria involve the scale penalty term.

Superiority of Adaptive Selection. Without loss of generality, we consider the case of selecting 2-way marginals. Assuming that before selecting (A_i, A_j) , we have already fitted marginals (A_i, A_1, \dots, A_k) and (A_j, A_1, \dots, A_k) , so that we can use this intermediate knowledge to estimate the distribution of (A_i, A_j) as

$$\hat{\Pr}[A_i, A_j] = \sum \Pr[A_1, \dots, A_k] \cdot \Pr[A_i|A_1, \dots, A_k] \Pr[A_j|A_1, \dots, A_k]. \quad (1)$$

Before selecting the marginal (A_i, A_j) , the non-adaptive methods do not use intermediate results and can only use independent 1-way marginals to measure its representativeness. This independent measurement can be written as $\mathbb{D}_{\text{KL}}(\Pr[A_i, A_j] \parallel \Pr[A_i] \Pr[A_j])$, where \mathbb{D}_{KL} is the KL Divergence [31]. The adaptive methods, because they select features based on intermediate synthesis results, will have a conditional estimation as $\mathbb{D}_{\text{KL}}(\Pr[A_i, A_j] \parallel \hat{\Pr}[A_i, A_j])$, which is the true KL divergence when choosing marginal (A_i, A_j) . Here $\hat{\Pr}[A_i, A_j]$ is defined in Equation (1). Now we give the following theorem to prove the superiority of adaptive selection.

THEOREM 2. *For any pair of attributes (A_i, A_j) , the KL divergence error of conditional estimation is always no larger than that of independent estimation. Formally, we have*

$$\mathbb{D}_{\text{KL}}(\Pr[A_i, A_j] \parallel \hat{\Pr}[A_i, A_j]) \leq \mathbb{D}_{\text{KL}}(\Pr[A_i, A_j] \parallel \Pr[A_i] \Pr[A_j])$$

Here $\hat{\Pr}[A_i, A_j]$ is defined in Equation (1).

The proof of Theorem 2 is deferred to the appendix. In essence, this theorem demonstrates that non-adaptive methods tend to over-estimate the representativeness of features under KL divergence, whereas adaptive methods can correct this error by leveraging intermediate results.

7 DATA SYNTHESIS MODULE COMPARISON

Previous sections have shown that current solutions utilize various techniques to generate data on selected features. This raises critical questions about their utility and efficiency. Thus, we will analyze this problem in this section. For data-fitting methods, there are GUM, Genetic algorithm, and Relaxed Projection. For the model-fitting type, we consider PGM and (deep) generative network.

GUM. GUM, used by PrivSyn, is an iterative adjustment method that modifies values in the initial dataset to align with the selected marginals. We assume that GUM merges marginals to k cliques of sizes $\{c_1, \dots, c_k\}$. Since the maximum number of operations to fit each value in the marginals is no more than the size of the synthetic dataset, the time complexity of GUM is $\mathcal{O}(\sum_{i=1}^k T c_i n)$, where T is the number of update iterations, n represents the synthetic dataset size, respectively. Because GUM strictly controls the clique size c_i , ensuring we have small cliques and simplifying the update process (e.g., we can choose a small T to reach convergence). This guarantees the efficiency of GUM. However, it fits marginals one by one, overlooking overall correlations within the dataset, which may limit its utility.

Genetic Algorithm. Genetic algorithms [53] use mutation and crossover operations to adjust datasets. The complexity depends on the number of mutations and crossovers per iteration. For instance, the algorithm by Liu et al. [35] has a complexity of $\mathcal{O}(T(P_m + P_c))$, where T is the number of iterations, and P_m and P_c represent mutations and crossovers per iteration, respectively. The execution time of genetic algorithms is strongly influenced by the number of tuning rounds, which often exceeds the dataset size when handling complex datasets with diverse values.

Relaxed Projection. Relaxed projection mechanism [3, 62] treats the dataset as a trainable model, optimizing it to match marginals. Given T optimization rounds, synthetic data size n , and data dimension d , the time complexity is $\mathcal{O}(Tnd)$. This method can be regarded as both a model-fitting and data-adjusting approach. The complexity is driven by the size of the synthetic data and the number of optimization rounds. High-dimensional datasets with large attribute domains can result in an inflated encoded data dimension, making optimization difficult to converge and more time-intensive.

PGM. PGM [40] constructs a junction tree of marginal cliques C_1, C_2, \dots, C_k , ensuring that the intersection set S_i of any two cliques appears only in those marginals. The overall distribution is approximated as:

$$\Pr[A_1, \dots, A_d] \approx \Pr[C_1] \cdot \prod_{i=2}^k \Pr[C_i \setminus S_i \mid S_i]. \quad (2)$$

Let T be the number of training iterations, and let k cliques have sizes $\{c_1, \dots, c_k\}$. The total complexity is $\mathcal{O}(\sum_{i=1}^k T c_i + nk)$, which includes model training and data generation complexity. PGM's efficiency depends on constructing reasonable cliques, which are automatically determined by the junction tree. Densely selected marginals can lead to large cliques, greatly increasing time costs.

Generative Network. Generative networks [36] rely on parameters m , batch size b , and training iterations T . The training complexity is $\mathcal{O}(Tmb)$. Generating n synthetic records results in a total

complexity of $O(Tmb + mn)$. The efficiency is affected by network design and training hyperparameters, which can be challenging for high-dimensional datasets.

In summary, we find that model-fitting type methods heavily depend on the construction of the generative model. While this approach can achieve high efficiency, it also risks encountering the curse of dimensionality. In contrast, data-fitting methods are often limited by the complexity of the data itself, which can significantly impact algorithm efficiency.

8 EXPERIMENTAL SETUP

Our evaluations aim to (1) compare all well-established methods within our unified framework; (2) explore and verify the importance of preprocessing in DP tabular synthesis tasks; (3) investigate feature selection and synthesis modules of these methods for a more fine-grained comparison. The comparison results can validate previous theoretical analysis and guide method selection for different modules within the framework.

To achieve this, we design three main experiments as follows, (1) *Overall Evaluation*: We evaluate method performance across metrics and datasets; (2) *Preprocessing Investigation*: These experiments focus on the preprocessing investigation by comparing preprocessed datasets with raw ones; (3) *Module Comparison*: We evaluate the effectiveness of different feature selection and synthesis modules by reconstructing algorithms using them.

8.1 Datasets

The selection of datasets is crucial in comparison. To make a comprehensive comparison, we would expect that (1) the datasets should vary in record size, attributes’ domain sizes, and attribute distribution to better demonstrate the capacity of methods in synthesizing datasets of different complexity; (2) the proportion of numerical attributes and categorical attributes should be different to evaluate different preprocessing methods better. Therefore, we choose five datasets used in previous work [30, 35, 39, 69] as our datasets, which are ACSincome (INC), ACSemploy (EMP), Bank (BK), Higgs-small (HIG) and Loan (LN). The detailed information about these datasets is provided in the appendix.

8.2 Implementations

In our experiments, we consider PrivSyn, PrivMRF, RAP++, AIM, Private-GSD, GEM, DP-MERF, and TabDDPM. We do not include RAP in the experiments because RAP++ directly improves upon it. Moreover, notice that Private-GSD is a synthesis algorithm, we use its one-shot 2-way marginal version in the overall evaluation, which is also used in their original work, and pay more attention to its performance in module comparison. Finally, we train TabDDPM under DP-GSD by opacus [66]. We repeat all evaluations five times and report the average results. The DP parameter δ is set to be 10^{-5} by default. PrivSyn and AIM are executed on CPUs, while the other methods are executed with GPUs. We by default employ uniform binning and rare category merging preprocessing methods. The detailed hyperparameters setting of studied methods, including preprocessing algorithms, is provided in the appendix.

8.3 Evaluation Metrics

Various evaluation approaches have been proposed [15, 39, 69]. However, most current works focus on comparing synthesis utility, or how similar synthetic data is to real data. While utility is important, algorithm efficiency is also a critical factor in algorithm selection. Therefore, we mainly consider the following metrics.

Machine Learning Efficiency (higher is better). A widely accepted metric for evaluating generated data is its performance on downstream tasks. A common approach involves training Machine Learning (ML) models on the generated data and assessing their performance on test data.

The previous works [15, 30, 69] typically select one or more ML models as downstream tasks. It’s important to note that a large number of models do not necessarily lead to a fairer evaluation. Generally, simpler models have weaker data-fitting capabilities, which cannot reach fair conclusions. In our comparison, we therefore selected four machine learning models known for strong performance across various datasets: MLP, CatBoost, XGBoost, and Random Forest. We report the average F1 score on held-out test data as our metric value, and other related metrics, such as AUC and accuracy, are provided in the appendix.

Another justification is that we use the test dataset instead of the training dataset for evaluation for this and the remaining metrics. While both approaches have been employed in previous works, we opt to use test data for evaluation due to the belief that it better reflects an algorithm’s generalization ability.

Query Error (lower is better). Making queries [9] is a commonly used data analysis technique, which can also be conducted to measure relatively high-dimensional similarity due to its high efficiency. Here, we consider using the 3-way marginal query method employed by Du et al. and McKenna et al. [15, 40], which utilizes the statistical ℓ_1 error of frequency query result to reflect the magnitude of the error. Formally, the query error can be expressed as, $\mathbb{E}_{r \in R} |q_r(D_{syn}) - q_r(D_{test})|$, where q_r refers to the query function, which is a combination of range query (for numerical attributes) and point query (for categorical attributes). \mathbb{E} is the mathematical expectation, and R refers to the set of all 3-way marginals.

Fidelity Error (lower is better). Marginal Fidelity is precise in evaluating low-dimensional similarity, such as average 2-way marginal discrepancy. For example, the work by Du et al. [15] suggests using Wasserstein distance as the fidelity measurement, but this method can be infeasible in terms of computation time when dealing with complex data. Alternatively, total variation distance (TVD) can be used for measurement, which has also been utilized in some studies [57, 60]. We define the TVD as $\frac{1}{2} \sum_{1 \leq i \leq j \leq d} |M_{i,j}^{syn} - M_{i,j}^{test}|$, where $M_{i,j}^{syn}$ and $M_{i,j}^{test}$ are the real 2-way marginals determined by the synthetic dataset and test dataset, respectively.

