



A supervised generative optimization approach for tabular data

Fadi Hamad*

University of Pittsburgh
Pittsburgh, Pennsylvania, USA
fah33@pitt.edu

Saheed Obitayo

J.P. Morgan AI Research
New York, USA
saheed.o.obitayo@jpmorgan.com

Shinpei Nakamura-Sakai*

Yale University
New Haven, Connecticut, USA
s.nakamura.sakai@yale.edu

Vamsi K. Potluru

J.P. Morgan AI Research
New York, USA
vamsi.k.potluru@jpmchase.com

ABSTRACT

Synthetic data generation has emerged as a crucial topic for financial institutions, driven by multiple factors, such as privacy protection and data augmentation. Many algorithms have been proposed for synthetic data generation but reaching the consensus on which method we should use for the specific data sets and use cases remains challenging. Moreover, the majority of existing approaches are “unsupervised” in the sense that they do not take into account the downstream task. To address these issues, this work presents a novel synthetic data generation framework. The framework integrates a supervised component tailored to the specific downstream task and employs a meta-learning approach to learn the optimal mixture distribution of existing synthetic distributions.

KEYWORDS

Synthetic data generation, Bayesian Optimisation, Meta-Learning, Fraud Detection

ACM Reference Format:

Fadi Hamad, Shinpei Nakamura-Sakai, Saheed Obitayo, and Vamsi K. Potluru. 2023. A supervised generative optimization approach for tabular data. In *4th ACM International Conference on AI in Finance (ICAIF '23)*, November 27–29, 2023, Brooklyn, NY, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3604237.3626907>

1 INTRODUCTION

Synthetic data generation is vital in various industries, like finance, telecommunication, and healthcare where data-driven decision-making is crucial [11]. It resolves data scarcity and quality concerns by providing synthetic data that preserves statistical properties and relations with the original data. Synthetic data enables testing new ideas without compromising real data, blending multiple sources, and protecting individual privacy [25]. However, using synthetic data may cause performance degradation in modeling [10] where utility degradation depends on the fidelity of the data generation

process and the downstream task. To address this issue and maintain synthetic data quality, the development of a framework to mitigate the degradation is indispensable.

The majority of existing approaches are “unsupervised” in the sense that they do not take into account the downstream task. For instance, the methods discussed in [19] treat the output variable used in the downstream task like other covariates. Their primary focus is creating models that merely “resemble” the original data sets, with variations in how they quantify this resemblance. However, in some use cases, the primary objective is not just achieving similarity to the actual data distribution but optimizing for downstream predictions using the generated data. While research suggests [10] that a closer match to the original data distribution generally leads to better downstream performance, there is an alternative approach worth exploring, incorporating a “supervised” component and optimizing directly on the downstream loss function. By doing so, the generated data can be tailored more effectively to improve the performance of the specific downstream task.

Recent studies show that there are various approaches to model tabular data distribution and sample rows from the distribution[5]: neural network-based[18], machine learning-based[2], and statistical-based[12] generative models. Each of these synthetic data generation methods has its distinct capabilities and features. Many algorithms have been proposed for synthetic data generation but reaching the consensus on which method we should use for the specific data sets and use cases remains challenging. For instance, [26] found that, while table variational autoEncoder (TVAE) generally outperforms Conditional Tabular Generative Adversarial Networks (CTGAN) in most experiments, CTGANs are preferred when prioritizing differential privacy. Furthermore, Gaussian Copula, being a parametric model, proves advantageous when dealing with known marginal distributions. More importantly, the appropriate method depends on many factors such as the distribution of the observed data or the objective of generating the synthetic data.

This work proposes a novel synthetic data generation framework, the Supervised and Composed Generative Optimization Approach for Tabular data (SC-GOAT), to address the aforementioned issues. The framework comprises two key steps. Firstly, we incorporate a supervised component customized for the specific downstream task. To achieve this, we leverage a Bayesian optimization approach to fine-tune the hyperparameters related to the neural networks. Additionally, apart from optimizing the traditional loss of GAN or TVAE through their parameters, we also tune the hyperparameters based on a loss function specific to the downstream task on the

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICAIF '23, November 27–29, 2023, Brooklyn, NY, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0240-2/23/11.
<https://doi.org/10.1145/3604237.3626907>

validation set. Subsequently, we evaluate the model’s performance based on the test error to accurately assess its effectiveness.

In the second step, we adopt a meta-learning approach, leveraging Bayesian optimization, to identify the optimal mixture distribution of existing synthetic data generation methods. To the best of our knowledge, this approach is the first to generate synthetic data based on a mixture of multiple synthetic data generation methods. From each method we learned in the first step, we explore multiple data generation techniques and tune the proportion of data sets sampled. This approach is motivated by the quest to discover the projection of the true underlying data distribution onto the set encompassing various synthesizers. Employing supervised Bayesian optimization, we search for the ideal mixture that optimizes the downstream performance metric. By dynamically combining the strengths of different data generation methods, we aim to enhance the overall synthetic data quality and its suitability for downstream tasks.

Our contributions

- (1) Our approach incorporates supervised components, granting us the flexibility to customize the metric of interest. Whether it’s efficacy, fidelity, or privacy, we can tailor the approach accordingly to boost its performance.
- (2) We introduce a meta-learning framework that leverages various methods to learn the optimal mixture distribution, improving our metric of interest. Additionally, our approach remains robust, even with inefficient synthetic data from certain models.
- (3) Our proposed methodology consistently outperforms existing methods, exhibiting a statistically significant improvement with a p-value of less than 1% in the majority of cases.

