

RESEARCH ARTICLE

Trading Off Scalability, Privacy, and Performance in Data Synthesis

XIAO LING¹, TIM MENZIES¹, (Fellow, IEEE), CHRISTOPHER HAZARD²,
JACK SHU², AND JACOB BEEL²

¹Computer Science Department, North Carolina State University, Raleigh, NC 27695, USA

²Howso, Raleigh, NC 27603, USA

Corresponding author: Xiao Ling (xling4@ncsu.edu)

ABSTRACT Synthetic data has been widely applied in the real world recently. One typical example is the creation of synthetic data for privacy concerned datasets. In this scenario, synthetic data substitute the real data which contains the privacy information, and is used to public testing for machine learning models. Another typical example is the unbalance data over-sampling which the synthetic data is generated in the region of minority samples to balance the positive and negative ratio when training the machine learning models. In this study, we concentrate on the first example, and introduce (a) the Howso engine, and (b) our proposed random projection based synthetic data generation framework. We evaluate these two algorithms on the aspects of privacy preservation and accuracy, and compare them to the two state-of-the-art synthetic data generation algorithms DataSynthesizer and Synthetic Data Vault. We show that the synthetic data generated by Howso engine has good privacy and accuracy, which results in the best overall score. On the other hand, our proposed random projection based framework can generate synthetic data with highest accuracy score, and has the fastest scalability.

INDEX TERMS Synthetic data generation, privacy preservation, regression, classification.

I. INTRODUCTION

To successfully apply artificial intelligent approaches (e.g. machine learning & deep learning algorithms) in the real world application, data has become the most important part to support these algorithms. However, many types of data have privacy concerns and limited in terms of publicity [1], [2]. This causes the issue on training the machine learning models on these regions since these models require support from huge training data.

To mitigate this issue, synthetic data has been widely studied to substitute the real data. Synthetic data is the fake data points which generated by the generative model with the information from the real data points [3], [4], [5]. To evaluate the real world application of synthetic data, privacy preservation is one of the critical measurements which checks if identities from the original dataset can be detected or recognized in the synthetic data [6], [7], [8].

The associate editor coordinating the review of this manuscript and approving it for publication was Porfirio Tramontana¹.

Similarity is another measurement since the synthetic data need to capture the information from the original data [9], [10]. Moreover, another important measurement is to check if the synthetic data can be used to substitute the real data on training the models [11], [12].

In this study, we evaluate a new data synthesis method called recursive random projections. Initially developed for optimization, recursive random projection uses the FASTMAP [13], [14] technique to recursively bi-cluster the data into numerous small leaf clusters. An optimizer then samples just $N = 2$ points per leaf. But after reading the data synthesis literature [3], [4], [5], [6], [7], [8], [9], [10], [11], [12] we began to wonder if recursive random projection could also be a data synthesis algorithm just by sampling much more than $N = 2$.

To check this, we perform the study described in this paper. As described in §IV, random projection is augmented with mutation and crossover operators to generate synthetic data points per leaf cluster. This is then compared to state-of-the-art synthetic data generation algorithms such as (a) the

DataSynthesizer [15], (b) the Synthetic Data Vault [16], and (c) the Howso engine [17].

To structure this inquiry, we ask these questions:

- RQ1: When considering the privacy, which synthetic data generation algorithm can generate the synthetic data with the highest privacy preservation score?
- RQ2: Which synthetic data generation algorithm can generate data that has higher similarity to the original data?
- RQ3: When the machine learning model is trained on the synthetic data, can the model achieve compatible performance with those trained on the original data?
- RQ4: Which algorithm has the best scalability?
- RQ5: What suggestions can we provide from analyzing the conclusions in RQ1 to RQ4?

The contributions of this paper are

- We proposed a random projection based synthetic data generation framework.
- We made an empirical experiment to compare our proposed method, Howso engine, and two state-of-the-art synthetic data generation algorithms in different aspects such as (a) privacy preservation, (b) statistical measurements, (c) marginal probability, and (d) performance on training the machine learning models.
- We found random projection based framework and Howso engine outperform state-of-the-art methods in some of the aspects.
- The random projection based framework, in terms of scalability, can run significantly faster than Howso engine, and have similar runtime comparing to state-of-the-art methods, while outperform on more metrics than state-of-the-art methods.

The rest of this paper is constructed as follow: Section II illustrates the background of synthetic data generation, privacy & accuracy, and literature review on the synthetic data generation. Section III presents the Howso engine. Section IV illustrates our proposed random projection based synthetic data generation framework. Section V shows two state-of-the-art synthetic data generation algorithms DataSynthesizer and Synthetic Data Vault. Section VI presents the summary of benchmarks, evaluation metrics, and statistical analysis used in our experiment. Section VII shows our experimental results and our analysis to the results. We also discuss the threat to validity of our experiment in Section VIII, and make a conclusion to this study in Section IX.

II. BACKGROUND

A. SYNTHETIC DATA GENERATION

Synthetic data is used to substitute the real data which cannot be shared to public due to privacy concerns. It is usually generated by the generative model which learns the patterns of data from the original dataset. More specifically, the data synthesize process usually can be described as follow:

Given a dataset \mathcal{D} which some of the features $\{f_i, f_j, \dots, f_k\}$ have privacy concerns, the generative model \mathcal{M} learns the statistical properties p for each feature

and the correlation c between different features from the original dataset \mathcal{D} . Then the generative model $\mathcal{M}(p, c)$ will generate the synthetic dataset \mathcal{D}' that similar to \mathcal{D} , but no original identity can be detected by the synthetic value in the features $\{f_i, f_j, \dots, f_k\}$.

Researchers mainly focus on the (a) privacy preservation [6], [33], [34] and (b) model accuracy [35], [36] to judge if the synthesized data is both informative and safe to share. As we will discuss later in this paper, we evaluate the synthetic data in our study through both privacy preservation score and model performance score, as well as the statistical distribution score.

B. RELATED WORK

Synthetic data generation and evaluation has been widely studied in the past few years. The literature mainly split in two directions. One direction is the development of new synthetic data generation algorithms such as [15] and [16]. Another direction is applying well-established generation methods to different datasets in various domains, and evaluating the results through different metrics [30].

We search the literature in Google Scholar for what has been published in top venues¹ then summarize the related work in Table 1. To our best knowledge, recent literature on synthetic data can be mainly split into two regions. The studies in the first category develops the new generation algorithm, and the studies in the second category mainly evaluates the utility of synthetic data generated by different models through different metrics. In our study, we introduce Howso engine and our proposed random projection based synthetic data generation framework, so we compare these two methods to the algorithms proposed in the literature in the first category. We choose our comparison objects with following rules:

- First, the study goal should focus on the tabular data since all benchmarks in our study are tabular based.
- Second, the implementation should be based on Python to compatible with our inputs.

This means we focused on (a) the DataSynthesizer [15] and (b) the Synthetic Data Vault [16]. Hence, in our study, we compare Howso engine and random projection based framework to DataSynthesizer and Synthetic Data Vault.

III. HOWSO ENGINE

Howso Engine is developed by Howso.² It is an AI engine which supports the synthetic data generation. Specifically, Howso Engine utilizes the k-nearest neighbors to synthesize data with both global and local distributions [37]. The algorithm contains three parameters to control the search, and the best combination of the parameters is found by the grid search. Since there is only two parameters that need to be explored, the grid search is very fast.

¹As defined by Google Scholar metrics.

²<https://www.howso.com/>

TABLE 1. Literature review on recent synthetic data studies. The literature can be described in two categories: (a) The design of new synthetic data generation algorithms (synthetic data generation), and (b) The implementation of synthetic data generation algorithms and the evaluation on different utility metrics (utility validation). Some of long titles are not fully shown in the first column. For the entire title please refer to the specific reference.