Running Time (lower is better). A straightforward measurement of algorithm efficiency is the execution time when generating the same amount of data. The running time does not include the preprocessing step and only counts the time spent in feature selection and data synthesis modules.

In addition to our primary evaluation criteria, we also use some secondary metrics, such as marginal size. These metrics are simple enough, so we omit the discussion of them here.

Table 2: Overall utility of synthetic data under different methods. The results with the best performance are highlighted in bold (due to the limited number of digits displayed, some data may appear equal in the table, but there is actually an order in their values). Ground Truth is obtained by comparing real data with test data.

Dataset	ACSincome			ACSEmploy			Bank			Higgs-small			Loan		
ML Efficiency \uparrow	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn	0.38	0.39	0.41	0.42	0.43	0.39	0.47	0.47	0.47	0.40	0.43	0.43	0.25	0.26	0.26
PrivMRF	0.73	0.78	0.78	0.72	0.80	0.81	0.62	0.69	0.71	0.50	0.64	0.64	0.52	0.52	0.52
RAP++	0.66	0.73	0.77	0.74	0.77	0.80	0.65	0.69	0.67	0.52	0.53	0.54	0.45	0.42	0.43
AIM	0.76	0.78	0.78	0.78	0.80	0.81	0.67	0.71	0.71	0.63	0.64	0.67	0.52	0.52	0.52
Private-GSD	0.76	0.77	0.77	0.72	0.73	0.72	0.47	0.48	0.49	0.48	0.47	0.49	0.25	0.24	0.25
GEM	0.70	0.68	0.66	0.67	0.70	0.69	0.51	0.56	0.53	0.51	0.52	0.52	0.50	0.49	0.51
DP-MERF	0.65	0.67	0.71	0.58	0.71	0.66	0.60	0.57	0.55	0.53	0.56	0.57	0.32	0.18	0.18
TabDDPM	0.41	0.41	0.39	0.48	0.42	0.51	0.47	0.47	0.47	0.36	0.35	0.34	0.24	0.24	0.24
Ground Truth	0.79	0.79	0.79	0.81	0.81	0.81	0.76	0.76	0.76	0.72	0.72	0.72	0.54	0.54	0.54
Query Error \downarrow	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn	0.003	0.002	0.002	0.006	0.004	0.004	0.007	0.004	0.003	0.009	0.004	0.003	0.006	0.005	0.004
PrivMRF	0.002	0.001	0.001	0.004	0.002	0.002	0.005	0.003	0.003	0.005	0.005	0.003	0.005	0.005	0.004
RAP++	0.019	0.005	0.003	0.029	0.009	0.003	0.014	0.006	0.005	0.035	0.029	0.028	0.020	0.014	0.011
AIM	0.002	0.001	0.001	0.004	0.002	0.001	0.007	0.002	0.002	0.005	0.003	0.003	0.005	0.005	0.004
Private-GSD	0.004	0.003	0.002	0.026	0.026	0.026	0.044	0.044	0.043	0.044	0.044	0.044	0.038	0.037	0.036
GEM	0.014	0.017	0.016	0.010	0.006	0.006	0.118	0.021	0.022	0.066	0.065	0.065	0.030	0.030	0.029
DP-MERF	0.019	0.018	0.024	0.039	0.037	0.036	0.038	0.035	0.036	0.039	0.035	0.034	0.006	0.006	0.006
TabDDPM	0.066	0.064	0.060	0.088	0.067	0.079	0.074	0.071	0.088	0.106	0.106	0.107	0.067	0.066	0.070
Ground Truth	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Fidelity Error \downarrow	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn	0.15	0.12	0.12	0.12	0.09	0.09	0.24	0.10	0.21	0.29	0.20	0.15	0.34	0.35	0.36
PrivMRF	0.11	0.07	0.05	0.07	0.04	0.03	0.13	0.07	0.06	0.21	0.16	0.16	0.31	0.24	0.23
RAP++	0.52	0.24	0.19	0.30	0.13	0.07	0.43	0.36	0.36	0.69	0.65	0.65	0.66	0.58	0.55
AIM	0.09	0.06	0.05	0.05	0.03	0.02	0.11	0.09	0.09	0.19	0.17	0.14	0.35	0.32	0.29
Private-GSD	0.23	0.21	0.20	0.22	0.22	0.22	0.52	0.52	0.52	0.65	0.65	0.65	0.67	0.66	0.66
GEM	0.26	0.28	0.27	0.15	0.08	0.09	0.76	0.21	0.23	0.57	0.57	0.36	0.53	0.52	0.52
DP-MERF	0.52	0.51	0.50	0.34	0.34	0.32	0.47	0.48	0.48	0.60	0.56	0.54	0.91	0.92	0.93
TabDDPM	0.78	0.77	0.71	0.60	0.51	0.57	0.79	0.78	0.82	0.70	0.81	0.88	0.95	0.95	0.94
Ground Truth	0.05	0.05	0.05	0.02	0.02	0.02	0.03	0.03	0.03	0.12	0.12	0.12	0.10	0.10	0.10

9 EXPERIMENTAL RESULTS

9.1 Overall Evaluation

Utility Comparison. Table 2 shows the detailed utility metrics of different methods respectively. To make the result more straightforward, we utilize the t-SNE [61], a dimensionality reduction technique, to visualize the synthesis results of all methods on the Bank dataset in Figure 2. By reducing the dimensionality of the dataset with t-SNE technique and plotting the scatter distribution, these visualizations reveal the distribution overlap between synthetic and real data, demonstrating synthesizers’ ability to generate data resembling the original.

The first clear conclusion is that PrivMRF and AIM achieve the best utility metrics among all methods. Another straightforward finding is that statistical methods generally perform better than machine learning approaches. Both AIM and PrivMRF rely on graphical models, and the quantitative evaluations of GEM, DP-MERF, and TabDDPM are worse than most other statistical methods in some metrics, such as RAP++ and PrivSyn in fidelity error.

Utility & Efficiency Trade-off. We present the relationship between utility and efficiency on different algorithms in Figure 3. Although AIM and PrivMRF exhibit excellent utility performance, they require more execution time, representing a trade-off in utility.

PrivMRF is slightly more efficient than AIM because it involves initialized marginal selection, which reduces the number of selection rounds and accelerates the algorithm. Additionally, all machine learning-based methods, despite having lower generation utility, demonstrate high efficiency, partly due to the use of GPUs in the model-fitting steps.

TabDDPM has a weak performance in synthesis utility. Among all the methods, TabDDPM performs poorly in all three dimensions, and in Figure 2, the visualization of the synthesis result does not match the real data distribution. We track the training loss with DP-SGD and compare it with the loss without DP-SGD. In Figure 4a, we can observe that the loss under DP-SGD does not converge well when training. This raises our suspicion of DP-SGD’s efficiency. However, the diffusion model still demonstrates excellent performance when generating images [14, 33] under DP-SGD. Therefore, we believe the failure of TabDDPM is potentially caused by its unsuitability for tabular data synthesis due to tabular data’s high sensitivity to exact attribute values and the relation between records (e.g., overall distribution), which is hard to achieve under DP-SGD.

Two GAN-based methods, DP-MERF and GEM, perform differently. Even though GEM and DP-MERF both use deep generative networks for synthesizing, compared to GEM, DP-MERF

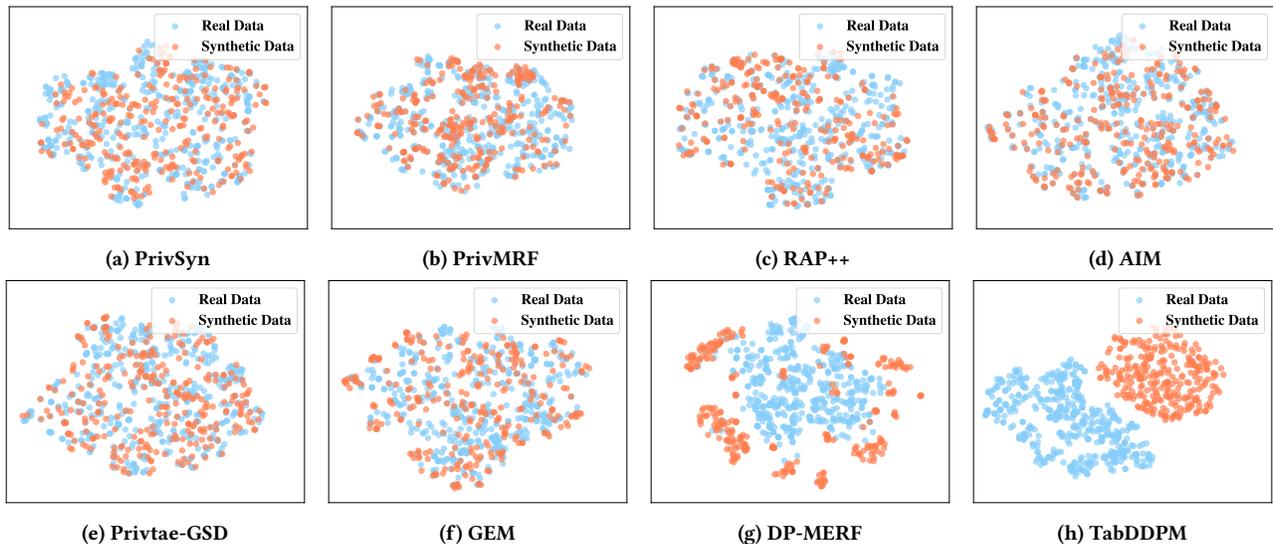


Figure 2: T-SNE scatter plots of synthesis results on Bank dataset under $\epsilon = 1.0$

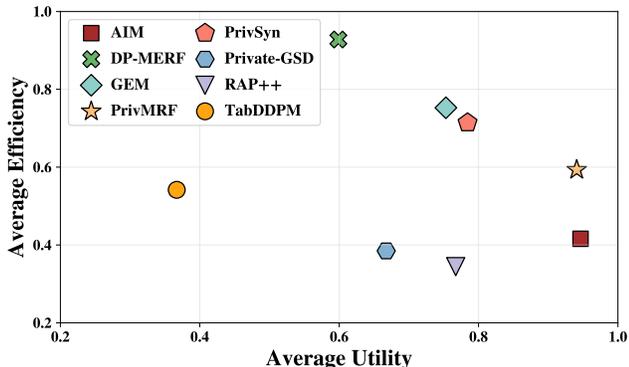


Figure 3: Scaled average utility and efficiency of different algorithms. Utility is obtained by the average performance of three utility metrics and efficiency is measured by average logarithmic running time. All measurements are scaled to $[0, 1]$, and the higher value means better performance.

performs poorly regarding query and fidelity error. One key reason lies in the construction of features. Firstly, random Fourier features embedding, used by DP-MERF, is still approximating the joint distribution of numerical attributes, which inherently introduces approximation error. Furthermore, when dealing with categorical variables, DP-MERF focuses solely on the marginal distributions between categorical variables and the label variable, which overlooks other marginals. While GEM employs an adaptive marginal selection strategy, continuously adjusting the fitting target.