2 RELATED WORK

Recent studies have revealed diverse approaches to modeling tabular data distribution and sampling from it [5]. These approaches include Neural network-based methods [18], Machine Learning-based techniques [2], and Statistical-based generative models [12]. Each of these methods for synthetic data generation possesses unique capabilities and features. For the purpose of this paper, our focus centers on exploring Neural network-based and Statistical-based approaches.

The Synthetic Data Vault (SDV) project, utilized for conducting most of the experiments [19], offers two Generative Adversarial Networks (GAN)-based models for data generation from single tables: Conditional Tabular GAN (CTGAN) and CopulaGAN. GANs represent a powerful generative modeling approach employing deep learning methods like convolutional neural networks. Since the original GAN formulation [9], ongoing research has led to the proposal of new optimization strategies and modifications to address GAN limitations. One notable model that builds upon prior successes is CTGAN, which employs mode-specific normalization to capture non-Gaussian and multimodal distributions [26]. It also introduces a conditional generator and training by sampling to tackle challenges posed by highly imbalanced categorical columns and the sparsity of one-hot-encoded vectors, limitations observed in previous GAN architectures. Another neural network approach

from SDV is known as TVAE [26], which is the variational encoders adapted for tabular data.

Beyond neural networks, synthetic data generation can also be achieved by treating each table column as a random variable, modeling a multivariate probability distribution, and sampling from it. SDV presents a synthesizer using this approach called Gaussian Copula [15], which leverages copula functions. These mathematical functions allow for describing the joint distribution of multiple random variables by analyzing the dependencies between their marginal distributions [19]. In the SDV project, univariate marginals are learned using a Gaussian mixture model, while the multivariate copula is learned as a Gaussian copula. The Gaussian Copula approach is valuable for modeling both the covariances between features and their distributions [14].

On the other hand, Bayesian optimization is a powerful and efficient technique used in various fields to optimize complex and costly functions [21, 23]. This methodology is particularly valuable when exploring black-box functions, where the underlying mathematical form is unknown or computationally expensive to evaluate. Additionally, Bayesian optimization is used when the hyper-parameters space isn’t continuous or the loss function isn’t differentiable. At its core, Bayesian optimization employs a probabilistic model, typically a Gaussian Process, to capture the surrogate representation of the objective function. By iteratively selecting the next sampling point based on a trade-off between exploration and exploitation, it intelligently navigates the search space, efficiently narrowing down the region likely to contain the global optimum [8]. This approach has shown remarkable success in tasks like hyperparameter tuning and parameter optimization in machine learning, engineering [8, 23], and other domains.

3 SYNTHETIC DATA GENERATION

Let $M = \{GC, CTGAN, C-GAN, TVAE\}$ be the set of the synthetic data generation methods being utilized. For each method $m \in M$, we have a corresponding synthetic data generation function $S_m(N; \omega_m; \theta_m)$ where N is the number of rows to simulate, ω_m is the set of parameters, and θ_m is the set of hyper-parameters. Note that $\theta_{GC} = \emptyset$ as Gaussian Copula does not use neural networks.

Let D_{real} represent the real data set and D_m denote the synthetic data generated by model $m \in M$. Additionally, we have three data sets: D_{train} , D_{val} , and D_{test} , representing the training, validation, and testing data sets, respectively. All D_* has an outcome vector and covariate matrix which could be represented as duplet $D_* = (X_*, Y_*)$. The downstream loss function is defined as $\mathcal{L}(Y, \hat{Y})$ where \hat{Y} is the outcome predicted by the downstream prediction function $\mu = f(Y \sim X)$, where $f(Y \sim X)$ is the notation for a regression estimator but f can be any machine learning estimator and μ denotes the learned function. Additionally, when μ is learned from the synthetic data generated by $S_m(N; \omega_m; \theta_m)$ we denote it as $\mu_{\omega(\theta_m)} = f(Y_m, X_m)$ where $(X_m, Y_m) = S_m(N; \omega_m; \theta_m)$.

SC-GOAT consists of two steps, supervising (Algorithm 1) and composing (Algorithm 2). In both steps of SC-GOAT, our approach follows standard optimization procedures by optimizing a loss function. However, unlike traditional methods, we adopt a Bayesian approach that constructs a probabilistic model around the involved parameters. Subsequently, we update these parameters based on

the evaluation performance of the loss function. To establish the prior/posterior distribution over the objective function, we employ the Parzen-Tree Estimator. This allows us to effectively locate the parameter space's optimal region, maximizing the expected improvement in the loss function. By employing Bayesian optimization in this manner, we can efficiently fine-tune the synthesizer models and enhance the overall performance in generating data that closely resembles the real data set boosting the downstream performance.

3.1 Supervising Synthesizers

The first step of SC-GOAT involves tuning the hyperparameters using an optimization approach that is supervised by the downstream performance metrics. The optimization formulation is given in (1) and (2) as a bi-level optimization problem. To solve this hyperparameter tuning optimization problem, we employ a Bayesian optimization approach [7]. The flexibility of Bayesian hyperparameter tuning allows for easy switching of the target function to optimize. Moreover, we have the option to incorporate privacy or fidelity regularization in addition to the downstream task. The pseudo-code for hyper-parameter tuning the given model m is presented in Algorithm 1.