Paper	Year	Algorithm(s)	Technique	Metric(s)	Region
Synthetic Data Generation					
The Synthetic Data Vault [16]	2016	SDV	Gaussian Copula	Accuracy & Qualitative findings	N/A
synthpop: Bespoke creation of synthetic data in R [18]	2016	SynthPop	CART tree based method	statistical metrics	N/A
DataSynthesizer: Privacy-Preserving Synthetic Datasets [15]	2017	DataSynthesizer	Greedy Bayes	Similarity measure (e.g feature distribution)	Urban science
LoPub: High-Dimensional Crowdsourced Data Publication... [19]	2018	LoPub	LDP & Expectation maximization	Accurate rate	N/A
Locally Private High-Dimensional Crowdsourced Data... [20]	2019	LoCop & DR_LoCop	Univariate marginal distribution & attribute dependence	Accurate rate	N/A
RPA and L-System Based Synthetic Data Generator for Cost-efficient Deep Learning Model Training [11]	2021	Lindenmayer Systems	RPA & L-system	Machine learning metrics	Image
Synthetic data generation using building information models [12]	2021	CycleGAN	GAN	Average precision	Image
Fedsyn: Synthetic data generation using federated learning [21]	2022	FedSyn	Federated learning	Accuracy & subjective analysis	Image
LDP-IDS: Local Differential Privacy for Infinite Data Stream [22]	2022	LDP-IDS	LDP	Privacy metric & MSE based utility metric	N/A
Generation of synthetic tympanic membrane images... [23]	2023	GAN	GAN	Human review	Medical
Utility Validation					
Utility of synthetic microdata generated using tree-based methods [10]	2015	SynthPop	N/A	Propensity score, Coefficient estimates, Mean overlap in the 95% confidence intervals	N/A
The validity of synthetic clinical data: a validation study of... [24]	2019	Synthea	N/A	Medical quality measurement	Medical
On the Utility of Synthetic Data: An Empirical Evaluation on Machine Learning Tasks [25]	2019	DS, SDV	N/A	Distribution, Correlation coefficient, Distance of nearest neighbors, Accuracy	N/A
Empirical evaluation on synthetic data generation with generative adversarial network [26]	2019	GAN	N/A	Correlation matrix, Accuracy, Privacy metrics	N/A
Generation and evaluation of synthetic patient data [27]	2020	Sampling from marginal, Bayesian network, GAN, Gaussian process	N/A	LK divergence, pairwise correlation difference, log-cluster, cross-classification	Medical
Can synthetic data be a proxy for real clinical trial data? A validation study [28]	2021	Sequential decision tree	N/A	Bivariate analysis, multivariate analysis	Medical
Fake It Till You Make It: Guidelines for Effective Synthetic Data Generation [2]	2021	DS, SDV, Synthpop	N/A	Propensity score pMSE, Accuracy	N/A
Generating and evaluating cross-sectional synthetic electronic healthcare data... [29]	2021	N/A	N/A	Univariate/Multivariate distance, Privacy preservation	Medical
Synthetic data use: exploring use cases to optimise data utility [9]	2021	N/A	N/A	Distribution comparison, Hellinger distance, Accuracy, Bivariate correlation, Area under the receiver operating characteristic	N/A
Utility Metrics for Evaluating Synthetic Health Data Generation Methods: Validation Study [30]	2022	N/A	N/A	Maximum mean discrepancy, Hellinger distance, Wasserstein distance	Medical
A multi-dimensional evaluation of synthetic data generators [31]	2022	N/A	N/A	Attribute fidelity, Bivariate fidelity, Population difelity, Application fidelity	N/A
How Faithful is your Synthetic Data? Sample-level Metrics for Evaluating and Auditing Generative Models [32]	2022	GAN, VAE, WGAN-GP, ADS-GAN	N/A	Alpha-precision, Beta-recall, Authenticity	N/A

- The Minkowski coefficient $p \in [0.1, 2.0]$ which controls the calculation of distance.

$$d(x, y, \Delta) = \left(\sum_{i=1}^n \Delta(x_i, y_i)^p \right)^{1/p} \quad (1)$$

- Iteration parameter $l = 6$ which is used to find the parameter for distance calculation.
- The number of neighbors $k \in [5, 22]$ during the execution of k -nearest neighbors algorithm.

In the distance calculation, Δ is the specific function to measure the distance between two variables. Howso Engine implements the Lukaszuk-Karmowski metric (LK metric) for the Laplace distribution [17]. Specifically,

$$\Delta(x_i, y_i) = |x_j - x_i| + \frac{1}{2} e^{-\frac{|x_j - x_i|}{b}} (3b + |x_j - x_i|) \quad (2)$$

The Laplace distribution is preferred here since it makes entropy-minimizing assumptions about the underlying data, and more performant than Gaussian distribution. In the above calculation, b is the surprisal that needs to be found through multiple iterations for each dataset. Specifically, the initial b is set to $1/k$ where k is the number of neighbors used in k -nearest neighbor algorithm. The first iteration finds the local neighbors by traditional Minkowski distance since no information in the initial round. In the end of the initial round, we can update the surprisal b :

- For the numerical feature k , Howso Engine uses mean absolute deviation (MAE):

$$b = \frac{1}{n} \sum_{i=1}^n |x_i - \mu_k| \quad (3)$$

- For the symbolic feature, Howso Engine uses mode and accuracy instead of mean and MAE.

Once the value of b is stabilized (which we found b tends to stable in 6 iterations), Howso Engine iteratively picks one random feature to synthesize. Specifically,

- The values in the first picked feature will be synthesized on the basis of the global histogram (for nominal features) or the global Laplace distribution (for continuous features).
- For all subsequent features, values are generated based on the distribution in the k nearest neighbors to the partially synthesized cases.

The above process is done m times per feature where m is the number of synthetic instances that request to be generated.

IV. RECURSIVE RANDOM PROJECTION

Recursive random projection, as its name shows, projects the high-dimensional data into different low-dimensional clusters by using the random pivot selection procedure recursively [38]. In the synthetic data generation, the relationships between different features are hard to capture. Previous literature uses causal graph [15] or covariance metrics [16] to explore the connections between different features. In this study, we adopt the random projection to split the data into different clusters, which each cluster capture the data points that have similar feature patterns. In this section, we will introduce our design of random projection framework from the following two aspect:

- First, the **cluster** algorithm which split the data into different clusters.
- Second, the **mutation & crossover** operators which mutate the data points in the same cluster to generate synthetic data points that have strong connection to the original data points.

The **cluster** algorithm aims to split the data points into different clusters, which each cluster will contain data points with similar patterns. To achieve that, we utilize the FASTMAP random projection algorithm [13]. With a set of data points, FASTMAP uses the cosine rule to project the data points into the hyperplane formed by two farthest points. More specifically, with two farthest points a and b , any third point c from the set of data points can be mapped into the line connecting a and b by

$$x = (a^2 + c^2 - b^2)/(2c) \quad (4)$$

Algorithm 1 shows the recursive clustering procedure, which in each function call, the algorithm first check if the number of current candidates in \mathcal{D} less than the threshold t (line 1). If the number does not reach the threshold, then stop the recursive. Otherwise, the **split** function will bi-cluster the current candidates, and perform the **cluster** algorithm again to the two sub-clusters (line 2-5). Algorithm 2 illustrates in detail how the **split** function works. More specifically, it firstly picks a random pivot, and finds two farthest points based on the random pivot (line 1-4). After that, as we stated above, it uses the cosine rule to map all other data points to the line formed by two farthest point (line 5-10). Finally, it returns two subsets based on the distance calculated by cosine rule (line 11-13).

Algorithm 1 cluster: The Overall Recursive Clustering Structure Inside the Random Projection. \mathcal{D} is the Set of Candidates and t is the Stopping Criteria. It Outputs a Tree Structure \mathcal{T} That the Nodes in the i^{th} Depth are Bi-Clustered into the $(i + 1)^{\text{th}}$ Depth

Require: $\mathcal{D}, t, \mathcal{T}, \mathcal{N}$

- 1: **if** $|\mathcal{D}| > t$ **then**
 - 2: $\mathcal{D}_{east}, \mathcal{D}_{west} = \mathbf{split}(\mathcal{D})$ ▷ bi-cluster
 - 3: $\mathcal{N}_{left.value}, \mathcal{N}_{right.value} = \mathcal{D}_{east}, \mathcal{D}_{west}$
 - 4: **cluster**($\mathcal{D}_{east}, t, \mathcal{T}, \mathcal{N}_{left.value}$) ▷ left node recursive
 - 5: **cluster**($\mathcal{D}_{west}, t, \mathcal{T}, \mathcal{N}_{right.value}$) ▷ right node recursive
 - 6: **end if**
-

Algorithm 2 split: Split a Set of Candidates \mathcal{D} into Two Subsets by Using the FASTMAP Technology [13], [14]

Require: \mathcal{D}

- 1: $rand = \mathbf{random}(1, |\mathcal{D}|)$
 - 2: $pivot = \mathcal{D}(rand)$ ▷ pick a random point as pivot
 - 3: $p_E = \mathbf{mostDistance}(\mathcal{D}, pivot)$ ▷ farthest point to pivot
 - 4: $p_W = \mathbf{mostDistance}(\mathcal{D}, p_E)$ ▷ farthest point to east
 - 5: $c = \mathbf{distance}(p_E, p_W)$ ▷ Similarity measure as distance
 - 6: **for** $idx = 1 : |\mathcal{D}|$ **do**
 - 7: $a = \mathbf{distance}(\mathcal{D}(idx), p_E)$
 - 8: $b = \mathbf{distance}(\mathcal{D}(idx), p_W)$
 - 9: $\mathcal{D}(idx).d = (a^2 + c^2 - b^2) / (2c)$ ▷ cosine rule
 - 10: **end for**
 - 11: $sorted = \mathbf{sort}(\mathcal{D}.d)$ ▷ Sort all points via distance
 - 12: $\mathcal{D}_E = \mathcal{D}[0.5*\mathbf{size}(sorted)]$
 - 13: $\mathcal{D}_W = \mathcal{D}[0.5*\mathbf{size}(sorted):]$
-

Traditional clustering algorithms require $O(N^2)$ calculations to fully split the data. However, random projection can achieve that in $O(2N)$, which is much faster than traditional clustering algorithms. This is the intuition we use random projection to reduce the scalability on generating synthetic data.