To verify our hypothesis, we track the fitting errors of different categories of marginals in Figure 4b and Figure 4c. During the DP-MERF training process, the total variation distance (TVD) of most marginals, except those for “categorical-categorical” marginals, does not effectively converge. This, to some extent, supports our thinking. By comparison, GEM achieves better convergence across

marginals. Moreover, in Figure 2g, we observe that the data points generated by DP-MERF concentrate in several areas, which we believe is caused by unbalanced feature construction. In contrast, as shown in Figure 2f, the data distribution of GEM is better aligned with the target dataset.

PrivSyn works poorly on machine learning efficiency. Another finding is that PrivSyn performs well on range query error and fidelity error metrics, but underperforms in terms of machine learning efficiency. This scenario aligns with our previous analysis in Section 7, which expresses the concern that GUM focuses more on local marginal cliques and may ignore global relevance. The experiments on reconstructed algorithms in Section 9.3 can also support this statement.

9.2 Preprocessing Investigation

In this Section, we will explore the influence of preprocessing on algorithms’ performance, and how different preprocessing strategies will influence the quality of synthesis data.

Influence of Preprocessing on Dataset Information. Directly comparing the algorithms’ performances with and without preprocessing is difficult because running algorithms on raw datasets is often too time-consuming and computationally complex (e.g., some methods even require more than 24 hours to run on raw datasets). Therefore, we consider comparing the marginal sizes and utility metrics calculated on preprocessed and raw datasets. The detailed results are shown in Figure 5.

A straightforward finding is that preprocessing decreases the complexity of data, with the introduction of only a small error. In Figure 5, the average marginal size significantly decreases after preprocessing (from 10^8 to 10^3 in Higgs-small and Loan datasets). Meanwhile, its negative influences on utility are small enough (change on query error < 0.003 , TVD < 0.1). We infer this is because binning can preserve most numerical characteristics, and the low-frequency categorical values only contribute a small proportion of the overall correlation between attributes in the dataset.

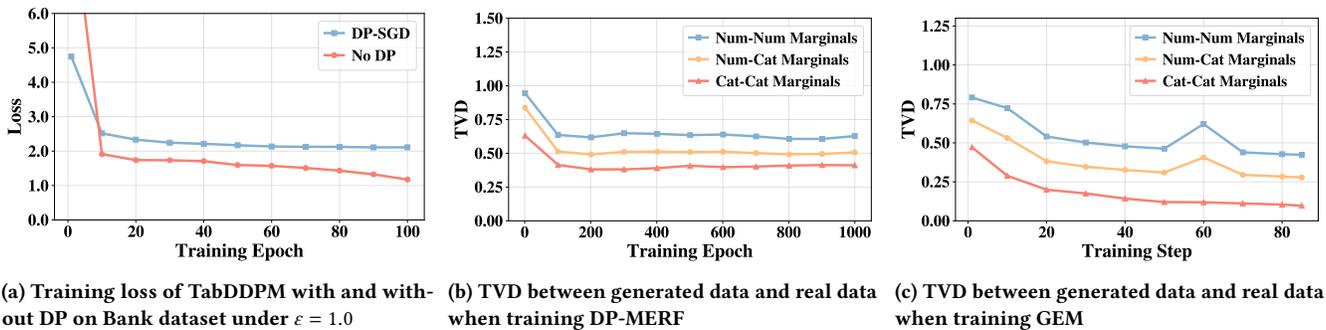


Figure 4: Figures for analyzing different methods in Section 9.1

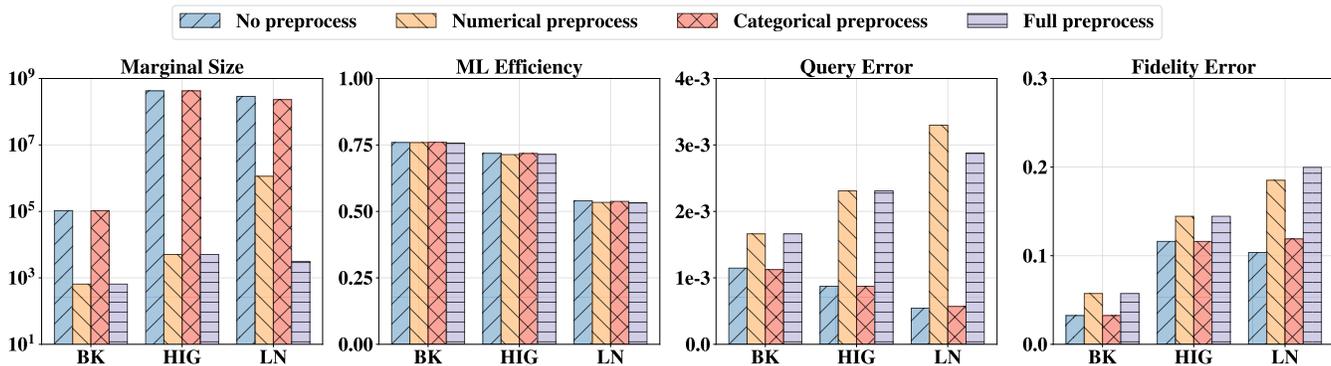


Figure 5: Metrics under different preprocessing settings. These metrics are obtained by comparing preprocessed raw data with test data. Numerical preprocess means only conducting numerical discretization, and categorical preprocess means only conducting categorical rare merging. By default, the results are under the setting that $\epsilon = 1.0$ and 10% of the budget is used for preprocessing (if needed).

Table 3: Preprocessed numerical attributes’ domain sizes under different discretization algorithms. The result is obtained on Higgs-small dataset and under setting $\epsilon = 1.0$.

Dataset	Min Domain Size			Max Domain Size		
	Raw	PrivTree	Uniform	Raw	PrivTree	Uniform
Bank	505	25	100	6024	28	100
Higgs-small	4870	6	100	60696	18	100
Loan	101	9	100	93995	32	100

Numerical Discretization Ablation. Firstly, we summarize the domain sizes under different discretization results in Table 3. It is obvious that PrivTree always generates significantly fewer bins to represent numerical attributes, even for attributes with highly complex value distributions. This approach not only reduces noise in feature measurements but also simplifies the features, raising concerns about losing information. To better demonstrate the influence of discretization, we conduct the ablation study on Higgs-small dataset, which contains the most complex numerical variables among all datasets. For this ablation, we focus on $\epsilon = 1.0$, as is standard practice [15, 57]. The results are shown in Table 4.

In most cases, PrivTree performs better on machine learning tasks but is slightly worse in query errors and fidelity errors compared with uniform binning. We believe fewer bins help capture

overall correlations by introducing less noise, which benefits machine learning tasks. However, query tasks and fidelity are more sensitive to exact values, and more bins lead to higher accuracy. Two exceptions are Private-GSD and GEM, where PrivTree outperforms uniform binning. We hypothesize that this is because PrivTree reduces the dimensionality of variables, thereby facilitating network training and improving the convergence of the genetic algorithm.

These results do not completely align with observations in Tao et al.’s work [57]. We would propose three possible reasons. Firstly, Higgs-small dataset is more value-rich than the datasets used in their work, which can lead to high sensitivity to binning methods. Secondly, the evaluation metrics are different, while ours focus more on measurements on higher-dimensional marginals (e.g., 3-way marginals). This can influence the results. Finally, the investigated algorithms vary in the two works, which may lead to different comparison conclusions. Another unusual finding is that under PrivTree binning, AIM shows a significantly worse runtime. We believe this is because smaller domain sizes will cause AIM to allocate budget to more marginals. This may lead to large cliques, slowing down the execution speed of PGM.

We can conclude that different discretization methods have their own advantages and weaknesses. It is important to choose an appropriate method for different algorithms. For example, for those

Table 4: Marginal-based methods’ performance based on different discretization methods. The result is obtained on the Higgs-small dataset and under $\epsilon = 1.0$.

Method	ML Efficiency \uparrow		Query Error \downarrow		Fidelity Error \downarrow		Running Time \downarrow	
	PrivTree	Uniform	PrivTree	Uniform	PrivTree	Uniform	PrivTree	Uniform
PrivSyn	0.43	0.43	0.005	0.004	0.19	0.20	2 min	6 min
PrivMRF	0.65	0.64	0.005	0.003	0.19	0.16	15 min	7 min
AIM	0.67	0.65	0.005	0.003	0.20	0.19	690 min	18 min
RAP++	0.55	0.53	0.030	0.029	0.55	0.65	71 min	37 min
Private-GSD	0.50	0.47	0.043	0.044	0.57	0.65	57 min	57 min
GEM	0.56	0.54	0.019	0.061	0.29	0.55	0.4 min	6 min

methods based on generative networks, PrivTree could be a potentially better preprocessing method because it can significantly decrease the dimension of models by reducing domain complexity. For PGM and GUM, uniform binning could be a better choice due to their stronger ability to fit marginals with large domains.

Category Merging Ablation. Finally, we briefly discuss the necessity of introducing a fixed merging threshold. We compare the maximum domain size of preprocessed categorical attributes under different fixed merging thresholds in Figure 6. We can conclude that domain size cannot be reduced efficiently under large ϵ if we do not apply a fixed merging threshold (0.0%). Furthermore, we may over-merge categories if this threshold is large (e.g., 1%), which may cause potential synthesis errors.

9.3 Module Comparison

In this subsection, we decompose our experiments into two parts: one focusing on different feature selection algorithms and the other on different synthesis modules. In other words, we fix either the selection or the synthesis approach and then evaluate how different algorithms perform under that fixed condition. Since some features are not compatible with other synthesis methods (e.g., random Fourier features), we only focus on marginal-based methods, namely PrivSyn, RAP++, Private-GSD, AIM, PrivMRF, and GEM.