The supervising synthesizer optimization problem using bi-level formulation is given by:

$$\begin{aligned} \theta^* = \arg \min_{\theta} \quad & \mathcal{L}(Y_{val}, \hat{Y}_{val}) \\ \text{s.t.} \quad & \hat{Y}_{val} := \hat{\mu}_{\omega^*}(\theta)(X_{val}) \end{aligned} \quad (1)$$

$$\text{s.t.} \quad \omega^*(\theta) = \arg \min_{\omega} \mathcal{F}(\omega, \theta, D_{train}) \quad (2)$$

where the outer optimization problem (1) is minimizing the loss function on the validation set D_{val} and the inner optimization problem (2) is minimizing the loss function denoted by \mathcal{F} on the training set D_{train} for synthesizer model S . Note that these functions, \mathcal{L} and \mathcal{F} , are not necessarily the same and may be measured on different models. For instance, \mathcal{F} always refers to the loss function used during the training of the synthesizer model S , whereas \mathcal{L} refers to the model's performance on the validation set, possibly employing a different evaluation metric. Alternatively, \mathcal{L} could also refer to the loss function for the downstream task performed by model f , as it is the case in our approach.

3.2 Composing Synthesizers

The second step of SC-GOAT is the composing process. Here, our objective is to utilize a meta-learning approach to determine the mixture distribution among the synthesizers in M . We refer to it as a meta-learning approach because we learn the final model from the models obtained in the previous step. For each synthesizer $m \in M$, we define $\alpha_m \in [0, 1]$ as the proportion of the total observations sampled from S_m . The final synthetic data comprises $[\alpha_m N]$ observations for each m , where $[\cdot]$ denotes the closest integer function. The formulation of this meta-learning approach using an optimization framework is given in (3) while the pseudo-code for this step is presented in Algorithm 2. Note that the θ^m we use could be the default parameters of each $m \in M$ or the tuned parameter obtained in the Supervising step 1.

Algorithm 1: Supervising Step - Generative Optimization Approach for Tabular data (S-GOAT)

Input: $D_{real} = (X_{real}, Y_{real}), N$
 Create partition $\{D_{train}, D_{val}, D_{test}\} = \{(X_{train}, Y_{train}), (X_{val}, Y_{val}), (X_{test}, Y_{test})\}$
 Initialize $\hat{\theta}_m^0$
for $k = 1, \dots, K$ **do**
 Fit S_m using D_{train}
 Generate $D_m^k = (X_m^k, Y_m^k)$ where
 $(X_m^k, Y_m^k) = S(N; \omega_m; \theta_m^k)$
 Train $\hat{\mu}_{\theta_m^k}^k = f(Y_m^k \sim X_m^k)$
 Compute $\hat{Y}_{val}^k = \hat{\mu}^k(X_{val})$
 Compute $l^k = \mathcal{L}(\hat{Y}_{val}^k, Y_{val})$
 Suggest $\hat{\theta}_m^{k+1}$ using Bayesian Optimization based on
 $\{\theta_m^0, \dots, \theta_m^k\}$ and $\{l^0, \dots, l^k\}$
return $\theta_m^{k^*}$ where $k^* = \arg \min_k l^k$

The meta-learning optimization formulation is given by:

$$\begin{aligned} \alpha^* = \arg \min_{\alpha \in \mathcal{R}^{|M|}} \quad & \mathcal{L}(Y_{val}, \hat{Y}_{val}) \\ \text{s.t.} \quad & (X_{syn}^m, Y_{syn}^m) = S_m([\alpha_m N], \theta_m) \quad \forall m \in M \\ & X_{syn} = [(X_{syn}^m)^T : m \in M]^T \\ & Y_{syn} = [Y_{syn}^m : m \in M] \\ & \hat{\mu} = f(Y_{syn}, X_{syn}) \\ & \hat{Y}_{val} := \hat{\mu}(X_{val}) \end{aligned} \quad (3)$$

where \mathcal{L} refers to the loss function of the downstream task on the validation set D_{val} , which we use to evaluate the quality of the α 's generated by the Bayesian optimization at each iteration. This involves evaluating the downstream task performance on the combined synthetic data generated using different methods as highlighted in Algorithm 2.

3.3 Evaluation

When generating synthetic data, one common concern is assessing the quality of the generated data. To evaluate synthetic generation models for tabular data, various benchmarking approaches are available, allowing flexibility in adapting the loss function to suit the specific objectives of synthetic data generation.

To evaluate the accuracy of preserving individual attributes and attribute pairs in synthetic data, the KS-Test and CS-Test are valuable tools. The KS-Test compares continuous column distributions using the empirical CDF [6], while the CS-Test compares discrete column distributions using the Chi-Squared test [19]. Additionally, fidelity can be assessed by building a machine learning classifier to differentiate between real and synthetic data [19].