After the random projection returns the clusters, we then utilize the **mutation and crossover** operators used in the differential evolution algorithm to generate the synthetic data [39]. In the differential evolution algorithm, a better solution can be found in the current region of data points by mutating the data points as follow

$$y_{\text{new}_i} = \begin{cases} x_{1_i} + F * (x_{2_i} - x_{3_i}), & \text{if } r_i < CR \text{ or } k = R \\ x_{\text{old}_i}, & \text{O.W.} \end{cases} \quad (5)$$

where x_{old_i} is the point from the original set of candidates, and x_1, x_2 , and x_3 are three other random points from the set of candidates. F is the difference scaling factor from 0 to 1, and CR is the crossover probability that also from 0 to 1. Large F indicates more scaling on the difference of two data points during the mutation, and large CR indicates more probability the new candidate has a new value in each index. R is a random index such that the value in that index must be mutated. This can prevent the duplicated new candidate when CR is very small. We hypothesis that the synthetic data generated by this mutation and crossover operator can capture the feature information from the original data in each cluster since each cluster includes the data points that are close to each other.

V. EXPERIMENTAL METHODS

In this section, we will briefly explain two state-of-the-art synthetic data generation algorithms we compared to.

A. DATASYNTHESIZER

DataSynthesizer is proposed by Ping et al. [15]. Their framework can handle two different attribute mode. One is *independent attribute mode*, which each feature is treated individually. Another one is *correlated attribute mode*, which the causal graph is used to describe the relationships between each feature. First of all, DataSynthesizer implements a module called DataDescriber to capture the data type for each feature, as well as its distribution and correlation. Also, it will add the noise to the data distribution to preserve the privacy. Secondly, DataGenerator module will generate the synthetic data based on the attribute mode.

More specifically, for the *independent attribute mode*, DataDescriber performs the frequency-based estimation of the unconditioned probability distributions of each attribute [15]. To preserve the privacy, the noise $Lap(\frac{1}{n\epsilon})$ will be added to the distribution where n is the size of the inputs and ϵ is set to 0.1 by default. DataGenerator will then uses the distribution to generate the synthetic data for each feature.

On the other hand, the *correlated attribute mode* is quite different to the *independent attribute mode*. The

GreedyBayes algorithm is utilized to construct the causal graph for all the features. GreedyBayes is a kind of greedy selection algorithm that select the highest correlated feature which maximize the mutual information to the subsets of features that have been visited. The noise $Lap(\frac{4(|A|-k)}{n\epsilon})$ will also be added to preserve the privacy. After the causal graph is generated, the algorithm will then use the knowledge from causal graph and the distribution to generate the synthetic data.

B. SYNTHETIC DATA VAULT

Synthetic Data Vault (SDV) is proposed by Patki et al. [16]. The basic intuition behind SDV is to use the Gaussian Copula as the generative model to synthesize the data based on the distribution and the covariance of the features. More specifically, with a dataset that has columns $\{c_1, c_2, \dots, c_n\}$, we use $\{f_1, f_2, \dots, f_n\}$ to express the cumulative distribution function for those columns. The Gaussian Copula will calculate the inverse cumulative distribution functions ϕ of the Gaussian distribution applied to the original cumulative distribution function $\{f_1, f_2, \dots, f_n\}$. Algorithm 3 shows how to use the Gaussian Copula to calculate the covariance matrix when given the dataset.

Algorithm 3 Gaussian Copula: The Gaussian Copula Algorithm for Analyzing the Distribution and Covariance of the Dataset [16]. Return the Covariance Matrix Σ

Require: $\mathcal{D} = \{d_1, d_2, \dots, d_p\}$, $\mathcal{F} = \{f_1, f_2, \dots, f_n\}$

1: **for** d_i in \mathcal{D} **do**

2: $Y_i = [\Phi^{-1}(f_0(d_{i0})), \Phi^{-1}(f_1(d_{i1})), \dots, \Phi^{-1}(f_n(d_{in}))]$

3: **end for**

4: $\Sigma = \text{computeCovariance}(\{Y_1, Y_2, \dots, Y_n\})$

To generate synthetic data, for a given row, SDV will generate the synthetic value based on the feature distributions. Moreover, if there is information on the other features, then the covariance information calculated from Gaussian Copula will also be used along with the feature distribution to generate the synthetic data.

We use the SDV public API³ to implement the SDV in our study. Note that their online API also includes the Machine Learning based generative model and the GAN based generative model. Hence, in our study, we compare our methods to both three SDV generative models.

C. COMPUTATIONAL COST

As to the computational cost of the above, this divides into (a) the generation cost and (b) the cost of using the generated data during machine learning. In all our studies, we generated data of similar size to the original data. This explains why, in these studies, the learning cost was observed to be similar for the original and the synthetic data.

As to the generation cost, that varied wildly depending on the synthesis method. For example, recursive random projects

³<https://github.com/sdv-dev/SDV/tree/main>

TABLE 2. Summary of benchmarks used in our experiment.

Benchmark	# Rows	# Cols	Task
glass	203	10	multi-classification
596_fri_c2_250	250	6	regression
breast_cancer	286	10	classification
cars	392	9	multi-classification
irish	500	6	classification
522_pm10	500	8	regression
profb	673	10	classification
tic_tac_toe	958	10	classification
churn	5000	21	classification
adult	48842	15	classification

ran in nearly constant time across all these data sets while anything using GAN exhibited a runtime that was exponential on model size.

For more details on the computational cost of generation, see the results of **RQ3**.

VI. EXPERIMENTAL SETUP

In this section we will illustrate following things: (a) the benchmarks, (b) the evaluation metrics, and (c) the statistical analysis procedure.

A. BENCHMARK

Table 2 shows 10 machine learning datasets from the Penn Machine Learning Benchmark (PMLB) datasets⁴ [40]. To select those ten datasets, we firstly counted the number of instances in all datasets in PMLB. After that, we group those datasets in the 10 clusters based on the number of instances, and randomly pick one dataset from each cluster. This step can ensure that the benchmarks used in our experiment have different size, and thus we can analyze the runtime of each synthetic data generator more empirically. Moreover, we manually inspect the selected datasets, and replace some of them to the one with more privacy concerns. Please note that the “privacy concerns” means some features are very informative to identify the individuals. In this case, the datasets with those features will not be able to share in the real world. As we stated in the Introduction section, the synthetic data is designed to replace those datasets with sensitive information, and that is the reason we will make this replacement to our selected benchmarks. However, some datasets we used may not have sensitive features, and there is no equivalent one to replace with, then we will keep using these datasets and assume they have the privacy concerns as other datasets have. The last point is that we also keep some datasets which the task is either multi-class classification or regression, to compare the synthetic data generation performance in both regression and classification task. The summary of our selected benchmarks is shown in Table 2.

B. EVALUATION METRICS

The **Metric(s)** column in Table 1 shows different metrics have been used in the past literature. In our study, we consider the

⁴<https://github.com/EpistasisLab/pmlb>

most general case of data synthesis which the original data is used to train the machine learning models but cannot be shared due to the privacy issue. In this scenario, we consider privacy metric to evaluate privacy issue, information coefficients metric and distribution metric to evaluate the data distribution, and model comparison metric to evaluate the quality of synthetic data.

- **Privacy metric** which evaluates the privacy preservation of the synthetic data.
- **Informational coefficients metric** which evaluates the statistical information of the synthetic data.
- **Distribution metric** which evaluates the synthetic data through joint distribution.
- **Model performance metric** which build the machine learning models on the original data and the synthetic data and test their performance on the original test data.

