

Research Article

SampleLLM: Optimizing Tabular Data Synthesis in Recommendations

Jingtong Gao¹, Zhaocheng Du², Xiaopeng Li¹, Yichao Wang², Xiangyang Li², Huifeng Guo², Ruiming Tang², Xiangyu Zhao¹

1. City University of Hong Kong, Hong Kong; 2. Huawei Noah's Ark Lab, China

Tabular data synthesis is crucial in machine learning, yet existing general methods—primarily based on statistical or deep learning models—are highly data-dependent and often fall short in recommender systems. This limitation arises from their difficulty in capturing complex distributions and understanding complicated feature relations from sparse and limited data, along with their inability to grasp semantic feature relations. Recently, Large Language Models (LLMs) have shown potential in generating synthetic data through few-shot learning and semantic understanding. However, they often suffer from inconsistent distribution and lack of diversity due to their inherent distribution disparity with the target dataset. To address these challenges and enhance tabular data synthesis for recommendation tasks, we propose a novel two-stage framework named SampleLLM to improve the quality of LLM-based tabular data synthesis for recommendations by ensuring better distribution alignment. In the first stage, SampleLLM employs LLMs with Chain-of-Thought prompts and diverse exemplars to generate data that closely aligns with the target dataset distribution, even when input samples are limited. The second stage uses an advanced feature attribution-based importance sampling method to refine feature relationships within the synthetic data, reducing any distribution biases introduced by the LLM. Experimental results on three recommendation datasets, two general datasets, and online deployment illustrate that SampleLLM significantly surpasses existing methods for recommendation tasks and holds promise for a broader range of tabular data scenarios.

1. Introduction

Tabular data is integral to a wide array of machine learning applications across sectors like e-commerce and healthcare, underscoring its foundational importance^{[1][2][3][4][5]}. This widespread reliance amplifies the urgent demand for high-quality synthetic tabular data, particularly in recommender systems, where data sparsity poses significant challenges^{[6][7][4]}. Insufficient quality and volume of datasets critically impair the performance and efficiency of machine learning models^[8].

Traditional tabular data synthesis methods^{[9][10][11]}, focusing on a wide range of general tasks, primarily rely on statistical and deep learning models, which are heavily data-dependent and effective with abundant data availability^[12]. However, when applied to recommendation tasks, these methods not only struggle in scenarios characterized by inherent data sparsity and scarcity but also fail to capture the semantic relationships between features, which is proven to be more and more crucial in recommendation modeling^{[13][14]}. Consequently, there is a pressing need for innovative

approaches capable of generating high-quality synthetic tabular data from minimal input while understanding semantic feature relations for recommendation tasks.

The advent of Large Language Models (LLMs) has marked a shift in possibilities^[15], offering new capabilities through few-shot learning and deep semantic understanding^{[16][17][18]}. Despite their promise in synthetic data generation, LLMs frequently encounter challenges with inconsistent distributions and inaccurately modeled feature relationships^{[19][20]}, both of which are critical for effective recommendation modeling^{[21][22][23]}. These challenges arise from an intrinsic distribution mismatch between LLMs' inherent knowledge and target datasets. Moreover, when using limited exemplars, a simplistic random selection for few-shot learning might overlook important regions of the original dataset, leading to reduced output diversity, as demonstrated in Figure 1.

To address these challenges, we propose SampleLLM, a novel two-stage framework designed to enhance the quality of LLM-based tabular data synthesis in recommendation tasks. The first stage utilizes LLMs with Chain-of-Thought prompts^[24] and a curated selection of diverse exemplars to improve semantic understanding and feature relation modeling^[25] in generating synthetic data with limited input samples. The second stage further integrates a feature attribution-based importance sampling technique to refine the synthetic data, reducing any distribution biases introduced by LLMs by employing a semi-independence assumption for feature interactions that could streamline computation while preserving the essential characteristics and relationships within the dataset.




Figure 1. Visualization of LLM-generated and original tabular samples on the HELOC dataset reveals that the synthetic data lacks distribution alignment and diversity, clustering around a few centers within the original data distribution.