While evaluating the distribution of synthetic and real data is crucial, we must also address privacy protection at an individual level. Like many machine learning models, synthetic generative approaches are susceptible to privacy attacks [24], including Membership Inference Attacks (MIA) [22], Reconstruction attacks [17], and

Algorithm 2: Composing Step - Generative Optimization Approach for Tabular data (C-GOAT)

Input: $D_{real}, \theta^m \forall m \in M$;
 Create partition $\{D_{train}, D_{val}, D_{test}\} = \{(X_{train}, Y_{train}), (X_{val}, Y_{val}), (X_{test}, Y_{test})\}$
 Initialize $\alpha = \{\alpha_m\}_{m \in M}$
for $k = 1, \dots, K$ **do**
 Sample $D_m^k = S_m([\alpha_m^k N; \theta^m]) \quad \forall m \in M$
 Create $D_{syn}^k = \begin{pmatrix} D_{GC}^k \\ \vdots \\ D_{TVAE}^k \end{pmatrix} = \begin{pmatrix} X_{GC}^k & Y_{GC}^k \\ \vdots & \vdots \\ X_{TVAE}^k & Y_{TVAE}^k \end{pmatrix}$
 Train $\hat{\mu}^k = f(Y_{syn}^k \sim X_m^k)$
 Compute $\hat{Y}_{val}^k = \hat{\mu}^k(X_{val})$
 Compute $l^k = \mathcal{L}(\hat{Y}_{val}^k, Y_{val})$
 Suggest α^{k+1} using Bayesian optimization approach based on $\{\alpha^0, \dots, \alpha^k\}$ and $\{l^0, \dots, l^k\}$
return $D_{syn}^{k^*}$ where $k^* = \arg \min_k l^k$

Property inference attacks [13]. Addressing these privacy vulnerabilities is crucial to preserving the utility and integrity of synthetic data.

In this paper, we aim to evaluate our synthetic generative models through the lens of the downstream classification model's accuracy, which serves as a robust metric to assess the models' overall performance.

4 DATA

Adult. This data set is a sample from the US Census Bureau Database that contains the census result of the year 1994¹. This data set includes 48,842 records and 14 attributes. Each record contains the following features such as age, gender, education, relationship, occupation, race, and native country of a representative individual in the census record. These attributes are a mixture of numerical, ordinal, and categorical data types. The data set has a binary target label which indicates whether the income of an individual is less or greater than fifty thousand dollars. Therefore, the data set has a classification task which is to predict if a person makes over 50K a year based on the census attributes

Credit Card Fraud. To showcase the usefulness of synthetic tabular data, we use the credit card fraud data set². This data set contains transactions collected in the span of two days made by credit cards by European cardholders for the month of September 2013. From an analysis, it can be observed that the data set is highly imbalanced containing 492 frauds out of 28,4807 total transactions. The positive class of fraud accounts for 0.172% of all transactions. The credit card fraud data set contains only numerical input variables with 31 features. With respect to confidentiality and privacy, 28 of the features - V1 to V28 are principal components obtained by the means of PCA. 'Time', 'Amount', and 'Class' are the only

features not to be transformed with PCA. The feature 'Time' contains the seconds elapsed between each transaction and the first transaction in the data set. The feature 'Class' is the target variable which takes the value of 0 for cases of no fraud and 1 for cases of fraud. The feature 'Amount' is the transaction amount. Given the class imbalance ratio of the credit fraud data set, we processed the data set by oversampling the minority class with random under-sampling of the majority class, leading to a more balanced data set. This involved duplicating examples in the minority class in order to reach an equal balance between the minority and majority class. This process will reduce the number of data points available. This technique is called Synthetic Minority Oversampling Technique (SMOTE)[3]. Applying this technique will lead us to have two separate data sets, an original imbalanced credit data set and a new balanced credit data set.

This data set will be helpful in the context of fraud detection for machine learning utility. We can answer whether synthetic data generation can help with downstream tasks in the fraud management process. The utility of models trained on fraud data sets allows us to measure the effectiveness of detecting and predicting potential fraudulent operations. This provides guidance to fraud practitioners interested in utility using synthetic data to train fraud detection models.

5 EXPERIMENTAL RESULTS

We evaluate our approach on three diverse data sets discussed in the previous section: the adult data set¹, the balanced credit card data set³, and the imbalanced credit card data set². We chose these data sets as they are widely utilized in previous works for evaluating tabular synthetic data generation methods. The adult data set contains both numerical and categorical variables allowing us to showcase the applicability of our approach in generating different types of data.

Furthermore, by selecting both balanced and imbalanced data sets, we can demonstrate the robustness of our approach across various data distributions. The data set's descriptions are summarized in Table 1. For the adult data set, we utilized all available records, totaling 48.842K. However, for the credit card data set, we sampled 50K records from the available 28.407K records similar to [27]. Through this evaluation, we gain valuable insights into the generalizability and performance of our approach, enhancing its credibility as a powerful tool for generating high-quality synthetic data across a diverse range of scenarios.

Our method is implemented as an open-source Python package that will be available on GitHub. The implementation utilizes four generative methods, namely Gaussian Copula, CTGAN, Copula GAN, and TVAE, available from the SDV [20] python package. For the downstream task evaluation, we utilize the XGBoost classifier python package [4], and for the Bayesian optimization, we use the hyperopt python package [1].

All experiments were conducted using Python 3.10. The code repository provides comprehensive instructions for replicating the experiments and includes detailed result tables. For further insights

¹Data available on the UCI platform at <https://archive.ics.uci.edu/dataset/2/adult>

²Data available on the Kaggle platform at <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>

³The credit card data set with oversampling the minority class as mentioned in Section 4.