In our study, we evaluate each algorithm through these four metrics. The details of each metric is explained in the following subsections. All these metrics are used to validate the synthetic data.

1) PRIVACY PRESERVATION

Privacy preservation evaluates the privacy of synthetic data. It checks if the distance from a synthetic data point to the density of the region of its nearest original neighbor is small or not. More specifically, to evaluate the privacy preservation score of a synthetic data point x_{syn} , we first find its nearest original neighbor x_{ori} , and calculate the distance from x_{syn} to x_{ori} as d . After that, we find k -nearest original neighbors of x_{ori} , and calculate the minimum distance d_{min} between any two of the points in the group of that k -nearest neighbors. The final score of privacy preservation is the minimum distance ratio of distance d and the distance d_{min} . Though larger score indicates better privacy preservation, the minimum distance ratio of 1 already indicates that the synthetic point is located outside of the density region of its nearest original neighbor. Hence any value greater or equal than 1 can indicate a good privacy preservation.

2) STATISTICAL SIMILARITY

Statistical similarity compares the synthetic data to the original data through statistical measurements. The statistical measurements are split into three parts.

- **Central tendency** which includes *mean, mode, median, 25th percentile, 75th percentile, minimum, and maximum*.
- **Variability of dispersion** which includes *entropy, Kurtosis, mean absolute deviation, standard deviation, skew, and variance*.
- **Frequency distribution** which describes the *uniqueness* of the data.

All those metrics are calculated in both the synthetic dataset and the original dataset. For all pairs of score in $\{S^{ori}, S^{syn}\}$ for a certain feature, we calculate the **SMAPE**. Specifically, SMAPE is the symmetric mean absolute

percentage error which is an accuracy measure based on relative error [41], and is calculated as follow

$$SMAPE = \frac{1}{n} \sum_{t=1}^n \frac{|s_t^{syn} - s_t^{ori}|}{(|s_t^{ori}| + |s_t^{syn}|)/2} \quad (6)$$

where s_t^{syn} and s_t^{ori} are the scores of statistical measurement t in synthetic data S^{syn} and original data S^{ori} . The final statistical similarity score is the average SMAPE of all features. The smaller score in this metric means less difference between synthetic data and original data on the statistical measurements.

3) MARGINAL DISTRIBUTION SIMILARITY

Marginal distribution similarity evaluates the synthetic data through the marginal distribution. The marginal distribution of the numeric feature is estimated through the KNN density estimation and the marginal distribution of the nominal feature is estimated using normalized value counts. The score of marginal distribution similarity is then calculated by comparing the estimated distributions using *Jensen-Shannon divergence*. Specifically, Jensen-Shannon divergence measures the similarity between two probability distribution [42], and is calculated as follow

$$\delta_{JS}(\mathcal{P}, \mathcal{Q}) = \delta_{KL}(\mathcal{P}||\mathcal{M}) + \delta_{KL}(\mathcal{Q}||\mathcal{M}) \quad (7)$$

where $\mathcal{M} = (\mathcal{P} + \mathcal{Q})/2$. In the above formula, δ_{KL} is the Kullback-Leibler divergence [43] which measures the distance of probability distribution P to the reference probability distribution Q by

$$\delta_{KL}(\mathcal{P}||\mathcal{Q}) = \int_{-\infty}^{\infty} p(x) \log\left(\frac{p(x)}{q(x)}\right) dx \quad (8)$$

The smaller Jensen-Shannon divergence value indicates that two probability distributions are similar.

4) MODEL COMPARISON

Model comparison evaluates the quality of synthetic data on training the machine learning models. More specifically, both the original data and synthetic data is split into 80% of training set and 20% of test set. After that, the LightGBM classification model or regression model is trained on the **synthetic training set**, and test on the **original test set**. For the regression task, we report *RMSE*, R^2 , and *Spearman correlation coefficient*. Specifically

- **RMSE** is the root mean square error which measures the average difference between the predicted values from a regression model and the actual values [44]. It is calculated as follow

$$RMSE = \sqrt{\frac{\sum_{i=1}^N (\hat{y}_i - y_i)^2}{N}} \quad (9)$$

- R^2 (or say coefficient of determination) represents how well the data fit the regression model [45]. Specifically, let \bar{y} be the mean of all observation y , and \hat{y} be the

predicted value from the model, R^2 is calculated as follow

$$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2} \quad (10)$$

- **Spearman correlation coefficient** is the statistical measure which check the linear correlation between two populations [46]. Given a pair of same feature from original dataset and synthetic dataset X_{ori} and X_{syn} , the score of Spearman is calculated by

$$\rho_{X_{ori}, X_{syn}} = \frac{\mathbb{E}((X_{ori} - \mu_{X_{ori}})(X_{syn} - \mu_{X_{syn}}))}{\sigma(X_{ori})\sigma(X_{syn})} \quad (11)$$

where μ is the mean and σ is the standard deviation of the population. The overall score of Spearman is the geometric mean of all features. We expect higher Spearman value since score 1 means two populations perfectly fit the linear correlation.

And for the classification task, we report *accuracy*, *precision*, *recall*, and *Matthews correlation coefficient*. Specifically, in the classification task, we notate TP , TN , FP , FN as the value of *true positive*, *true negative*, *false positive*, and *false negative* returned from the confusion matrix,

- **Accuracy** evaluates the ratio of number of *correct predictions* over the *total number of predictions*.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN} \quad (12)$$

- **Precision** evaluates the ratio of number of *correct positive predictions* over the total number of *predicted positive cases*.

$$Precision = \frac{TP}{TP + FP} \quad (13)$$

- **Recall** evaluates the ratio of number of *correct positive predictions* over the total number of *actual positive cases*.

$$Recall = \frac{TP}{TP + FN} \quad (14)$$

- **Matthews correlation coefficient** evaluates the prediction performance by summarizing the entire confusion matrix.

$$MCC = \frac{TN * TP - FN * FP}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (15)$$

For all those classification metrics, the value close to 1 will indicate the better performance.

C. SCOTT-KNOTT ANALYSIS

To perform significant test, we utilize the Scott-Knott Analysis. We choose Scott-Knott analysis since (a) it is fully non-parametric and (b) it can reduce the potential error during the analysis with only at most $O(\log_2(N))$ statistical tests for the $O(N^2)$ analysis.

TABLE 3. Privacy preservation score. Higher values are better. The dark gray cell marks the algorithm in the highest rank by Scott-Knott analysis, and the light gray cell marks the algorithm in the second highest rank.

Algorithm	Benchmarks										Wins
	glass	596_fri_c2_250	breast_cancer	cars	irish	522_pm10	profb	tic_tac_toe	churn	adult	
DataSynthesizer-Ind	1.48	0.53	1.05	3.95	6.00	1.34	1.10	0.00	1.89	17.31	9
DataSynthesizer-Cor	2.24	1.22	2.00	3.78	1.00	1.48	1.07	0.00	1.91	1.68	10
Synthetic Data Vault-ML	0.88	0.60	0.00	1.17	0.15	0.80	0.50	0.00	0.57	0.44	2
Synthetic Data Vault-GC	1.44	0.64	0.00	0.89	0.00	0.53	0.33	0.00	1.71	0.38	3
Synthetic Data Vault-GAN	1.44	0.74	0.00	1.68	0.00	0.80	0.54	0.00	0.76	0.00	3
Recursive Random Projection	0.45	0.46	0.00	0.29	0.00	0.56	0.55	0.00	0.70	0.00	1
Howso Engine	1.00	1.00	0.00	1.00	1.25	1.00	0.45	0.00	0.82	0.01	6

TABLE 4. Statistical similarity score. The overall score is evaluated through SMAPE in equation 6. Smaller values are better. The dark gray cell marks the algorithm in the highest rank by Scott-Knott analysis, and the light gray cell marks the algorithm in the second highest rank.