In summary, our contribution could be summarized as follows:

- To the best of our knowledge, this is the first method to consider distribution alignment for LLM-based tabular data synthesis in recommendations.

- We propose a comprehensive framework SampleLLM that combines a feature attribution-based importance sampling strategy for distribution alignment with advanced few-shot LLM generation techniques. This integration significantly enhances the utility and distribution consistency of the generated tabular data.
- Experiments conducted on three widely used recommendation datasets, two general datasets, and through online deployment reveal that the synthetic tabular data generated by SampleLLM not only surpasses existing baselines in recommendation tasks but also shows promise for broader application in general tasks.

2. Framework

In this section, we first describe the problem formulation of the tabular data synthesis, and then provide an overview of SampleLLM and detail its key components.

2.1. Problem Formulation

The task of tabular data synthesis involves generating synthetic tabular data that closely resemble those in a given dataset. The synthetic data is used for various purposes, including augmenting training datasets and conducting experiments without exposing sensitive information. Formally, let $\mathcal{D}_o = \{x_1, x_2, \dots, x_N\}$ represent the original dataset with N samples, where each x_i is a sample with K attributes and one label y_i , i.e., $x_i = \{x_i^1, x_i^2, \dots, x_i^K, y_i\}$. The objective is to develop a generative model G that produces a synthetic dataset with T samples $\mathcal{D}_s = \{x'_1, x'_2, \dots, x'_T\}$ such that \mathcal{D}_s exhibits similar statistical properties and distributional characteristics as \mathcal{D}_o and could be applied to various tasks with similar performance. These requirements can be formalized as follows:

- **Distributional Similarity:** The synthetic dataset \mathcal{D}_s should follow the similar distribution as the original dataset \mathcal{D}_o . This can be expressed as:

$$P_{\mathcal{D}_s}(X) \approx P_{\mathcal{D}_o}(X), \quad (1)$$

where $P_{\mathcal{D}_s}(X)$ and $P_{\mathcal{D}_o}(X)$ denote the probability distributions of the synthetic and original datasets, respectively.

- **Utility:** The synthetic dataset \mathcal{D}_s should be useful for the same downstream tasks as the original dataset \mathcal{D}_o . If \mathcal{D}_s is used to train a machine learning model \mathcal{M} , the model's performance on downstream tasks should closely match that of a model trained on the original dataset \mathcal{D}_o . Formally, this can be expressed as:

$$U(\mathcal{D}_s, \mathcal{M}) \approx U(\mathcal{D}_o, \mathcal{M}), \quad (2)$$

where $U(\mathcal{D}_o, \mathcal{M})$ denotes the utility of \mathcal{D}_o for model \mathcal{M} . For simplicity, this aspect is referred to as **Machine Learning Efficacy (MLE) utility**^{[26][27]} in this paper. In addition, in terms of data enhancement^{[10][28]}, we aim for the machine learning model \mathcal{M} trained on the combined dataset $\mathcal{D}_o + \mathcal{D}_s$ to achieve performance on downstream tasks that is comparable to or surpasses that of a model trained exclusively on the original dataset \mathcal{D}_o . Formally, this can be expressed as:

$$U(\mathcal{D}_o + \mathcal{D}_s, \mathcal{M}) \gtrsim U(\mathcal{D}_o, \mathcal{M}) \quad (3)$$

For simplicity, this aspect is referred to as **augmentation utility** in this paper.

By ensuring these properties, the synthetic dataset \mathcal{D}_s can effectively supplement or substitute for the original dataset \mathcal{D}_o in various applications, providing a valuable tool for data argumentation and analysis while preserving privacy and

confidentiality.

2.2. SampleLLM Overview

As illustrated in Figure 2, SampleLLM consists of two stages. In the first stage, the designed instruction and sampled exemplars from the original dataset \mathcal{D}_o are selected to serve as the input for the LLM, producing synthetic tabular data $\mathcal{D}_{s'}$. In the second stage, a novel feature attribution-based importance sampling operation is performed to achieve further feature relation modeling and distributional alignment between the synthetic and original datasets, resulting in the final synthetic dataset \mathcal{D}_s .