⁴This is the original adult data set with target encoder transformation for the categorical features.

Table 1: Description of data sets

Data set	Label	Observation	Continuous	Binary	Multi-class	Label = 0	Label = 1
Adult ¹	'income'	48,842	6	2	7	76.07%	23.93%
Credit Balanced ³	'Class'	50,000	30	1	0	66.70%	33.3%
Credit Imbalanced ²	'Class'	50,000	30	1	0	99.82%	0.18%

Table 2: Description of synthetic data sets generated using each model.

Data set	Method	Untuned		Tuned	
		Label = 0	Label = 1	Label = 0	Label = 1
Adult	Gaussian Copula [20]	80.40%	19.60%	82.14%	17.9%
	CTGAN [20]	83.66%	16.34%	77.15%	22.9%
	CopulaGAN [20]	74.96%	25.04%	74.96%	25.0%
	TVAE [20]	76.44%	23.56%	77.63%	22.4%
	SC-GOAT	78.47%	21.53%	77.41%	22.59%
Adult Transformed ⁴	Gaussian Copula	94.29%	5.71%	80.36%	19.6%
	CTGAN	68.63%	31.37%	77.45%	22.6%
	CopulaGAN	74.85%	25.15%	74.85%	25.2%
	TVAE	78.70%	21.30%	76.32%	23.7%
	SC-GOAT	70.19%	29.81%	72.93%	27.07%
Credit Balanced	Gaussian Copula	70.07%	29.93%	63.66%	36.3%
	CTGAN	68.33%	31.67%	63.73%	36.3%
	CopulaGAN	91.97%	8.03%	91.97%	8.0%
	TVAE	63.90%	36.10%	63.90%	36.1%
	SC-GOAT	63.90%	36.10%	63.90%	36.1%
Credit Imbalanced	Gaussian Copula	0.00%	100.00%	0.00%	100.0%
	CTGAN	100.00%	0.00%	100.00%	0.0%
	CopulaGAN	99.82%	0.18%	99.82%	0.2%
	TVAE	100.00%	0.00%	100.00%	0.0%
	SC-GOAT	99.65%	0.35%	99.65%	0.35%

into the generated synthetic data sets, a summary is provided in Table 2 for one experiment.

Each experiment was repeated 10 times and for these experiments, 70% of the real data was used for training, 20% for validation, and the remaining 10% for testing. For the untuned setup, we set $K = 350$ in Algorithm 1, while for the tuned setup, we used $K = 150$ in Algorithm 2. For Algorithm 2, we generate the alphas using the uniform distribution $Uniform(0, 1)$ and then we scale those alphas to add up to 1 as follows:

$$\alpha_m = \frac{\alpha_m}{\sum_{j \in M} \alpha_j} \quad \forall m \in M$$

For the first iteration, instead of randomly generating the alphas, we use a warm start. The warm start also addresses a weakness in the Bayesian optimization using the Tree-Structured Parzen-Estimator, as in the majority of cases, it fails to converge to the optimal solution if it lies on one of the corner points. By "corner points" here, we mean that the optimal solution only considers the best method and neglects the others, which can be mathematically defined as $\alpha = 1$ for the best model and $\alpha = 0$ for the other methods. This decision is based on the evaluation metric for the validation data set, which, in our implementation, is the AUC score. Therefore, our approach considers five initial starting points. Four of them represent the corner points that correspond to each model $m \in M$. For example, the initial starting point that only represents the Gaussian Copula model will be given by $[1.0, 0.0, 0.0, 0.0]$. The last initial starting point is initialized based on the AUC validation of each individual model:

$$\alpha_m = \frac{auc_{val}^m - auc_{val}^*}{\sum_{j \in M} (auc_{val}^j - auc_{val}^*)} \quad \forall m \in M$$

where $auc_{val}^* = \min_{i \in M} auc_{val}^i$.

The decision for which initial point to pick depends on the point that gives the best validation AUC score for the first iteration. This initialization scheme ensures that in case our optimal solution lies on one of the corner points, we will converge to it. Additionally, if the optimal solution represents a mixture of each individual model, we will also be able to capture it, as the algorithm will generate alpha vectors that continuously improve the loss function.

However, to prevent the chance of overfitting, we implemented early stopping by adding a condition to stop the algorithm once the AUC score on the validation set doesn't improve for the last k iterations. For algorithm 1, we used $k = 10$, and for algorithm 2 we used $k = 15$.

Our results include fitting all the individual models from the SDV package [20] without any hyper-parameter tuning as well as tuning these models as mentioned in Algorithm 1. For CTAB-GAN+, the details of the experiment are mentioned in Subsection 5.3. We also reported the results of our method using both tuned and untuned setups where for the untuned setup we only use Algorithm 2 with untuned models from the SDV package [20] while for the tuned setup we use both Algorithm 1 and Algorithm 2.

Table 3: Average test AUC for the XGBoost baseline model fitted only on real data for each data set for 10 experiments.

Method	Adult	Credit Balanced	Credit Imbalanced
XGBoost [4]	90.51%	99.99%	96.02%

5.1 Performance Evaluation

To optimize the loss function, we aim to maximize the AUC score for the downstream classification task. This is achieved by training an XGBoost classifier [4] on the training data set and subsequently evaluating its performance on a separate validation data set. To ensure a fair comparison between the different methods, the XGBoost classifier is utilized with its default parameters. By focusing on the maximization of the AUC score in the downstream task, we can accurately evaluate the synthetic data's quality. Our primary objective is to generate data that closely resembles the real data. Therefore, by emphasizing the AUC score, we ensure that the synthetic data is as representative as possible, enabling it to capture essential characteristics and patterns present in the real data set.