Algorithm	Benchmarks										Wins
	glass	596_fri_c2_250	breast_cancer	cars	irish	522_pm10	profb	tic_tac_toe	churn	adult	
DataSynthesizer-Ind	0.77	0.84	0.50	0.35	0.55	0.50	0.56	0.40	0.60	0.43	0
DataSynthesizer-Cor	0.74	0.75	0.14	0.35	0.44	0.55	0.45	0.17	0.66	0.44	0
Synthetic Data Vault-ML	0.40	0.51	0.10	0.22	0.07	0.39	0.32	0.17	0.32	0.33	2
Synthetic Data Vault-GC	0.45	0.40	0.10	0.13	0.15	0.47	0.13	0.02	0.32	0.32	4
Synthetic Data Vault-GAN	0.32	0.66	0.12	0.22	0.09	0.45	0.25	0.05	0.20	0.16	0
Recursive Random Projection	0.13	0.38	0.12	0.12	0.17	0.29	0.26	0.00	0.32	0.11	4
Howso Engine	0.16	0.45	0.06	0.08	0.16	0.26	0.15	0.00	0.19	0.13	5

In our experiment, we repeat each synthetic data generation algorithm 10 times since they are stochastic. With a list of candidates C , where each candidate c_i is a list of results from 10 repeats for a certain synthetic data generation algorithm, Scott-Knott recursively partitions C into two sub-lists C_1 and C_2 . The split is based on the expected mean value before and after the division, which the goal is to maximize the expected mean value [47], [48], [49]. The delta of expected mean value before and after the split is calculated as follow:

$$E(\Delta) = \frac{l(C_1) \cdot |\mu(C_1) - \mu(C)| + l(C_2) \cdot |\mu(C_2) - \mu(C)|}{l(C)} \quad (16)$$

where l is the length function to count the length of each list.

After the split is finished, Scott-Knott will then utilize the Cliff's Delta procedure to check if two sub-lists differ significantly by

$$\Delta = \frac{\sum_{x \in C_1} \sum_{y \in C_2} \begin{cases} +1, & \text{if } x > y \\ -1, & \text{if } x < y \\ 0, & \text{if } x = y \end{cases}}{|C_1||C_2|} \quad (17)$$

More specifically, Cliff's Delta estimates the probability that a value in the sub-list C_1 is greater than a value in the sub-list C_2 , and then minus the reverse probability [50]. If Cliff's Delta is equal or greater than 0.147 (i.e. $\Delta \geq 0.147$) then it is not a small effect [51].

VII. RESULTS

In this section, we present our experimental results, and answer RQs based on the results.

RQ1: When considering the privacy, which synthetic data generation algorithm can generate the synthetic data with the highest privacy preservation? To evaluate the synthetic data in terms of privacy, we implement the *Privacy Preservation* metric to evaluate the privacy score. As we described in Section VI-B, privacy preservation evaluates the distance from the synthetic data point to the density of the region of the k -nearest neighbors of the closest original data point. We calculate such distance ratio for all synthetic data points, and take the geometric mean as the final privacy preservation score for the synthetic dataset. As mentioned in Section VI-B1, higher score refers to better privacy. However, the score of 1 already indicates that the synthetic data lies out of the density region of its closest original data point. Hence, in this research question, we consider an approach with good privacy preservation if its score is greater or equal than 1. Table 3 presents the experimental results. At first glance, DataSynthesizer has higher raw values than any other algorithms. However, those methods with score equal or greater than 1 will not induce the privacy issue in the real world application. Hence, if a method is not in the highest rank through statistical testing but its raw value is equal or greater than 1, we still consider it as a good method and marked them with grey color in Table 3.

As we can see, synthetic data generated by DataSynthesizer has the highest privacy score. After that, Howso engine has promising performance on 6 case studies. SDV and recursive random projection do not perform well in terms of privacy preservation. Hence, our answer to RQ1 is

DataSynthesizer and Howso engine are two promising algorithms that can generate synthetic data with good privacy preservation score. If we only consider privacy, the result would recommend DataSynthesizer.

TABLE 5. Marginal distribution similarity score. The overall score is the mean of the Jensen-Shannon divergence scores calculated by equation 7 through all features. The dark gray cell marks the algorithm in the highest rank by Scott-Knott analysis, and the light gray cell marks the algorithm in the second highest rank.

Algorithm	Benchmarks										Wins
	glass	596_fri_c2_250	breast_cancer	cars	irish	522_pm10	profb	tic_tac_toe	churn	adult	
DataSynthesizer-Ind	0.37	0.05	0.34	0.34	0.39	0.39	0.13	0.10	0.07	0.40	2
DataSynthesizer-Cor	0.36	0.07	0.35	0.31	0.34	0.35	0.12	0.10	0.07	0.52	2
Synthetic Data Vault-ML	0.34	0.04	0.23	0.39	0.25	0.47	0.10	0.09	0.04	0.48	6
Synthetic Data Vault-GC	0.30	0.03	0.23	0.30	0.25	0.42	0.10	0.09	0.04	0.46	7
Synthetic Data Vault-GAN	0.26	0.05	0.23	0.29	0.24	0.45	0.11	0.09	0.04	0.46	7
Recursive Random Projection	0.42	0.04	0.23	0.40	0.24	0.67	0.10	0.09	0.07	0.54	4
Howso Engine	0.27	0.05	0.23	0.39	0.24	0.54	0.10	0.09	0.04	0.53	6

However, as we will discuss in the RQ5, the evaluation of synthetic data cannot only concentrate on the privacy, and therefore, we recommend Howso engine more than DataSynthesizer.

RQ2: Which synthetic data generation algorithm can generate data that has higher similarity to the original data? Similarity is another important measurement on checking if synthetic data can provide same information as original data does. We implement *statistical similarity score* and *marginal distribution similarity score* to evaluate the synthetic data. Specifically,

- Statistical similarity score evaluate the synthetic data by comparing its statistical measurements to the original data for each feature.
- Marginal distribution similarity score evaluates the synthetic data from the estimated marginal probability distribution of each feature.

Table 4 shows the statistical similarity score, which is calculated by evaluating the SMAPE of 14 different statistical measurements in synthetic data and original data (described in §VI-B2). From this table, we can find Howso engine and recursive random projection algorithm have better performance in more case studies than any other algorithms. This indicates that the synthetic data generated by Howso engine and recursive random projection algorithm can better capture the statistical patterns in the original data. More specifically, the large SMAPE scores on DataSynthesizer reduce its advantage on the privacy preservation since the synthetic data generated by it does not follow the statistical patterns in the original data.

Table 5 presents the marginal probability distribution similarity score, which is calculated by Jensen-Shannon divergence. Specifically, the marginal distribution is estimated through the KNN density estimation in numeric features and normalized value counts in categorical features. Different to the statistical similarity, this metric concentrates more on the estimated local probability distribution of each feature and check if the probability distribution of the same feature from synthetic data and original data is similar or not through Jensen-Shannon divergence. As we can see, most of the algorithms get good scores except DataSynthesizer, which is also shown to have worse score in statistical similarity score.




FIGURE 1. Runtime of different synthetic data approaches.

For all other algorithms except DataSynthesizer, we prefer Synthetic Data Vault with Gaussian Copula, recursive random projection, and Howso engine since they both perform well in two similarity metrics. Hence, our answer for RQ3 is

Considering two different similarity measurements, Synthetic Data Vault with Gaussian Copula, recursive random projection, and Howso engine achieve good performance in both two metrics. Thus, we recommend these three algorithms when evaluating the similarity.

RQ3: When the machine learning model is trained on the synthetic data, can the model achieve compatible performance with those trained on the original data? The model performance score is a very important indicator to evaluate the synthetic data. We split both the original data and the synthetic data to 80% training set and 20% test set. Then we use the model trained on the synthetic training data to predict the test set split from the original data. As stated in Section VI-B, for the classification task we collect accuracy, precision, recall, and Matthews correlation coefficient, and for the regression task we collect RMSE, R^2 ,

TABLE 6. Model comparison score. Classification tasks are evaluated by accuracy, precision, recall, and MCC, while regression tasks are evaluated by RMSE, R², and spearman. The dark gray cell marks the algorithm in the highest rank and the light gray cell marks the algorithm in the second highest rank. Due to space reason, DataSynthesizer is simplified to DS, synthetic data vault is simplified to SDV, and recursive random projection is simplified to RRP.