To make clearer comparisons, we divide the datasets into two groups. The first group, ACSincome and ACSemploy, is relatively low-dimensional with moderately sized attributes’ domains. We call them “small datasets” for description simplicity. The second group, Bank, Higgs-small, and Loan, is higher-dimensional or has larger attribute domains, which are called “large datasets”. To make a better comparison, we scale each metric and display the average performance on these two groups of datasets under different ϵ .

Comparison of Feature Selection Modules. Since some algorithms use the same or similar marginal selection methods, or have not specified them, we consider four marginal selection methods included in previous work. They are: PrivSyn selection, PrivMRF selection, RAP++ selection, and AIM selection. The synthesis processes for these four selection methods are all set to PGM. Figure 7 show the average scaled performances.

Among the selection methods, PrivSyn and RAP++ show weaker fitting utility, consistent with our analysis. PrivSyn fails to use intermediate information during selection, reducing its ability to capture necessary data features and leading to excessive selections. This also overwhelms PGM on larger datasets due to high memory demands. Similarly, RAP++ employs the Gumbel mechanism to select multiple marginals per iteration, without fully accounting

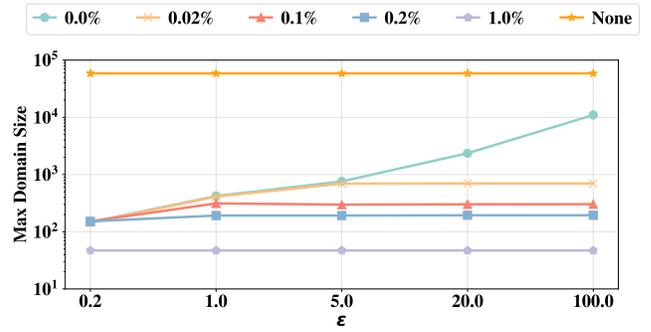


Figure 6: The maximum attribute domain size of Loan dataset under fixed merging thresholds and different ϵ . In the figure, “0.0%” means no fixed threshold (merge only by 3σ criteria), and “None” means no merging preprocessing (raw data).

for intermediate results. It also decreases the number of synthesis rounds required, thereby accelerating the algorithm’s execution. Moreover, its selection criterion focuses solely on marginal query error, ignoring noise error, which is unbalanced and potentially influences the algorithm’s performance.

The utility performances of AIM and PrivMRF, are very similar since both perform a single marginal refinement after updating the intermediate information. A notable observation is that PrivMRF has a shorter running time than AIM. We attribute this to the well-designed initial marginal set in PrivMRF, which decreases the need for further refinement and speeds up the total algorithm.

Comparison of Synthesis Modules. We compare the synthesis algorithms of all six algorithms: GUM, PGM, relaxed projection (RP), genetic algorithm (GA), and generative network (GN). PGM is employed by both AIM and PrivMRF. To avoid multiple fitting steps from adaptive selection and control the running time, we use the PrivSyn selection algorithm as the common marginal selector. The results are shown in Figure 8. GUM and PGM are conducted on CPUs, and the remaining works are on GPUs.

Among these methods, PGM and the generative network achieve the best overall accuracy and stability in fitting features on small datasets. PGM benefits from its expressive structure, which, as indicated by Equation (2), infers marginals without refitting existing ones, thus minimizing compounding errors. However, the PGM’s fitting process is quite slow, due to the densely selected marginals by PrivSyn, causing it to fail to deal with large datasets. Indeed, PGM is most efficient when coupled with a marginal selection procedure that is designed to make the resulting graphical model “tree-like” [7, 38, 39]. The generative network’s strong representational capacity also supports accurate feature fitting on small datasets. However, when it turns to larger and more complex datasets, its superiority in efficiency and utility will be diminished due to the greater difficulty of training high-dimensional models.

In contrast, GUM, the genetic algorithm, and relaxed projection demonstrate weaker fitting capabilities. GUM refines marginals individually, overlooking global correlations. However, it shows a superiority in running time, aligning with our analysis. The genetic algorithm’s reliance on randomness makes it unsuitable for complex, high-dimensional datasets with large attribute domains under

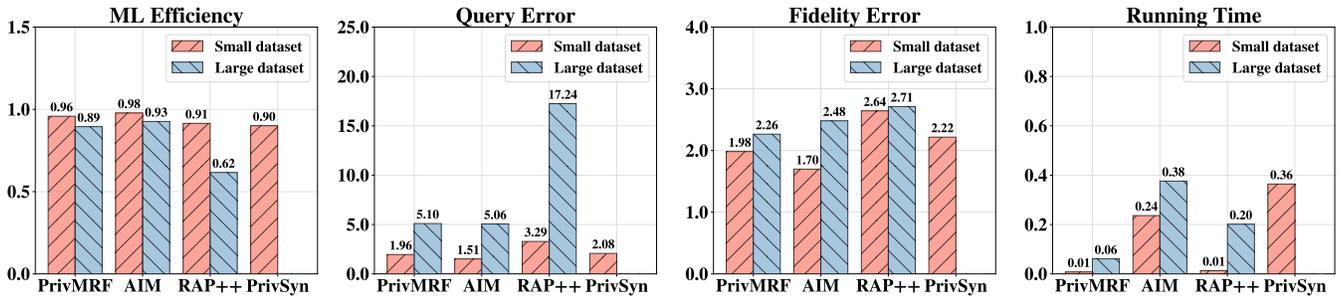


Figure 7: Average scaled metrics for different selection modules on different datasets. We use PGM as the common synthesizer and match PGM with different marginal selection modules. The machine learning efficiency, query error, and fidelity error are scaled using ground truth values provided in Table 2 while running time is scaled using a simple min-max linear normalization. The small dataset includes ACSincome and ACSemploy, whereas the large dataset consists of Bank, Higgs-small, and Loan.

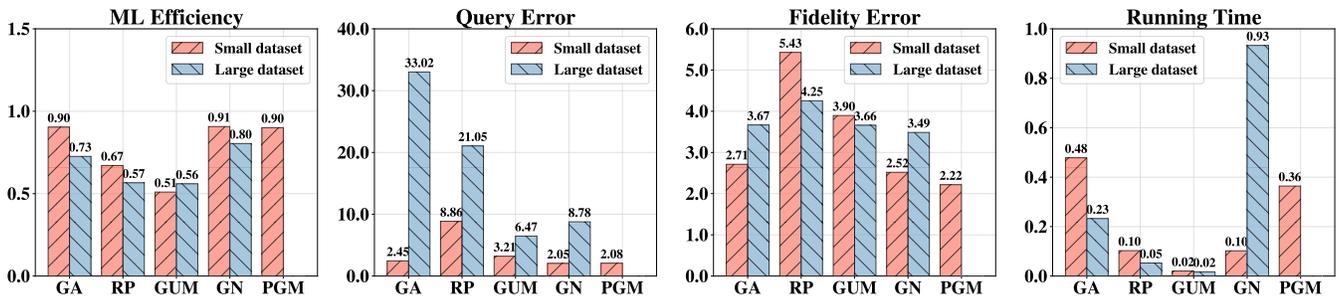


Figure 8: Average scaled metrics for different synthesis modules on different datasets. We use PrivSyn as the common marginal selector and try different synthesizers. The machine learning efficiency, query error, and fidelity error are scaled using ground truth values provided in Table 2 while running time is scaled using a simple min-max linear normalization. The small datasets include ACSincome and ACSemploy, whereas the large datasets consist of Bank, Higgs-small, and Loan.

time constraints, leading to poor performance on large datasets. Relaxed projection suffers similarly. In high-dimensional spaces with extensive attributes' domains, the encoded data dimension becomes excessively large, complicating optimization and making it harder to converge to the optimal point, and finally resulting in poor performance.

10 CONCLUSION AND DISCUSSION

Data synthesis under differential privacy remains a critical and active area of research. Despite a growing number of methods being proposed, the field still lacks a fair and comprehensive comparative analysis. In this paper, we conduct analysis, comparison, and evaluation of several methods on a unified framework. The primary findings of our study are as follows:

- A significant trade-off exists between utility and efficiency, resulting in no single algorithm dominating the others. Machine learning-based methods, though highly efficient, generally exhibit inferior performance compared to traditional statistical methods. Nonetheless, some statistical methods, such as AIM and PrivMRF, are often time-consuming because of their heavy computation requirement, which is a trade-off for superior utility.
- Data preprocessing strategy is important but algorithm-dependent. Data preprocessing is crucial in reducing algorithm complexity

without introducing too much error. The choice of data preprocessing methods should align with the working principles of the algorithm. Appropriate preprocessing techniques can enhance both the efficiency and effectiveness of the algorithms, whereas unsuitable methods may hinder their performance.

- Improvement on data synthesis module remains a promising direction. We find that some feature selection methods outperform others, but existing synthesis modules exhibit significant drawbacks in different ways. For example, PGM, even though demonstrates outstanding utility, cannot work on densely selected marginals (due to high memory requests and long execution time), while generative networks face significant challenges when dealing with high-dimensional data. There remains substantial room for improvement in developing algorithms that can efficiently, reliably, and accurately generate data.

Our work not only sheds light on the strengths and limitations of current methods but also identifies several promising directions for future research. These include tailoring preprocessing strategies to algorithmic needs, developing more effective feature selection techniques, and improving the utility of machine learning models or the efficiency of some statistical methods under differential privacy. Addressing these issues and challenges can help advance this field and better support differential private data analysis.