5.2 Baseline model

To comprehensively evaluate the effectiveness and improvements of our method, as well as the quality of the generated data compared to the original real data set, we fitted XGBoost on the original data and

assessed the model’s performance in terms of AUC. By comparing our results against the baseline XGBoost model fitted on real data, we gain valuable insights into the efficiency of our approach and the similarity between the generated synthetic data and the real data.

5.3 CTAB-GAN+

Given that the primary criteria for evaluating our approach rely on downstream losses, we conduct a thorough comparison against CTAB-GAN+ [27]. CTAB-GAN+ stands out as a novel conditional tabular GAN, surpassing existing state-of-the-art approaches by incorporating downstream losses into conditional GANs. This innovation results in higher utility synthetic data that proves beneficial in both classification and regression domains. The model introduces several other major improvements over existing methods [27]. As we compare our approach to this state-of-the-art alternative, we aim to showcase the strengths and competitive advantages of our synthetic data generation technique in practical scenarios.

We focus on the default version of CTAB-GAN+ without any fine-tuning. The decision is driven by time constraints, as tuning the model requires a significant amount of time compared to the models present in SDV. By considering the default CTAB-GAN+, we can still gain valuable insights and effectively evaluate the relative strengths of our method without the need for extensive fine-tuning efforts. Table 5 summarizes the results.

To ensure a fair comparison, we did not incorporate CTAB-GAN+ in the composition process. Our goal was to demonstrate that we could surpass the current state-of-the-art through meta-learning of preexisting methods.

5.4 Adult data set

We evaluated the efficiency of our method on the adult data set¹, focusing on its performance with categorical data. For encoding the categorical features, we employed two distinct approaches. The first approach relied on the implicit handling by the SDV python package [20] during the fitting of synthesizers on the real data. This implicit handling is implemented using a label encoder. The second approach involved implementing a target encoder, which outperforms traditional encoding schemes, especially for categorical features with high cardinality in the categorical columns [16]. We referred to the transformed data, utilizing the target encoder, as the ‘Adult Transformed’ data set.

Subsequently, we compared the performance of our approach against each individual synthesizer model $m \in M$ using two setups: untuned models and tuned models, as described in section 3.1. For the ‘income’ column, we mapped rows with values ‘ $\leq 50k$ ’ to 0 and rows with values ‘ $> 50k$ ’ to 1. The results of the comparison are summarized in Table 5.

5.5 Balanced credit card data set

We further demonstrate the effectiveness of our approach on the credit card data set², which exclusively comprises numerical features. Initially, the credit card data set exhibited an imbalanced distribution, as indicated in Table 1. However, as detailed in Section 4, we preprocessed the data set using a random undersampling

of the majority class. This resulted in a more balanced data set, facilitating a fairer evaluation.

Similarly to our approach for the adult data set, we conduct a thorough comparison of our method’s performance against that of each individual synthesizer model $m \in M$, employing two setups: untuned models and tuned models, as described in Section 3.1. The results of this comparison are summarized in Table 5.

5.6 Imbalanced credit card data set

Considering real-life scenarios often involve highly imbalanced data, we evaluated the performance of our approach on an imbalanced credit card data set. Initially, we ran our approach using four synthesizer methods in both ‘tuned’ and ‘untuned’ setups. However, we observed weaknesses in handling highly imbalanced data sets, as these methods struggled to generate data from both classes effectively. For instance, TVAE only generated data from the majority class, while CopulaGAN unexpectedly generated data solely from the minority class. To address this issue, we implemented conditional sampling, with only Gaussian Copula successfully generating data resembling the original data set. The results of this comparison are summarized in Table 5.

5.7 Results Analysis

To simplify the comparison between our approach and the other different methods, and to demonstrate the applicability of our method in various scenarios, we present the results averaged across all the experiments for all the methods on the test data in Figure 1. Here we only show the plot comparison for the untuned setup since we didn’t tune the CTAB-GAN+.

It can be observed that our method outperforms all the other approaches in most cases. Only for the imbalanced credit card data set, CTAB-GAN+ performs better than our approach. This is related to the way CTAB-GAN+ handles imbalanced data, as they implement a training-by-sampling strategy. The idea behind their approach is to resample classes, giving higher chances to minority classes to train the model [27]. This approach is somehow similar to making the data balanced as we did in Subsection 4 to handle the credit card data, and this similarity can be observed clearly on the balanced data set, as both approaches perform the same (refer to Figure 1). This means that comparing our approach with CTAB-GAN+ on imbalanced data isn’t fair due to the extra step they implement. Even without implementing any additional upsampling, our approach performed very close to CTAB-GAN+ on the imbalanced data (refer to Table 5)³.

A major part of our work is providing a metric to understand the most suitable synthesizer among multiple synthesizers for a specific objective. To demonstrate that, we present in Table 4 the values of alphas for each model from one of the experiments. We can clearly see from this table that the alpha weights are linked to the performance of each individual model, and as expected, the model with better performance has a higher α weight compared to the other models. In the balanced credit card data, the corner solution emerged as the winner, with all weights assigned to TVAE

³In Figure 1, we illustrate CTAB-GAN+ for the imbalanced data set using a partial filled color to emphasize the distinction between their approach and ours in handling imbalanced data, as described in Section 5.7.