Metric	Algorithm	Benchmarks										Wins
		glass	596_fri_c2_250	breast_cancer	cars	irish	522_pm10	prob	tic_tac_toe	churn	adult	
Accuracy	DS_Ind	0.15	-	0.31	0.44	0.57	-	0.33	0.35	0.80	0.65	0
	DS_Cor	0.32	-	0.64	0.20	0.66	-	0.53	0.66	0.70	0.75	0
	SDV_ML	0.34	-	0.67	0.70	0.79	-	0.67	0.63	0.85	0.76	1
	SDV_GC	0.41	-	0.71	0.65	0.77	-	0.67	0.65	0.85	0.74	2
	SDV_GAN	0.17	-	0.64	0.38	0.52	-	0.64	0.58	0.82	0.79	0
	RRP	0.71	-	0.72	0.84	0.96	-	0.64	0.77	0.87	0.82	5
	Howso Engine	0.76	-	0.74	0.86	0.94	-	0.67	0.85	0.90	0.81	8
Precision	DS_Ind	0.02	-	0.10	0.58	0.64	-	0.11	0.47	0.74	0.61	0
	DS_Cor	0.23	-	0.60	0.45	0.70	-	0.47	0.61	0.73	0.69	0
	SDV_ML	0.26	-	0.50	0.68	0.79	-	0.60	0.52	0.74	0.70	0
	SDV_GC	0.44	-	0.69	0.64	0.77	-	0.60	0.62	0.78	0.67	1
	SDV_GAN	0.22	-	0.58	0.42	0.56	-	0.60	0.58	0.81	0.79	0
	RRP	0.67	-	0.65	0.83	0.96	-	0.61	0.78	0.87	0.82	5
	Howso Engine	0.73	-	0.73	0.87	0.94	-	0.65	0.86	0.90	0.80	7
Recall	DS_Ind	0.15	-	0.31	0.44	0.57	-	0.33	0.35	0.80	0.65	0
	DS_Cor	0.17	-	0.64	0.20	0.66	-	0.53	0.66	0.70	0.75	0
	SDV_ML	0.34	-	0.67	0.70	0.79	-	0.67	0.63	0.85	0.76	1
	SDV_GC	0.41	-	0.71	0.65	0.77	-	0.67	0.65	0.85	0.74	2
	SDV_GAN	0.17	-	0.64	0.38	0.52	-	0.64	0.58	0.82	0.79	0
	RRP	0.71	-	0.72	0.84	0.96	-	0.64	0.77	0.87	0.82	5
	Howso Engine	0.76	-	0.74	0.86	0.94	-	0.67	0.85	0.90	0.81	8
MCC	DS_Ind	0.00	-	0.00	0.15	0.22	-	0.00	-0.13	-0.05	-0.06	0
	DS_Cor	0.06	-	0.03	-0.02	0.35	-	-0.08	0.10	-0.09	0.11	0
	SDV_ML	-0.02	-	0.00	0.31	0.51	-	0.07	-0.04	0.00	0.11	0
	SDV_GC	0.18	-	0.05	0.33	0.57	-	0.10	0.15	0.03	0.06	1
	SDV_GAN	0.07	-	0.00	-0.08	0.06	-	0.03	0.04	0.18	0.40	0
	RRP	0.60	-	0.12	0.65	0.92	-	0.13	0.44	0.31	0.50	6
	Howso Engine	0.64	-	0.34	0.70	0.86	-	0.17	0.67	0.50	0.45	6
RMSE	DS_Ind	-	1.23	-	-	-	1.42	-	-	-	-	0
	DS_Cor	-	1.13	-	-	-	1.05	-	-	-	-	0
	SDV_ML	-	0.91	-	-	-	0.84	-	-	-	-	0
	SDV_GC	-	0.87	-	-	-	2.35	-	-	-	-	0
	SDV_GAN	-	1.20	-	-	-	1.05	-	-	-	-	0
	RRP	-	0.53	-	-	-	0.73	-	-	-	-	1
	Howso Engine	-	0.46	-	-	-	0.72	-	-	-	-	2
R ²	DS_Ind	-	-0.59	-	-	-	-1.51	-	-	-	-	0
	DS_Cor	-	-0.33	-	-	-	-0.39	-	-	-	-	0
	SDV_ML	-	0.18	-	-	-	-0.03	-	-	-	-	0
	SDV_GC	-	0.14	-	-	-	-6.88	-	-	-	-	0
	SDV_GAN	-	-0.74	-	-	-	-0.61	-	-	-	-	0
	RRP	-	0.68	-	-	-	0.31	-	-	-	-	1
	Howso Engine	-	0.77	-	-	-	0.29	-	-	-	-	2
Spearman	DS_Ind	-	0.19	-	-	-	0.36	-	-	-	-	0
	DS_Cor	-	0.49	-	-	-	0.48	-	-	-	-	0
	SDV_ML	-	0.74	-	-	-	0.64	-	-	-	-	0
	SDV_GC	-	0.74	-	-	-	0.51	-	-	-	-	0
	SDV_GAN	-	0.48	-	-	-	0.47	-	-	-	-	0
	RRP	-	0.92	-	-	-	0.77	-	-	-	-	2
	Howso Engine	-	0.94	-	-	-	0.78	-	-	-	-	2

and Spearman correlation coefficient. Table 6 shows the scores of these metrics. As we can see, in all metrics, Howso engine and recursive random projection framework get significant higher performance than other two state-of-the-art algorithms. Therefore, we conclude that the synthetic data generated by these two algorithms can be applied to the real world machine learning models without loss on the information. Based on above, our answer to RQ3 is:

Howso engine and recursive random projection algorithm can generate synthetic data that does not loss any original information when training on the machine learning models.

RQ4: Which algorithm has the best scalability? To answer this question, we record the runtime of each generation algorithm. Figure 1 shows the runtime for each generation algorithm. To visualize the scalability, we first record the actual runtime for each algorithm in each case study. The scatters shows the point (x, y) where x is the size of the dataset, and y is the runtime. After that, we plot the best fitting polynomial curve, which presents the trending of the runtime when the size of the dataset is increasing. As we can see, when the size of the dataset is small, all algorithms have similar runtime. However, when the size of the dataset becomes

larger (e.g. 40k+ rows), the runtime of Howso engine and GAN based algorithm increases exponentially (i.e. the blue line and the pink line). To the contrary, our proposed recursive random projection algorithm, along with DataSynthesizer and Synthetic Data Vault with Gaussian Copula, have better efficiency, which the best fitting curve of runtime is close to linear with low slope even though the size of the dataset goes exponentially large. Hence, our answer to RQ4 is:

Our proposed recursive random projection framework, along with DataSynthesizer and Synthetic Data Vault with Gaussian Copula have the best scalability even when the size of the dataset is very large.

RQ5: What recommendation can we provide from analyzing the conclusions from RQ1 to RQ4? Evaluating synthetic data needs to consider multi-dimensional criteria. To empirically evaluate all performance from RQ1 to RQ4, we transfer each criteria to a 0-1 range and use the radar chart to visualize the performance. Specifically

- For *privacy preservation*, all scores greater than 1 are treated as 1, and all other values will not be modified. The overall score of an algorithm will be the geometric mean of its scores in 10 case studies (If an algorithm




FIGURE 2. Radar chart of scores in multi-dimensional criteria.

has score 0, then we use 0.001 when calculating the geometric mean).

- For *statistical similarity* and *marginal distribution similarity*, the scores are 1 minus the values in Table 4 since smaller is better in these two metrics. The overall score is also the geometric mean through 10 case studies.
- For *model comparison*, the negative correlation coefficient will be treated as 0.001 since negative value means no correlation. We first calculate the geometric mean of each algorithm for each individual metric, and then calculate the geometric mean again for 7 metrics.
- For *scalability*, we transfer the actual runtime by minmax scalar in each case study. Then the overall score will also be the geometric mean through 10 case studies for each individual algorithm.

We select 4 algorithms to be presented in our chart. The first one is the correlation mode of DataSynthesizer. We do not choose independent mode since it is worse than correlation mode in privacy preservation, and similar to correlation mode in other metrics. The same procedure is applied to Synthetic Data Vault and the Gaussian Copula mode is selected. Hence in the chart, we have above two algorithms plus recursive random projection and Howso engine.

Figure 2 presents the radar chart of (a) DataSynthesizer with correlated mode, (b) SDV with Gaussain Copula, (c) Recursive Random Projection, and (d) Howso Engine. As we can see, the area covered by DataSynthesizer and SDV (blue dash line and orange dash line correspondingly) are obviously less than the area covered by Recursive Random Projection and Howso Engine. Hence, we offer two recommendations based on the analysis, and answer RQ5:

- **As seen in Figure 1, runtime between different methods does not differ too much. Hence, if the scalability is not an issue, then we recommend Howso Engine since it can achieve higher accuracy score and promising privacy preservation score.**
- **However, if the dataset that needs to be synthesized is very large, and scalability becomes more important, we recommend our proposed Recursive Random**

Projection since it scales very fast and generate high accurate synthetic data.