REFERENCES

- [1] 2016. Higgs particle dataset. <https://www.openml.org/search?type=data&status=active&id=4532>
- [2] 2020. Loan dataset. <https://www.kaggle.com/datasets/devanshi23/loan-data-2007-2014>
- [3] Sergul Aydore, William Brown, Michael Kearns, Krishnaram Kenthapadi, Luca Melis, Aaron Roth, and Ankit A Siva. 2021. Differentially private query release through adaptive projection. In *International Conference on Machine Learning*. PMLR, 457–467.
- [4] Andrés F. Barrientos, Alexander Bolton, Tom Balmat, Jerome P. Reiter, John M. de Figueiredo, Ashwin Machanavajjhala, Yan Chen, Charley Kneifel, and Mark DeLong. 2018. Providing Access to Confidential Research Data Through Synthesis and Verification: An Application to Data on Employees of the U.S. Federal Government. arXiv:1705.07872 [stat.AP] <https://arxiv.org/abs/1705.07872>
- [5] Vincent Bindschaedler, Reza Shokri, and Carl A. Gunter. 2017. Plausible Deniability for Privacy-Preserving Data Synthesis. arXiv:1708.07975 [cs.CR] <https://arxiv.org/abs/1708.07975>
- [6] Mark Bun and Thomas Steinke. 2016. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of cryptography conference*. Springer, 635–658.
- [7] Kuntai Cai, Xiaoyu Lei, Jianxin Wei, and Xiaokui Xiao. 2021. Data synthesis via differentially private markov random fields. *Proceedings of the VLDB Endowment* 14, 11 (2021), 2190–2202.
- [8] Mark Cesar and Ryan Rogers. 2021. Bounding, concentrating, and truncating: Unifying privacy loss composition for data analytics. In *Algorithmic Learning Theory*. PMLR, 421–457.
- [9] Chris Chatfield. 2018. *Introduction to multivariate analysis*. Routledge.
- [10] Imre Csiszár. 1967. On information-type measure of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.* 2 (1967), 299–318.
- [11] Teddy Cunningham, Graham Cormode, and Hakan Ferhatosmanoglu. 2021. Privacy-Preserving Synthetic Location Data in the Real World. In *17th International Symposium on Spatial and Temporal Databases (SSTD '21)*. ACM, 23–33. <https://doi.org/10.1145/3469830.3470893>
- [12] Travis Dick, Cynthia Dwork, Michael Kearns, Terrance Liu, Aaron Roth, Giuseppe Vietri, and Zhiwei Steven Wu. 2023. Confidence-ranked reconstruction of census microdata from published statistics. *Proceedings of the National Academy of Sciences* 120, 8 (Feb. 2023). <https://doi.org/10.1073/pnas.2218605120>
- [13] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. 2021. Retiring Adult: New Datasets for Fair Machine Learning. *Advances in Neural Information Processing Systems* 34 (2021).
- [14] Tim Dockhorn, Tianshi Cao, Arash Vahdat, and Karsten Kreis. 2023. Differentially Private Diffusion Models. arXiv:2210.09929 [stat.ML] <https://arxiv.org/abs/2210.09929>
- [15] Yuntao Du and Ninghui Li. 2024. Towards principled assessment of tabular data synthesis algorithms. arXiv preprint arXiv:2402.06806 (2024).
- [16] Liyue Fan. 2020. A survey of differentially private generative adversarial networks. In *The AAAI Workshop on Privacy-Preserving Artificial Intelligence*, Vol. 8.
- [17] Mei Ling Fang, Devendra Singh Dhami, and Kristian Kersting. 2022. Dp-ctgan: Differentially private medical data generation using ctgans. In *International conference on artificial intelligence in medicine*. Springer, 178–188.
- [18] Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Zhiwei Steven Wu. 2015. Dual Query: Practical Private Query Release for High Dimensional Data. arXiv:1402.1526 [cs.DS] <https://arxiv.org/abs/1402.1526>
- [19] Georgi Ganev, Meenatchi Sundaram Muthu Selva Annamalai, Sofiane Mahiou, and Emiliano De Cristofaro. 2025. The Importance of Being Discrete: Measuring the Impact of Discretization in End-to-End Differentially Private Synthetic Data. arXiv:2504.06923 [cs.CR] <https://arxiv.org/abs/2504.06923>
- [20] Chen Gong, Kecen Li, Zinan Lin, and Tianhao Wang. 2025. DPImageBench: A Unified Benchmark for Differentially Private Image Synthesis. arXiv preprint arXiv:2503.14681 (2025).
- [21] Chen Gong, Kecen Li, Zinan Lin, and Tianhao Wang. 2025. DPImageBench: A Unified Benchmark for Differentially Private Image Synthesis. arXiv:2503.14681 [cs.CR] <https://arxiv.org/abs/2503.14681>
- [22] Frederik Harder, Kamil Adamczewski, and Mijung Park. 2021. Dp-merf: Differentially private mean embeddings with randomfeatures for practical privacy-preserving data generation. In *International conference on artificial intelligence and statistics*. PMLR, 1819–1827.
- [23] Moritz Hardt, Katrina Ligett, and Frank McSherry. 2012. A simple and practical algorithm for differentially private data release. *Advances in neural information processing systems* 25 (2012).
- [24] Moritz Hardt and Guy N Rothblum. 2010. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st annual symposium on foundations of computer science*. IEEE, 61–70.
- [25] Yuzheng Hu, Fan Wu, Qinbin Li, Yunhui Long, Gonzalo Munilla Garrido, Chang Ge, Bolin Ding, David Forsyth, Bo Li, and Dawn Song. 2023. SoK: Privacy-Preserving Data Synthesis. arXiv:2307.02106 [cs.CR] <https://arxiv.org/abs/2307.02106>
- [26] IPUMS. 2025. IPUMS: Integrated Public Use Microdata Series. <https://www.ipums.org/>
- [27] James Jordan, Jinsung Yoon, and Mihaela Van Der Schaar. 2018. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *International conference on learning representations*.
- [28] Adam Kalai and Santosh Vempala. 2005. Efficient algorithms for online decision problems. *J. Comput. System Sci.* 71, 3 (2005), 291–307.
- [29] Rashid Hussain Khokhar, Benjamin CM Fung, Farkhund Iqbal, Khalil Al-Hussaini, and Mohammed Hussain. 2023. Differentially private release of heterogeneous network for managing healthcare data. *ACM Transactions on Knowledge Discovery from Data* 17, 6 (2023), 1–30.
- [30] Akim Kotelnikov, Dmitry Baranchuk, Ivan Rubachev, and Artem Babenko. 2023. Tabddpm: Modelling tabular data with diffusion models. In *International Conference on Machine Learning*. PMLR, 17564–17579.
- [31] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.
- [32] Haoran Li, Li Xiong, and Xiaoqian Jiang. 2014. Differentially private synthesis of multi-dimensional data using copula functions. In *Advances in database technology: proceedings. International conference on extending database technology*, Vol. 2014. NIH Public Access, 475.
- [33] Kecen Li, Chen Gong, Zhixiang Li, Yuzhong Zhao, Xinwen Hou, and Tianhao Wang. 2024. {PrivImage}: Differentially Private Synthetic Image Generation using Diffusion Models with {Semantic-Aware} Pretraining. In *33rd USENIX Security Symposium (USENIX Security 24)*. 4837–4854.
- [34] Ninghui Li, Zhiyun Zhang, and Tianhao Wang. 2021. DPSyn: Experiences in the NIST Differential Privacy Data Synthesis Challenges. arXiv:2106.12949 [cs.CR] <https://arxiv.org/abs/2106.12949>
- [35] Terrance Liu, Jingwu Tang, Giuseppe Vietri, and Steven Wu. 2023. Generating private synthetic data with genetic algorithms. In *International Conference on Machine Learning*. PMLR, 22009–22027.
- [36] Terrance Liu, Giuseppe Vietri, and Steven Z Wu. 2021. Iterative methods for private synthetic data: Unifying framework and new methods. *Advances in Neural Information Processing Systems* 34 (2021), 690–702.
- [37] Ryan McKenna, Jerome Miklau, Michael Hay, and Ashwin Machanavajjhala. 2018. Optimizing error of high-dimensional statistical queries under differential privacy. *Proceedings of the VLDB Endowment* 11, 10 (June 2018), 1206–1219. <https://doi.org/10.14778/3231751.3231769>
- [38] Ryan McKenna, Jerome Miklau, and Daniel Sheldon. 2021. Winning the NIST Contest: A scalable and general approach to differentially private synthetic data. arXiv preprint arXiv:2108.04978 (2021).
- [39] Ryan McKenna, Brett Mullins, Daniel Sheldon, and Jerome Miklau. 2022. Aim: An adaptive and iterative mechanism for differentially private synthetic data. arXiv preprint arXiv:2201.12677 (2022).
- [40] Ryan McKenna, Daniel Sheldon, and Jerome Miklau. 2019. Graphical-model based estimation and inference for differential privacy. In *International Conference on Machine Learning*. PMLR, 4435–4444.
- [41] Ilya Mironov. 2017. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*. IEEE, 263–275.
- [42] Ilya Mironov, Kunal Talwar, and Li Zhang. 2019. Rényi differential privacy of the sampled gaussian mechanism. arXiv preprint arXiv:1908.10530 (2019).
- [43] S. Moro, P. Rita, and P. Cortez. 2014. Bank Marketing. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5K306>.
- [44] Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. How to Use Heuristics for Differential Privacy. arXiv:1811.07765 [cs.LG] <https://arxiv.org/abs/1811.07765>
- [45] NIST. 2018. 2018 Differential Privacy Synthetic Data Challenge. Available at <https://www.nist.gov/ctl/pscr/open-innovation-prize-challenges/past-prize-challenges/2018-differential-privacy-synthetic>.
- [46] NIST. 2020. 2020 Differential Privacy Temporal Map Challenge. Available at <https://www.nist.gov/ctl/pscr/open-innovation-prize-challenges/past-prize-challenges/2020-differential-privacy-temporal>.
- [47] Wei Pang, Masoumeh Shafieinejad, Lucy Liu, and Xi He. 2024. ClavaDDPM: Multi-relational Data Synthesis with Cluster-guided Diffusion Models. arXiv preprint arXiv:2405.17724 (2024).
- [48] Alexandre Sablayrolles, Yue Wang, and Brian Karrer. 2023. Privately generating tabular data using language models. arXiv:2306.04803 [cs.LG] <https://arxiv.org/abs/2306.04803>
- [49] S Sangeetha, G Sudha Sadasivam, and Ayush Srikanth. 2022. Differentially private model release for healthcare applications. *International Journal of Computers and Applications* 44, 10 (2022), 953–958.
- [50] Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal* 27, 3 (1948), 379–423.
- [51] Mani Srivastava and Moustafa Alzantot. 2019. Differentially Private Dataset Release using Wasserstein GANs. https://github.com/nsl/nist_differential_privacy_synthetic_data_challenge. Accessed 2023-xx-xx.
- [52] Theresa Stadler, Bristena Oprisanu, and Carmela Troncoso. 2022. Synthetic Data – Anonymisation Groundhog Day. arXiv:2011.07018 [cs.LG] <https://arxiv.org/abs/2011.07018>