Table 4: Contribution of each individual model (α) for the final synthetic data generated.

Data set	Method	Weights (untuned setup)	Weights (tuned setup)
Adult	Gaussian Copula	20.90%	17.79%
	CTGAN	29.50%	41.83%
	CopulaGAN	19.07%	14.26%
	TVAE	30.53%	26.12%
Adult Transformed	Gaussian Copula	26.69%	19.44%
	CTGAN	32.98%	42.72%
	CopulaGAN	8.22%	11.51%
	TVAE	32.11%	26.33%
Credit Balanced	Gaussian Copula	0%	0%
	CTGAN	0%	0%
	CopulaGAN	0%	0%
	TVAE	100%	100%
Credit Imbalanced	Gaussian Copula	99.33%	99.33%
	CTGAN	0.22%	0.22%
	CopulaGAN	0.22%	0.22%
	TVAE	0.22%	0.22%

owing to its superior validation AUC before early stopping. For the imbalanced credit card data, the warm-started weight achieved the highest downstream AUC in most cases. Lastly, in the adult dataset, a mixture of the four methods was used, and Algorithm 2 learned the optimal weight based on the downstream validation AUC.

Comparing against the baseline model fitted only on real data as explained in Subsection 5.2, our aim is to show that the generated synthetic data have a similar downstream performance. Comparing the values from Tables 3 and 5 shows that the AUC test scores for the XGBoost fitted only on real data are very close to the AUC test scores for the XGBoost fitted only on synthetic data. This validates our point that the distribution of the data generated by our approach can perform a high-quality downstream task.

We could observe that tuning the neural network-related hyperparameters through Algorithm 1 did not lead to a significant performance boost. This finding raises two possibilities: we may need to expand the hyperparameter grid and explore the space more extensively, or the three methods we tuned inherently exhibit robustness to hyperparameters. We recognize this as an open discussion for further investigation. If the latter holds true, practitioners could potentially skip the costly hyperparameter tuning step, and use out untuned set-up leading to a more efficient synthetic data creation process.

6 CONCLUSION AND FUTURE WORK

Our approach has shown great promise, consistently outperforming the majority of previous methods in terms of the downstream metric. For future work, one potential direction is to evaluate the performance of SC-GOAT concerning privacy or fidelity aspects. Moreover, this approach could be further explored for data augmentation purposes, aiming to surpass the downstream metric achieved with real data. Such investigations could provide valuable insights and advancements in the field.

Disclaimer This paper was prepared for informational purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates (“J.P. Morgan”), and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

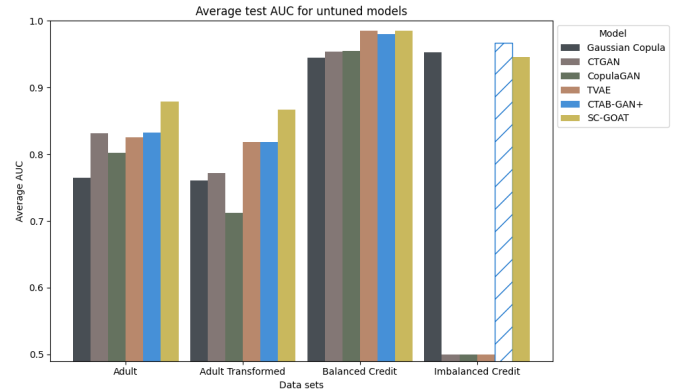


Figure 1: Average downstream test AUC score for 10 experiments using XGBoost fitted on the generated data by each model in the untuned setup⁵.

Table 5: Average, standard deviation, and one-sided paired t-test for the downstream test AUC score, using XGBoost fitted on the generated data by each method, on 10 experiments.

Data set	Method	Untuned				Tuned			
		average	std	test statistic	p-value	average	std	test statistic	p-value
Adult	Gaussian Copula	76.43%	0.04	10.07	0	76.43%	0.04	10.24	0
	CTGAN	83.12%	0.02	9.39	0	84.28%	0.01	7.84	0
	CopulaGAN	80.18%	0.03	7.59	0	79.58%	0.03	9.12	0
	TVAE	82.59%	0.01	14.25	0	82.92%	0.01	10.97	0
	CTAB-GAN+	83.25%	0.03	4.44	0.0002	-	-	-	-
	SC-GOAT	87.87%	0.00	-	-	88.12%	0.01	-	-
Adult Transformed	Gaussian Copula	76.06%	0.03	9.60	0	76.06%	0.03	9.95	0
	CTGAN	77.21%	0.04	7.20	0	78.62%	0.05	5.23	0
	CopulaGAN	71.24%	0.06	8.37	0	73.02%	0.05	9.56	0
	TVAE	81.86%	0.02	6.23	0	81.07%	0.03	5.49	0
	CTAB-GAN+	81.81%	0.03	4.76	0.0001	-	-	-	-
	SC-GOAT	86.67%	0.01	-	-	86.97%	0.01	-	-
Credit Balanced	Gaussian Copula	94.45%	0.01	14.10	0	94.45%	0.01	14.31	0
	CTGAN	95.34%	0.01	16.21	0	95.93%	0.01	13.15	0
	CopulaGAN	95.50%	0.01	14.18	0	96.41%	0.01	7.80	0
	TVAE	98.52%	0.00	0.00	0.5	98.48%	0.00	0.00	0.5
	CTAB-GAN+	98.04%	0.00	2.74	0.0068	-	-	-	-
	SC-GOAT	98.52%	0.00	-	-	98.48%	0.00	-	-
Credit Imbalanced	Gaussian Copula	95.32%	0.07	-0.23	0.5886	95.32%	0.07	0.47	0.3215
	CTGAN	50.00%	0.00	19.17	0	50.00%	0.00	28.47	0
	CopulaGAN	50.00%	0.00	19.17	0	50.00%	0.00	28.47	0
	TVAE	50.00%	0.00	19.17	0	50.00%	0.00	28.47	0
	CTAB-GAN+	96.70%	0.05	-0.74	0.7667	-	-	-	-
	SC-GOAT	94.60%	0.07	-	-	96.59%	0.05	-	-