VIII. THREAT TO VALIDITY

A. CONSTRUCT VALIDITY

mainly related to the different parameters setting and model construction which causes the different outcome. In our study, the threat of construct validity can happen in (a) the parameter choice in different generation models and (b) the choice of settings when evaluate the synthetic data. For example, the cluster size in our proposed random projection based framework can influence the performance of generating synthetic data. We empirically evaluate the different sizes of cluster and choose 12 as the parameter in our experiment. For another example, when we evaluate the synthetic data by training the machine learning model on it, we use the 80% train test split on both original data and synthetic data. Different train test split ratio may cause different final outcome. The train test split ratio we used is the default setting which highly used in other machine learning studies. To reduce the threat, we build our experimental scripts as a python package, and allow researchers to replicate our experiment with their own parameter choice.

B. CONCLUSION VALIDITY

refers to the threat that caused by applying different evaluation metrics when make the conclusion. To mitigate this threat, we apply four metrics (privacy preservation, descriptive statistics, marginal probability, and model comparison) which include most of the evaluation aspects of synthetic data in the past literature. Researchers may expect different conclusion when applying different metrics on our methods.

C. INTERNAL VALIDITY

focuses on the correctness of treatment caused the outcome. To reduce the effect caused by this threat, we collect ten highly used machine learning benchmarks from PMLB and run all algorithms on those ten benchmarks. Also, we control the size of synthetic data and make it equal to the size of the original data.

D. EXTERNAL VALIDITY

indicates the threat of applying this experiment to other fields. To mitigate this threat, the experimental goal of our study focuses on the machine learning object, which is one of the most well-known regions in the real world application. Moreover, our replication package can be applied to different datasets, which allows researchers to explore other real world application with our scripts.

IX. CONCLUSION

In this study, we explore the synthetic data generation algorithms and discuss the different validation metrics. We proposed recursive random projection based generator, and compare it to (a) two state-of-the-art generation

algorithms DataSynthesizer and Synthetic Data Vault, and (b) the Howso Engine from our industrial partner. We evaluate those synthetic data generators from (a) privacy preservation, (b) statistical similarity and marginal probability distribution similarity, (c) model performance comparison, and (d) scalability.

In privacy measurement, we find DataSynthesizer has the highest privacy preservation score through all case studies. Howso engine ranks behind the DataSynthesizer. However, when considering the similarity measurements and model comparison score, Howso engine and recursive random projection based framework get far more higher score than DataSynthesizer. Hence, we conclude that DataSynthesizer adds too much noise to the synthetic data which the fake points lie out of the original distribution and patterns.

By conducting the empirically analysis to five evaluation criteria, from the radar chart in RQ5, we offer two recommendations:

- If scalability is not an issue, then we recommend Howso Engine which has highest accuracy performance and promising privacy score.
- However, when the dataset is large enough, which will cause the scalability issue, we recommend recursive random projection based framework since it scales fast and can achieve highest accuracy performance.

In the future work, we will explore more synthetic data generation algorithms, as well as more benchmarks. Moreover, current recursive random projection based framework does not particularly add differential privacy operators. In the future, we will design such operators based on the condition in each cluster, and improve its privacy preservation score.

ACKNOWLEDGMENT

This work is a report of an industrial collaboration between Howso and NCState. In this work, Howso funded NCState to comparatively assess numerous data synthesis methods. The authors assert that the conclusions made here are the product of NCState and were not altered by their Howso collaborators.

REFERENCES

- [1] J. Dahmen and D. Cook, "SynSys: A synthetic data generation system for healthcare applications," *Sensors*, vol. 19, no. 5, p. 1181, Mar. 2019.
- [2] F. K. Dankar and M. Ibrahim, "Fake it till you make it: Guidelines for effective synthetic data generation," *Appl. Sci.*, vol. 11, no. 5, p. 2158, Feb. 2021.
- [3] D. M. Smith, G. P. Clarke, and K. Harland, "Improving the synthetic data generation process in spatial microsimulation models," *Environ. Planning A, Economy Space*, vol. 41, no. 5, pp. 1251–1268, May 2009.
- [4] M. Hernandez, G. Epelde, A. Alberdi, R. Cilla, and D. Rankin, "Synthetic data generation for tabular health records: A systematic review," *Neurocomputing*, vol. 493, pp. 28–45, Jul. 2022.
- [5] S. A. Assefa, D. Dervovic, M. Mahfouz, R. E. Tillman, P. Reddy, and M. Veloso, "Generating synthetic data in finance: Opportunities, challenges and pitfalls," in *Proc. 1st ACM Int. Conf. AI Finance*, New York, NY, USA, Oct. 2020, pp. 1–8, doi: 10.1145/3383455.3422554.
- [6] F. Liu, Z. Cheng, H. Chen, Y. Wei, L. Nie, and M. Kankanhalli, "Privacy-preserving synthetic data generation for recommendation systems," in *Proc. 45th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, Jul. 2022, pp. 1379–1389.
- [7] D. Rankin, M. Black, R. Bond, J. Wallace, M. Mulvenna, and G. Epelde, "Reliability of supervised machine learning using synthetic data in health care: Model to preserve privacy for data sharing," *JMIR Med. Informat.*, vol. 8, no. 7, Jul. 2020, Art. no. e18910.
- [8] A. Yale, S. Dash, R. Dutta, I. Guyon, A. Pavao, and K. P. Bennett, "Generation and evaluation of privacy preserving synthetic health data," *Neurocomputing*, vol. 416, pp. 244–255, Nov. 2020.
- [9] S. James, C. Harbron, J. Branson, and M. Sundler, "Synthetic data use: Exploring use cases to optimise data utility," *Discover Artif. Intell.*, vol. 1, no. 1, p. 15, Dec. 2021.
- [10] B. Nowok, "Utility of synthetic microdata generated using tree-based methods," *UNECE Stat. Data Confidentiality Work Session*, Oct. 2015.
- [11] E. Fiestas, O. E. Ramos, and S. Prado, "RPA and L-system based synthetic data generator for cost-efficient deep learning model training," in *Proc. IEEE 3rd Eurasia Conf. IoT, Commun. Eng. (ECICE)*, Oct. 2021, pp. 645–650.
- [12] Y. Hong, S. Park, H. Kim, and H. Kim, "Synthetic data generation using building information models," *Autom. Construct.*, vol. 130, Oct. 2021, Art. no. 103871.
- [13] C. Faloutsos and K.-I. Lin, "Fastmap: A fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 1995, pp. 163–174.
- [14] J. Platt, "Fastmap, metricmap, and landmark mds are all nystr," in *Proc. Int. Workshop Artif. Intell. Statist.*, 2005, pp. 261–268.
- [15] H. Ping, J. Stoyanovich, and B. Howe, "DataSynthesizer: Privacy-preserving synthetic datasets," in *Proc. 29th Int. Conf. Sci. Stat. Database Manage.*, Jun. 2017, pp. 1–5.
- [16] N. Patki, R. Wedge, and K. Veeramachaneni, "The synthetic data vault," in *Proc. IEEE Int. Conf. Data Sci. Adv. Anal. (DSAA)*, Montreal, QC, Canada, Oct. 2016, pp. 399–410, doi: 10.1109/dsaa.2016.49.
- [17] C. J. Hazard, C. Fusting, M. Resnick, M. Auerbach, M. Meehan, and V. Korobov, "Natively interpretable machine learning and artificial intelligence: Preliminary results and future directions," 2019, *arXiv:1901.00246*.
- [18] B. Nowok, G. M. Raab, and C. Dibben, "Synthpop: Bespoke creation of synthetic data in R," *J. Stat. Softw.*, vol. 74, no. 11, pp. 1–26, 2016.
- [19] X. Ren, C.-M. Yu, W. Yu, S. Yang, X. Yang, J. A. McCann, and S. Y. Philip, "LoPub: High-dimensional crowdsourced data publication with local differential privacy," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 9, pp. 2151–2166, Sep. 2018.
- [20] T. Wang, X. Yang, X. Ren, W. Yu, and S. Yang, "Locally private high-dimensional crowdsourced data release based on copula functions," *IEEE Trans. Services Comput.*, vol. 15, no. 2, pp. 778–792, Mar. 2022.
- [21] M. Raj Behera, S. Upadhyay, S. Shetty, S. Priyadarshini, P. Patel, and K. F. Lee, "FedSyn: Synthetic data generation using federated learning," 2022, *arXiv:2203.05931*.
- [22] X. Ren, L. Shi, W. Yu, S. Yang, C. Zhao, and Z. Xu, "LDP-IDS: Local differential privacy for infinite data streams," in *Proc. Int. Conf. Manage. Data*, New York, NY, USA, Jun. 2022, pp. 1064–1077.
- [23] K. Suresh, M. S. Cohen, C. J. Hartnick, R. A. Bartholomew, D. J. Lee, and M. G. Crowson, "Generation of synthetic tympanic membrane images: Development, human validation, and clinical implications of synthetic data," *PLOS Digit. Health*, vol. 2, no. 2, Feb. 2023, Art. no. e0000202.
- [24] J. Chen, D. Chun, M. Patel, E. Chiang, and J. James, "The validity of synthetic clinical data: A validation study of a leading synthetic data generator (Synthea) using clinical quality measures," *BMC Med. Informat. Decis. Making*, vol. 19, no. 1, pp. 1–9, Dec. 2019.
- [25] M. Hittmeir, A. Ekelhart, and R. Mayer, "On the utility of synthetic data: An empirical evaluation on machine learning tasks," in *Proc. 14th Int. Conf. Availability, Rel. Secur.*, Aug. 2019, pp. 1–6.
- [26] P.-H. Lu, P.-C. Wang, and C.-M. Yu, "Empirical evaluation on synthetic data generation with generative adversarial network," in *Proc. 9th Int. Conf. Web Intell. Mining Semantics*, Jun. 2019, pp. 1–6.
- [27] A. Goncalves, P. Ray, B. Soper, J. Stevens, L. Coyle, and A. P. Sales, "Generation and evaluation of synthetic patient data," *BMC Med. Res. Methodol.*, vol. 20, no. 1, pp. 1–40, Dec. 2020.
- [28] Z. Azizi, C. Zheng, L. Mosquera, L. Pilote, and K. El Emam, "Can synthetic data be a proxy for real clinical trial data? A validation study," *BMJ Open*, vol. 11, no. 4, Apr. 2021, Art. no. e043497.
- [29] Z. Wang, P. Myles, and A. Tucker, "Generating and evaluating cross-sectional synthetic electronic healthcare data: Preserving data utility and patient privacy," *Comput. Intell.*, vol. 37, no. 2, pp. 819–851, May 2021.