- [53] Felipe Petroski Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth O. Stanley, and Jeff Clune. 2018. Deep Neuroevolution: Genetic Algorithms Are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning. arXiv:1712.06567 [cs.NE] <https://arxiv.org/abs/1712.06567>
- [54] Arun Sai Suggala and Praneeth Netrapalli. 2020. Online non-convex learning: Following the perturbed leader is optimal. In *Algorithmic Learning Theory*. PMLR, 845–861.
- [55] Danyu Sun, Joann Qiongna Chen, Chen Gong, Tianhao Wang, and Zhou Li. 2024. Netdpsyn: synthesizing network traces under differential privacy. In *Proceedings of the 2024 ACM on Internet Measurement Conference*. 545–554.
- [56] Vasilis Syrgkanis, Akshay Krishnamurthy, and Robert Schapire. 2016. Efficient algorithms for adversarial contextual learning. In *International Conference on Machine Learning*. PMLR, 2159–2168.
- [57] Yuchao Tao, Ryan McKenna, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. 2022. Benchmarking Differentially Private Synthetic Data Generation Algorithms. arXiv:2112.09238 [cs.CR] <https://arxiv.org/abs/2112.09238>
- [58] MTCAJ Thomas and A Thomas Joy. 2006. *Elements of information theory*. Wiley-Interscience.
- [59] Reihaneh Torkzadehmahani, Peter Kairouz, and Benedict Paten. 2019. Dp-cgan: Differentially private synthetic data and label generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.
- [60] Toan V Tran and Li Xiong. 2024. Differentially Private Tabular Data Synthesis using Large Language Models. *arXiv preprint arXiv:2406.01457* (2024).
- [61] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [62] Giuseppe Vietri, Cedric Archambeau, Sergul Aydore, William Brown, Michael Kearns, Aaron Roth, Ankit Siva, Shuai Tang, and Steven Z Wu. 2022. Private synthetic data for multitask learning and marginal queries. *Advances in Neural Information Processing Systems* 35 (2022), 18282–18295.
- [63] Giuseppe Vietri, Grace Tian, Mark Bun, Thomas Steinke, and Zhiwei Steven Wu. 2020. New Oracle-Efficient Algorithms for Private Synthetic Data Release. arXiv:2007.05453 [cs.LG] <https://arxiv.org/abs/2007.05453>
- [64] Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. 2018. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739* (2018).
- [65] Mengmeng Yang, Chi-Hung Chi, Kwok-Yan Lam, Jie Feng, Taolin Guo, and Wei Ni. 2024. Tabular Data Synthesis with Differential Privacy: A Survey. *arXiv preprint arXiv:2411.03351* (2024).
- [66] Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao, Graham Cormode, and Ilya Mironov. 2021. Opacus: User-Friendly Differential Privacy Library in PyTorch. *arXiv preprint arXiv:2109.12298* (2021).
- [67] Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. 2017. Privbayes: Private data release via bayesian networks. *ACM Transactions on Database Systems (TODS)* 42, 4 (2017), 1–41.
- [68] Jun Zhang, Xiaokui Xiao, and Xing Xie. 2016. Privtree: A differentially private algorithm for hierarchical decompositions. In *Proceedings of the 2016 international conference on management of data*. 155–170.
- [69] Zhikun Zhang, Tianhao Wang, Ninghui Li, Jean Honorio, Michael Backes, Shibo He, Jiming Chen, and Yang Zhang. 2021. {PrivSyn}: Differentially private data synthesis. In *30th USENIX Security Symposium (USENIX Security 21)*. 929–946.

A DP PRELIMINARIES

A.1 RDP Properties

We briefly introduce RDP composition and post-processing properties in Theorem 3 and Theorem 4.

THEOREM 3 (COMPOSITION). *Let $f : \mathcal{D} \rightarrow \mathcal{R}_1$ be (α, ε_1) -RDP and $g : \mathcal{R}_1 \times \mathcal{D} \rightarrow \mathcal{R}_2$ be (α, ε_2) -RDP respectively. Then the mechanism defined as (X, Y) , where $X \sim f(D)$ and $Y \sim (D, f(D))$, satisfies $(\alpha, \varepsilon_1 + \varepsilon_2)$ -RDP.*

THEOREM 4 (POST-PROCESSING). *Let $f : \mathcal{D} \rightarrow \mathcal{R}_1$ is (α, ε) -RDP, and $g : \mathcal{R}_1 \rightarrow \mathcal{R}_2$ is an arbitrary randomized mapping. Then $g \circ f$ is also (α, ε) -RDP.*

A.2 DP Mechanism

Before introducing DP mechanisms, a key quantity needed to be defined is the sensitivity:

DEFINITION 3 (SENSITIVITY). *Let $f : \mathcal{D} \rightarrow \mathcal{R}^k$ be a vector-valued function of the input data, then the ℓ_2 sensitivity of f is defined as*

$$\Delta_f = \max_{D \sim D'} \|f(D) - f(D')\|_2$$

Gaussian Mechanism. Gaussian Mechanism (GM) [41], which adds noise sampled from Gaussian distribution, has widely been used to achieve (α, ε) -RDP. Specifically, Let f be a vector-valued function of the input data. The Gaussian Mechanism adds i.i.d. Gaussian noise with scale $\sigma \Delta_f$ to each entry of f .

$$\mathcal{A}(D) = f(D) + \sigma \Delta_f \mathcal{N}(0, \mathbf{I}), \quad (3)$$

where \mathcal{N} refers to Gaussian distribution. The RDP guarantee of GM is given by the Theorem 5.

THEOREM 5. *The Gaussian Mechanism defined above satisfies $(\alpha, \frac{\alpha}{2\sigma^2})$ -RDP.*

Exponential Mechanism. Let q_r be a score function for all $r \in \mathcal{R}$. Then the exponential mechanism (EM) [36, 39] outputs a candidate r according to the following distribution:

$$\Pr[\mathcal{A}(D) = r] \propto \exp\left(\frac{\varepsilon}{2\Delta} \cdot q_r(D)\right), \quad (4)$$

where $\Delta = \max_{r \in \mathcal{R}} \Delta(q_r)$. The RDP guarantee of EM is provided by Theorem 6

THEOREM 6. *The Exponential Mechanism defined above satisfies $(\alpha, \frac{\alpha \varepsilon^2}{8})$ -RDP for $\forall \alpha > 1$.*

Gumbel Noise. Reporting noise max with Gumbel noise is a derivative mechanism from the exponential mechanism used for marginal query selection by some previous work [62, 63]. Specifically, it outputs $i^* = \arg \max_{i \in [m]} \{|q_i(D) - a_i| + Z_i\}$, where q_i is the marginal query function defined in [62, 63]; a_i is the query answer and $Z_i \sim \text{Gumbel}(1/n\sqrt{2\rho})$. This mechanism has been proven to satisfy $(\alpha, \alpha\rho)$ -RDP for $\forall \alpha > 1$.

DP-SGD. Differential private stochastic gradient descent (DP-SGD) is the most popular way to train the model to satisfy DP. In the training process, we assume that \mathcal{L} is the loss function, and we

have a clipping function defined by $\text{Clip}_C(g) = \min\left\{1, \frac{C}{\|g\|_2}\right\} g$ and a Gaussian noise level σ [42]. The DP-SGD can be expressed as

$$\theta \leftarrow \theta - \eta \left(\frac{1}{|b|} \sum_{i \in b} \text{Clip}_C(\nabla \mathcal{L}(\theta, x_i)) + \text{CN}(0, \sigma^2 \mathbf{I}) \right)$$

Here η is the learning rate, $\nabla \mathcal{L}(\theta, x_i)$ is the gradient of the loss function \mathcal{L} in relation to model parameters θ and data point x_i in sample batch b . By clipping the gradient, we control the sensitivity and thus can apply the Gaussian mechanism to guarantee DP.

B DP GUARANTEE OF PREPROCESSING ALGORITHMS

B.1 PrivTree Binning

The privacy guarantee is given by the following lemma.

LEMMA 1. *For any $\alpha > 1$, Algorithm 2 satisfy $(\alpha, \alpha\rho_1)$ -Rényi differential privacy.*

This proof is organized by two steps: firstly we proof that for each attribute, Algorithm 2 can achieve a $(\alpha, \alpha\rho_1/K)$ -Rényi DP, then by composition theorem, we can draw the conclusion that Algorithm 2 can achieve $(\alpha, \alpha\rho_1)$ -Rényi DP totally.

By the privacy proof of PrivTree in Zhang et al.'s work [68], we have that if

$$\lambda' \geq \frac{2\beta - 1}{\beta - 1} \cdot \frac{1}{\varepsilon}$$

and

$$\delta' = \lambda' \cdot \ln \beta$$

single PrivTree algorithm satisfies ε -DP. Replacing ε and β by the definition in Algorithm 2, we have that each round of PrivTree binning satisfies $\frac{2\rho_1}{|V_n|}$ -DP

Then referring to Bun et al.'s work [6], we have that if a mechanism satisfies ε -DP, it also satisfies $(\alpha, \frac{1}{2}\varepsilon^2\alpha)$ -RDP. Therefore we obtain that PrivTree is $(\alpha, \frac{\rho_1}{|V_n|}\alpha)$ -RDP. Because we need to conduct PrivTree binning for $|V_n|$ rounds, by composition theorem of RDP, we have that Algorithm 2 satisfies $(\alpha, \rho_1\alpha)$ -RDP.

B.2 Rare Category Merging

For each attribute that needs preprocessing, we equally divide the privacy budget and use it to determine those categories whose frequency is lower than the threshold. These categories will be replaced by the same encoding category. The privacy guarantee can be formalized in the following lemma.

LEMMA 2. *For any $\alpha > 1$, Algorithm 1 satisfies $(\alpha, \alpha\rho_2)$ -Rényi differential privacy.*

The proof of this lemma can be easily obtained by the property of the Gaussian mechanism. Referring to Theorem 5, we know that adding Gaussian noise with $\sigma = \sqrt{\frac{1}{2\rho'}}$ satisfying $(\alpha, \alpha\rho')$ -Rényi differential privacy for any $\alpha > 1$. Then by combining the fact that $\rho' = \rho_2/|V_c|$ and the composition property of RDP, we have that the total algorithm satisfies $(\alpha, \alpha\rho_2)$ -Rényi differential privacy for any $\alpha > 1$.

C MISSING PROOF

In this section, we provide the proof of Theorem 2. Before we prove this theorem, we give a lemma as follows.