REFERENCES

- [1] James Bergstra, Daniel Yamins, and David Cox. 2013. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning*. PMLR, 115–123.
- [2] Gregory Caiola and Jerome P Reiter. 2010. Random Forests for Generating Partially Synthetic, Categorical Data. *Trans. Data Priv.* 3, 1 (2010), 27–42.
- [3] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [4] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [5] Josh Eno and Craig W Thompson. 2008. Generating synthetic data to match data mining patterns. *IEEE Internet Computing* 12, 3 (2008), 78–82.
- [6] Giovanni Fasano and Alberto Franceschini. 1987. A multidimensional version of the Kolmogorov–Smirnov test. *Monthly Notices of the Royal Astronomical Society* 225, 1 (1987), 155–170.
- [7] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. 2018. Bilevel programming for hyperparameter optimization and meta-learning. In *International conference on machine learning*. PMLR, 1568–1577.
- [8] Peter I Frazier. 2018. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811* (2018).
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).
- [10] Markus Hittmeir, Andreas Ekelhart, and Rudolf Mayer. 2019. On the utility of synthetic data: An empirical evaluation on machine learning tasks. In *Proceedings of the 14th International Conference on Availability, Reliability and Security*. 1–6.
- [11] James Jordon, Lukasz Szpruch, Florimond Houssiau, Mirko Bottarelli, Giovanni Cherubin, Carsten Maple, Samuel N Cohen, and Adrian Weller. 2022. Synthetic Data—what, why and how? *arXiv preprint arXiv:2205.03257* (2022).
- [12] Zheng Li, Yue Zhao, and Jialin Fu. 2020. Sync: A copula based framework for generating synthetic data from aggregated sources. In *2020 International Conference on Data Mining Workshops (ICDMW)*. IEEE, 571–578.
- [13] Zinan Lin, Shuaiqi Wang, Vyas Sekar, and Giulia Fanti. 2023. Summary Statistic Privacy in Data Sharing. *arXiv preprint arXiv:2303.02014* (2023).
- [14] Majlinda Llugiqi and Rudolf Mayer. 2022. An Empirical Analysis of Synthetic-Data-Based Anomaly Detection. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 306–327.
- [15] Guido Masarotto and Cristiano Varin. 2012. Gaussian copula marginal regression. (2012).
- [16] Daniele Micci-Barreca. 2001. A preprocessing scheme for high-cardinality categorical attributes in classification and prediction problems. *ACM SIGKDD Explorations Newsletter* 3, 1 (2001), 27–32.
- [17] Arvind Narayanan and Vitaly Shmatikov. 2006. How to break anonymity of the netflix prize dataset. *arXiv preprint cs/0610105* (2006).
- [18] Noseong Park, Mahmoud Mohammadi, Kshitij Gorde, Sushil Jajodia, Hongkyu Park, and Youngmin Kim. 2018. Data synthesis based on generative adversarial networks. *arXiv preprint arXiv:1806.03384* (2018).
- [19] Neha Patki, Roy Wedge, and Kalyan Veeramachaneni. 2016. The synthetic data vault. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE, 399–410.
- [20] Neha Patki, Roy Wedge, and Kalyan Veeramachaneni. 2016. The Synthetic data vault. In *IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. 399–410. <https://doi.org/10.1109/DSAA.2016.49>
- [21] Martin Pelikan, David E Goldberg, Erick Cantú-Paz, et al. 1999. BOA: The Bayesian optimization algorithm. In *Proceedings of the genetic and evolutionary computation conference GECCO-99*, Vol. 1. Citeseer.
- [22] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*. IEEE, 3–18.
- [23] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* 25 (2012).

- [24] Hui Sun, Tianqing Zhu, Zhiqiu Zhang, Dawei Jin, Ping Xiong, and Wanlei Zhou. 2021. Adversarial attacks against deep generative models on data: a survey. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [25] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing* 10, 3152676 (2017), 10–5555.
- [26] Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2019. Modeling tabular data using conditional gan. *Advances in neural information processing systems* 32 (2019).
- [27] Zilong Zhao, Aditya Kumar, Robert Birke, and Lydia Y Chen. 2022. Ctab-gan+: Enhancing tabular data synthesis. *arXiv preprint arXiv:2204.00401* (2022).