- [30] K. El Emam, L. Mosquera, X. Fang, and A. El-Hussuna, "Utility metrics for evaluating synthetic health data generation methods: Validation study," *JMIR Med. Informat.*, vol. 10, no. 4, Apr. 2022, Art. no. e35734.
- [31] F. K. Dankar, M. K. Ibrahim, and L. Ismail, "A multi-dimensional evaluation of synthetic data generators," *IEEE Access*, vol. 10, pp. 11147–11158, 2022.
- [32] A. Alaa, B. Van Breugel, E. S. Saveliev, and M. van der Schaar, "How faithful is your synthetic data? sample-level metrics for evaluating and auditing generative models," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 290–306.
- [33] T. Cunningham, G. Cormode, and H. Ferhatosmanoglu, "Privacy-preserving synthetic location data in the real world," in *Proc. 17th Int. Symp. Spatial Temporal Databases*, Aug. 2021, pp. 23–33.
- [34] R. Mayer, M. Hittmeir, and A. Ekelhart, "Privacy-preserving anomaly detection using synthetic data," in *Proc. Data Appl. Secur. Privacy, 34th Annu. IFIP WG 11.3 Conf. DBSec*, Regensburg, Germany, Springer, 2020, pp. 195–207.
- [35] D. R. Jeske, B. Samadi, P. J. Lin, L. Ye, S. Cox, R. Xiao, T. Younglove, M. Ly, D. Holt, and R. Rich, "Generation of synthetic data sets for evaluating the accuracy of knowledge discovery systems," in *Proc. 11th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2005, pp. 756–762.
- [36] F. Boutros, M. Huber, P. Siebke, T. Rieber, and N. Damer, "SFace: Privacy-friendly and accurate face recognition using synthetic data," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Oct. 2022, pp. 1–11.
- [37] A. Banerjee, C. J. Hazard, J. Beel, C. Mack, J. Xia, M. Resnick, and W. Goddin, "Surprisal driven k -NN for robust and interpretable nonparametric learning," 2023, *arXiv:2311.10246*.
- [38] S. S. Vempala, *The Random Projection Method*, vol. 65. Providence, RI, USA: Amer. Math. Soc., 2005.
- [39] R. Storn and K. Price, "Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces," *J. Global Optim.*, vol. 11, no. 4, pp. 341–359, 1997. doi: [10.1023/A:1008202821328](https://doi.org/10.1023/A:1008202821328).
- [40] R. S. Olson, W. La Cava, P. Orzechowski, R. J. Urbanowicz, and J. H. Moore, "PMLB: A large benchmark suite for machine learning evaluation and comparison," *BioData Min.*, vol. 10, no. 1, pp. 1–13, 2017.
- [41] M. V. Shcherbakov, A. Brebels, N. L. Shcherbakova, A. P. Tyukov, T. A. Janovsky, and V. A. Kamaev, "A survey of forecast error measures," *World Appl. Sci. J.*, vol. 24, no. 24, pp. 171–176, 2013.
- [42] B. Fuglede and F. Topsøe, "Jensen-Shannon divergence and Hilbert space embedding," in *Proc. Int. Symp. Inf. Theory, ISIT*, Jul. 2004, p. 31.
- [43] I. Csiszár, "I-divergence geometry of probability distributions and minimization problems," in *Proc. The Ann. Probab.*, 1975, pp. 146–158.
- [44] T. Chai and R. R. Draxler, "Root mean square error (RMSE) or mean absolute error (MAE)?—Arguments against avoiding RMSE in the literature," *Geoscientific Model Develop.*, vol. 7, no. 3, pp. 1247–1250, Jun. 2014.
- [45] A. Gelman, B. Goodrich, J. Gabry, and A. Vehtari, "R-squared for Bayesian regression models," *Amer. Statistician*, vol. 73, no. 3, pp. 307–309, Jul. 2019.
- [46] J. Hauke and T. Kossowski, "Comparison of values of pearson's and spearman's correlation coefficients on the same sets of data," *Quaestiones Geographicae*, vol. 30, no. 2, pp. 87–93, Jun. 2011.
- [47] T. Xia, R. Krishna, J. Chen, G. Mathew, X. Shen, and T. Menzies, "Hyperparameter optimization for effort estimation," 2018, *arXiv:1805.00336*.
- [48] H. Tu, Z. Yu, and T. Menzies, "Better data labelling with EMBLEM (and how that impacts defect Prediction)," *IEEE Trans. Softw. Eng.*, vol. 48, no. 1, pp. 278–294, Jan. 2022.
- [49] H. Tu, G. Papadimitriou, M. Kiran, C. Wang, A. Mandal, E. Deelman, and T. Menzies, "Mining workflows for anomalous data transfers," in *Proc. IEEE/ACM 18th Int. Conf. Mining Softw. Repositories (MSR)*, May 2021, pp. 1–12.
- [50] G. Macbeth, E. Razumiejczyk, and R. D. Ledesma, "Cliff's delta calculator: A non-parametric effect size program for two groups of observations," *Universitas Psychologica*, vol. 10, no. 2, pp. 545–555, Jun. 2010.
- [51] M. R. Hess and J. D. Kromrey, "Robust confidence intervals for effect sizes: A comparative study of cohen's d and cliff's delta under non-normality and heterogeneous variances," in *Proc. Annu. Meeting Amer. Educ. Res. Assoc.*, 2004, pp. 1–30.



XIAO LING is currently pursuing the Ph.D. degree in computer science with NC State University. His research interests include automated software testing, machine learning for software engineering, and landscape analysis for software analytics.



TIM MENZIES (Fellow, IEEE) received the Ph.D. degree from UNSW, in 1995. He is currently a Professor in computer science with NC State University, USA, researching SE and AI. For more information visit the link: <http://timm.fyi>.



CHRISTOPHER HAZARD received the Ph.D. degree from NC State University. He is currently the Co-Founder and the CTO of the understandable and privacy enhancing AI company Howso. He has a long career across software, AI, and gaming, including affiliations with Motorola, Hazardous Software, Kiva Systems (now Amazon Robotics), and NATO.



JACK SHU received the bachelor's degree from the University of Washington and the master's degree from the Johns Hopkins Whiting School of Engineering. He is currently the Director of Sales Engineer of the understandable and privacy enhancing AI company Howso.



JACOB BEEL received the B.S. degree in computer science from the University of California at Irvine, Irvine, and the M.S. degree in computer science from Georgia Tech. Since graduating, he deeply concerned about the social and ethical implications of AI. He has been with Howso (formerly Diveplane), where he aims to bring interpretability and attributability to AI and discover novel ways of accomplishing a variety of tasks.

• • •