LEMMA 3. *Assuming that*

$$P_1(z) = \sum_x p(x)P_1(z|x) \quad \text{and} \quad P_2(z) = \sum_x p(x)P_2(z|x),$$

we have

$$\mathbb{D}_{KL}(P_1(z) \parallel P_2(z)) \leq \sum_x p(x)\mathbb{D}_{KL}(P_1(z|x) \parallel P_2(z|x))$$

Firstly, we have

$$\ln\left(\frac{P_1(z)}{P_2(z)}\right) = \ln\left(\frac{\sum_x p(x)P_1(z|x)}{\sum_x p(x)P_2(z|x)}\right).$$

We already have “log-sum” inequality [10], which can be expressed as

$$\sum_{i=1}^n x_i \log\left(\frac{x_i}{y_i}\right) \geq \left(\sum_{i=1}^n x_i\right) \log\left(\frac{\sum_{i=1}^n x_i}{\sum_{i=1}^n y_i}\right).$$

Let $\alpha_x(z) = \frac{p(x)P_1(z|x)}{P_1(z)}$, by “log-sum” inequality, we have

$$\begin{aligned} \ln\left(\frac{P_1(z)}{P_2(z)}\right) &\leq \sum_x \alpha_x(z) \ln\left(\frac{p(x)P_1(z|x)}{p(x)P_2(z|x)}\right) \\ &= \sum_x \alpha_x(z) \ln\left(\frac{P_1(z|x)}{P_2(z|x)}\right) \end{aligned}$$

Taking mathematical expectations, we have

$$\mathbb{E}_{z \sim P_1} \left[\ln\left(\frac{P_1(z)}{P_2(z)}\right) \right] \leq \mathbb{E}_{z \sim P_1} \left[\sum_x \alpha_x(z) \ln\left(\frac{P_1(z|x)}{P_2(z|x)}\right) \right]$$

Notice that $\alpha_x(z) = P(x|z)$, we can rewrite the right-hand side of the above inequality as

$$\begin{aligned} &\sum_x \mathbb{E}_{x, z \sim p(x), P_1(z|x)} \left[\mathbb{I}(X=x) \ln\left(\frac{P_1(z|x)}{P_2(z|x)}\right) \right] \\ &= \sum_x p(x)\mathbb{D}_{KL}(P_1(z|x) \parallel P_2(z|x)) \end{aligned}$$

Combining all above, we have proved that

$$\mathbb{D}_{KL}(P_1(z) \parallel P_2(z)) \leq \sum_x p(x)\mathbb{D}_{KL}(P_1(z|x) \parallel P_2(z|x)).$$

Now we give the proof of Theorem 2. For the left-hand side of Equation (2), we have

$$\mathbb{D}_{KL}(\Pr[A_i, A_j] \parallel \Pr[A_i] \Pr[A_j]) = I(A_i, A_j), \quad (5)$$

where I refers to mutual information [50]. For the right-hand side of Equation (2), applying Lemma 3, we have the following property:

$$\begin{aligned} &\mathbb{D}_{KL}(\Pr[A_i, A_j] \parallel \\ &\quad \sum_{A_1, \dots, A_k} \Pr[A_1, \dots, A_k] \cdot \Pr[A_i|A_1, \dots, A_k] \Pr[A_j|A_1, \dots, A_k]) \\ &= \mathbb{D}_{KL} \left(\sum_{A_1, \dots, A_k} \Pr[A_1, \dots, A_k] \Pr[A_i, A_j|A_1, \dots, A_k] \parallel \right. \\ &\quad \left. \sum_{A_1, \dots, A_k} \Pr[A_1, \dots, A_k] \cdot \Pr[A_i|A_1, \dots, A_k] \Pr[A_j|A_1, \dots, A_k] \right) \\ &\leq \sum_{A_1, \dots, A_k} \Pr[A_1, \dots, A_k] \cdot \mathbb{D}_{KL}(\Pr[A_i, A_j|A_1, \dots, A_k] \parallel \\ &\quad \Pr[A_i|A_1, \dots, A_k] \Pr[A_j|A_1, \dots, A_k]) \\ &= I(A_i, A_j \mid A_1, \dots, A_k) \end{aligned} \quad (6)$$

where $I(\cdot \mid \cdot)$ is the conditional mutual information. Then by property of mutual information [58], we have

$$I(A_i, A_j \mid A_1, \dots, A_k) \leq I(A_i, A_j). \quad (7)$$

Combining Equation (5), Equation (6) and Equation (7), we can prove Theorem 2.

D MISSING EXPERIMENTAL SETTINGS

D.1 Dataset information

ACSIncome and ACSemploy [13] are both drawn from 2018 national census data. The Bank dataset is on the UCI open dataset website [43], which is related to direct marketing campaigns of a Portuguese banking institution. Higgs-small [1] dataset was produced using Monte Carlo simulations, which include different features of particles in the accelerator. Loan dataset [2] was derived from data LendingClub issued through 2007-2014.

We make a summary of these datasets’ basic statistics in Table 5. In this table, we represent the number of records, number of attributes/numerical attributes/categorical attributes, and attributes’ domain size range of these datasets. All of these datasets have been processed to ensure that there are no missing values.

Table 5: Summary of investigated datasets.

Name	#Records	#Attr	#Num	#Cat	Min/Max Domain
ACSIncome (INC) [13]	55320	10	2	8	2~93
ACSemploy (EMP) [13]	37881	17	1	16	2~92
Bank (BK) [43]	45211	16	6	10	2~6024
Higgs-small (HIG) [1]	98049	28	28	1	2~73715
Loan (LN) [2]	134658	42	25	17	2~93995

D.2 Hyperparameters for Preprocessing algorithms

We apply uniform discretization and rare category merging for all algorithms as the default aligned preprocessing methods. Moreover,

as shown in Table 5, some attributes only contain a few unique values, which are simple enough to handle without additional preprocessing. Applying preprocessing to such attributes is unnecessary and may introduce more errors. Therefore, we preprocess attributes only when their domain size exceeds 100. By default, the number of bins is set to 100, the fixed merging threshold θ to 0.2%, and the privacy budget proportion allocated to the preprocessing step to 10%.

D.3 Hyperparameters for Full Algorithms

We list the algorithms’ hyperparameters in Section 9.1. From here on out, unless otherwise specified, INC refers to the ACSincome dataset; EMP refers to the ACSemploy dataset; BK refers to the Bank dataset; HIG refers to the Higgs-small dataset; LN refers to the Loan dataset.

Table 6: PrivSyn Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Consistent iteration	501	501	501	501	501
Max update iteration	50	50	50	50	50

Table 7: AIM Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Max model size	100	100	100	100	100
Max iteration	1000	1000	1000	1000	1000
Max marginal size	$2.5e + 5$				

Table 8: Private-GSD Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Mutation number	50	50	50	50	50
Crossover number	50	50	50	50	50
Upsample number	$1e + 5$				
Genetic iteration	$1e + 6$				

Table 9: PrivMRF Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Graph construction parameter θ	6	6	6	6	6
Sample size k	400	400	400	400	400
Estimation iteration	3000	3000	3000	3000	3000
Size penalty	$1e - 8$				
Max marginal attributes number	6	6	6	6	6
Max clique size	$1e + 7$				

Table 10: GEM Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Synthesis size	1024	1024	1024	1024	1024
Learning rate	$1e - 3$				
Max iteration	500	500	500	500	500
Max selection round	50	85	80	140	210

Table 11: RAP++ Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Random projection number	$2e + 6$				
Categorical optimization rate	$3e - 3$				
Numerical optimization rate	$6e - 3$				
Top q	5	5	5	5	5
Categorical optimization step	1	1	1	1	1
Numerical optimization step	3	3	3	3	3
Upsample rate	10	10	20	20	40

Table 12: DP-MERF Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Random feature dimension	$2e + 3$				
Mini batch rate	$5e - 2$				
Epoch number	$1e + 3$				
Learning rate	$1e - 2$				

Table 13: TabDDPM Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Denoisier layer dimension	256	256	256	256	1024
Denoisier layer number	2	2	2	2	2
Epoch number	50	100	100	100	100
Batch size	512	512	512	1024	1024
Learning rate	$2e - 2$	$1e - 2$	$5e - 3$	$5e - 2$	$5e - 4$
Diffusion steps	100	100	100	1000	100

D.4 Hyperparameters for Different Feature Selection Algorithms

The PrivMRF and AIM selection algorithms are set to completely the same as the original work in PrivMRF and AIM, respectively. Therefore, we omit the description of them here and provide the detailed hyperparameter setting of RAP++ selection and PrivSyn selection.

Table 15: RAP++ Selection Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Top q	3	3	3	3	3
Selection step	4	6	6	7	10
Selection budget rate	0.5	0.5	0.5	0.5	0.5
Marginal budget rate	0.5	0.5	0.5	0.5	0.5

Table 16: PrivSyn Selection Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Selection budget rate	0.1	0.1	0.1	0.1	0.1
1-way marginal budget rate	0.1	0.1	0.1	0.1	0.1
2-way marginal budget rate	0.8	0.8	0.8	0.8	0.8

Table 14: Supplementary overall results of synthetic data under different methods. ML AUC and ML Accuracy are metrics obtained by downstream ML tasks. Running time is the total average execution time of the algorithm. Because Loan dataset is a multi-classification problem, thus it does not have AUC result.

Dataset	ACSincome			ACSEmploy			Bank			Higgs-small			Loan		
ML AUC	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn	0.53	0.50	0.51	0.47	0.47	0.42	0.44	0.52	0.49	0.50	0.50	0.50	-	-	-
PrivMRF	0.82	0.87	0.87	0.80	0.86	0.89	0.71	0.90	0.92	0.53	0.71	0.70	-	-	-
RAP++	0.75	0.82	0.85	0.81	0.84	0.87	0.75	0.85	0.88	0.54	0.56	0.56	-	-	-
AIM	0.85	0.87	0.87	0.85	0.88	0.88	0.87	0.89	0.91	0.68	0.72	0.74	-	-	-
Private-GSD	0.85	0.85	0.85	0.78	0.80	0.79	0.68	0.67	0.67	0.52	0.52	0.52	-	-	-
GEM	0.77	0.73	0.72	0.75	0.77	0.77	0.59	0.68	0.69	0.55	0.57	0.59	-	-	-
DP-MERF	0.75	0.78	0.79	0.73	0.78	0.74	0.72	0.64	0.65	0.54	0.60	0.60	-	-	-
TabDDPM	0.54	0.49	0.53	0.56	0.56	0.55	0.48	0.51	0.45	0.51	0.53	0.53	-	-	-
ML Accuracy	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn	0.59	0.58	0.59	0.52	0.49	0.49	0.88	0.88	0.88	0.53	0.52	0.52	0.54	0.54	0.54
PrivMRF	0.75	0.79	0.79	0.72	0.79	0.81	0.89	0.90	0.90	0.53	0.65	0.64	0.76	0.75	0.76
RAP++	0.68	0.75	0.78	0.74	0.78	0.80	0.85	0.88	0.89	0.53	0.55	0.55	0.62	0.64	0.65
AIM	0.78	0.79	0.79	0.78	0.80	0.81	0.90	0.90	0.90	0.63	0.65	0.67	0.75	0.75	0.75
Private-GSD	0.77	0.77	0.78	0.71	0.73	0.71	0.87	0.87	0.88	0.51	0.51	0.52	0.54	0.54	0.54
GEM	0.71	0.69	0.66	0.68	0.70	0.69	0.86	0.86	0.87	0.55	0.55	0.56	0.71	0.68	0.72
DP-MERF	0.69	0.70	0.72	0.64	0.71	0.67	0.84	0.79	0.74	0.53	0.58	0.58	0.35	0.37	0.37
TabDDPM	0.57	0.59	0.59	0.55	0.54	0.53	0.88	0.88	0.88	0.50	0.53	0.51	0.55	0.55	0.55
Running Time (min)	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn	0.28	0.16	0.18	0.30	0.25	0.34	0.57	0.53	0.55	5.39	5.95	4.54	13.29	14.04	13.44
PrivMRF	1.26	0.89	1.40	5.59	5.80	4.76	3.06	4.48	3.24	5.06	7.10	6.25	7.12	13.34	12.62
RAP++	25.86	24.95	25.72	26.28	27.04	28.30	24.40	23.94	23.27	37.79	36.66	35.88	544.25	2348.59	1761.60
AIM	1.96	6.29	126.89	5.37	10.59	443.70	3.60	10.23	23.57	12.62	18.46	184.57	31.03	187.27	645.13
Private-GSD	22.33	24.86	26.17	10.89	11.12	11.11	19.38	20.01	19.99	54.35	57.49	57.93	304.42	306.04	311.63
GEM	0.26	0.10	0.09	0.12	0.12	0.12	0.27	0.25	0.25	6.05	6.09	10.75	9.82	9.47	9.82
DP-MERF	0.40	0.07	0.06	0.30	0.06	0.06	0.09	0.06	0.06	0.09	0.06	0.06	0.39	0.83	0.34
TabDDPM	4.34	4.12	3.73	4.52	4.50	13.31	8.25	4.40	3.87	4.74	6.34	4.78	27.54	31.93	33.34

D.5 Hyperparameters for Different Synthesis Algorithms

Most synthesis algorithms we used are set to be the same as their original works, while relaxed projection and generative network methods need hyperparameter tuning to guarantee performance.

Table 17: Relaxed Projection Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Random projection number	$2e + 6$				
Optimization rate	$5e - 3$				
Optimization step	100	170	160	280	420

Table 18: Generative Network Hyperparameters

Hyperparameter	INC	EMP	BK	HIG	LN
Learning rate	$1e - 3$				
Synthesis size	1024	1024	1024	1024	1024
Max training iteration	50	50	100	100	1500

E SUPPLEMENTARY EXPERIMENT RESULTS

E.1 More Results of Algorithm Comparison

We have provided some important metric results of algorithm utility in Section 9.1. Here we list some other detailed evaluations such as

AUC, accuracy for machine learning efficiency and running time for time efficiency. Notably, these results do not necessarily influence our conclusion, thus we present them in the appendix.

E.2 Detailed Results of Preprocessing Influence

The detailed results of different preprocessing methods (used to plot Figure 5) are shown in Table 19. Similar to other experiments, these metrics are calculated by comparing preprocessed datasets with test datasets to demonstrate the error caused by preprocessing.

E.3 Detailed Results of Reconstruction Experiment

The detailed results of different preprocessing methods (used to plot Figure 7 and Figure 8) are shown in Table 20 and Table 21. Here, there are some “-” in the table. This is because PrivSyn tends to select as many marginals as possible, which will form large cliques. This will cause the size of the graphical model to be too large, requiring extremely large memory.

Table 19: Influences of preprocessing on different datasets. By default, the results are obtained under the setting that $\epsilon = 1.0$ and 10% of the budget is used for preprocessing.

Preprocessing Method	Marginal Size			ML Efficiency			Query Error			Fidelity Error		
	Bank	Higgs-small	Loan	Bank	Higgs-small	Loan	Bank	Higgs-small	Loan	Bank	Higgs-small	Loan
Raw	$1.05e+5$	$4.24e+8$	$2.88e+8$	0.76	0.72	0.54	0.001	0.001	0.001	0.003	0.001	0.001
Numerical preprocessing	$6.48e+2$	$5.05e+3$	$1.14e+6$	0.76	0.71	0.53	0.002	0.002	0.003	0.006	0.002	0.003
Categorical Preprocessing	$1.05e+5$	$4.24e+8$	$2.31e+8$	0.76	0.72	0.54	0.001	0.001	0.001	0.003	0.001	0.001
Full Preprocessing	$6.48e+2$	$5.05e+3$	$3.14e+3$	0.76	0.71	0.53	0.002	0.002	0.003	0.006	0.002	0.003

Table 20: Results of synthetic data under different feature selection methods. By default, we use PGM as the synthesis method. In this table, “-” means unable to execute due to time or memory limitation.

Dataset	ACSIncome			ACSEmploy			Bank			Higgs-small			Loan		
ML efficiency	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn selection	0.67	0.65	0.68	0.67	0.80	0.75	-	-	-	-	-	-	-	-	-
PrivMRF selection	0.73	0.78	0.78	0.81	0.80	0.81	0.62	0.70	0.71	0.50	0.64	0.64	0.52	0.52	0.52
RAP++ selection	0.62	0.74	0.74	0.74	0.78	0.79	0.51	0.51	0.47	0.49	0.49	0.50	0.30	0.26	0.26
AIM selection	0.76	0.78	0.78	0.78	0.80	0.81	0.67	0.71	0.70	0.63	0.65	0.67	0.52	0.52	0.52
Query Error	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn selection	0.003	0.001	0.001	0.003	0.002	0.003	-	-	-	-	-	-	-	-	-
PrivMRF selection	0.002	0.001	0.001	0.003	0.002	0.002	0.005	0.003	0.003	0.005	0.003	0.003	0.005	0.005	0.004
RAP++ selection	0.006	0.003	0.002	0.008	0.005	0.004	0.012	0.005	0.003	0.047	0.016	0.006	0.015	0.009	0.008
AIM selection	0.002	0.001	0.001	0.003	0.002	0.001	0.007	0.002	0.002	0.005	0.003	0.003	0.005	0.005	0.004
Fidelity Error	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
PrivSyn selection	0.14	0.08	0.07	0.07	0.04	0.04	-	-	-	-	-	-	-	-	-
PrivMRF selection	0.11	0.07	0.05	0.07	0.04	0.03	0.13	0.06	0.04	0.36	0.19	0.19	0.31	0.24	0.23
RAP++ selection	0.17	0.09	0.08	0.04	0.04	0.04	0.14	0.09	0.08	0.42	0.23	0.17	0.32	0.26	0.25
AIM selection	0.09	0.06	0.05	0.05	0.03	0.02	0.11	0.09	0.06	0.21	0.16	0.16	0.35	0.32	0.29

Table 21: Results of synthetic data under different synthesis methods. By default, we use PrivSyn’s InDif selection as the selection method. In this table, “-” means unable to execute due to time or memory limitation.

Dataset	ACSIncome			ACSEmploy			Bank			Higgs-small			Loan		
ML efficiency	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
GUM	0.39	0.40	0.42	0.45	0.45	0.40	0.47	0.47	0.47	0.40	0.43	0.43	0.25	0.26	0.26
PGM	0.76	0.64	0.67	0.68	0.80	0.75	-	-	-	-	-	-	-	-	-
Relaxed Projection	0.39	0.59	0.58	0.51	0.69	0.67	0.47	0.47	0.47	0.47	0.44	0.42	0.25	0.25	0.25
Genetic Algorithm	0.67	0.63	0.58	0.62	0.58	0.62	0.61	0.57	0.52	0.59	0.58	0.62	0.35	0.36	0.26
Generative Network	0.74	0.77	0.76	0.72	0.70	0.69	0.66	0.63	0.62	0.63	0.53	0.59	0.46	0.41	0.36
Query Error	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
GUM	0.003	0.002	0.002	0.006	0.004	0.004	0.007	0.004	0.003	0.009	0.004	0.003	0.006	0.005	0.004
PGM	0.003	0.001	0.001	0.003	0.002	0.003	-	-	-	-	-	-	-	-	-
Relaxed Projection	0.040	0.028	0.023	0.081	0.025	0.022	0.047	0.013	0.014	0.026	0.013	0.012	0.020	0.009	0.008
Genetic Algorithm	0.053	0.047	0.050	0.039	0.040	0.036	0.009	0.005	0.003	0.027	0.018	0.016	0.039	0.037	0.039
Generative Network	0.002	0.002	0.002	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.004	0.004
Fidelity Error	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$	$\epsilon = 0.2$	$\epsilon = 1$	$\epsilon = 5$
GUM	0.15	0.12	0.12	0.12	0.09	0.09	0.24	0.10	0.21	0.29	0.20	0.20	0.34	0.35	0.36
PGM	0.13	0.08	0.07	0.07	0.04	0.04	-	-	-	-	-	-	-	-	-
Relaxed Projection	0.61	0.52	0.41	0.51	0.26	0.21	0.41	0.16	0.16	0.31	0.22	0.21	0.41	0.29	0.28
Genetic Algorithm	0.59	0.59	0.56	0.30	0.30	0.30	0.13	0.09	0.08	0.35	0.28	0.27	0.56	0.55	0.55
Generative Network	0.08	0.08	0.08	0.06	0.07	0.06	0.16	0.16	0.15	0.24	0.23	0.23	0.47	0.38	0.27