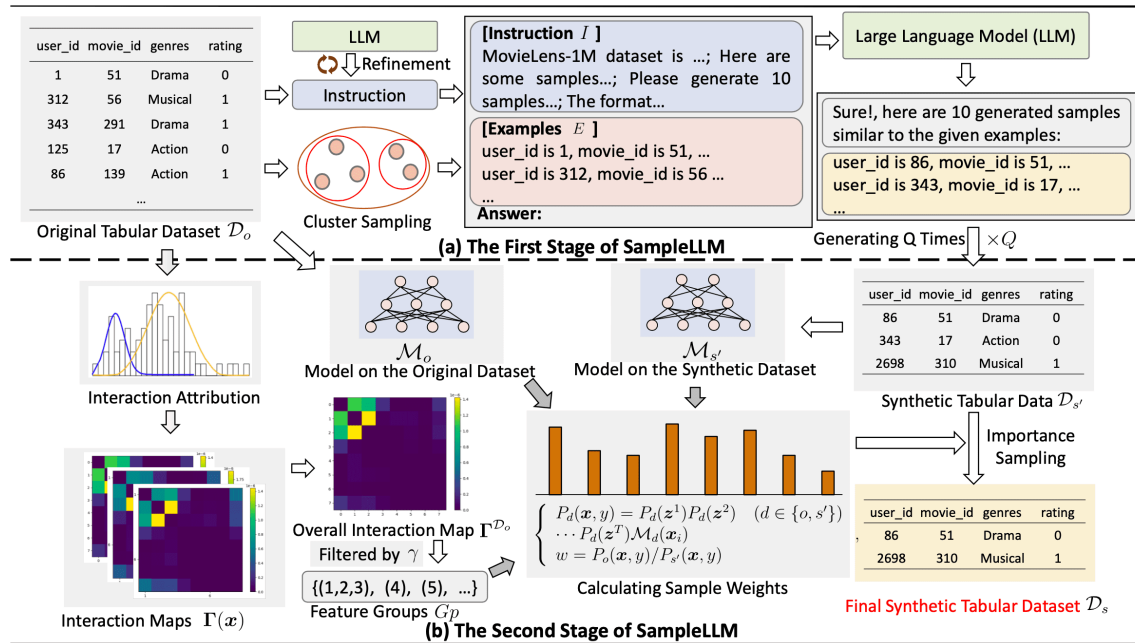


Figure 2. The overall structure of SampleLLM. (a) In the first stage, a manually designed instruction and a samples extracted with clustering sampling are used as inputs to the LLM, which generates b synthetic samples. This process is repeated Q times. (b) In the second stage, a novel feature attribution-based importance sampling method is employed on the synthetic samples.

Specifically, in the first stage, an instruction I is designed to explain the data generation task to the LLM. This instruction is first selected from a set of manually crafted instructions based on performance criteria, and then refined by LLM with official documents and a small number of samples in a Chain-of-Thought (CoT) manner to extract key information. Following this, a cluster sampling method is applied to the original tabular data to generate a exemplars represented by E . This ensures that the data generated by the LLM is diverse and closely resembles the original distribution. By combining the instruction and exemplars as inputs, the LLM generates b synthetic samples using a few-shot learning strategy. This process is repeated Q times to produce a total of bQ synthetic samples, where a , b , and Q are hyper-parameters. In the second stage, an innovative importance sampling method is applied to the synthetic dataset $\mathcal{D}_{s'}$ to further align its distribution with that of the original data \mathcal{D}_o . The weight of each sample is determined by the joint distribution probability of the features and labels under a semi-independent hypothesis where the non-independent

feature groups are obtained through the overall interaction map of the original dataset and the discriminant probabilities of different labels are given by a predictive model trained on the corresponding dataset.

An simple example of the instruction

The MovieLens 1M movie rating dataset is a stable benchmark dataset with 1 million ratings from 6000 users on 4000 movies. Here are some samples of this dataset: *{a Exemplars}*

According to the format, distribution, content, column name and column relation of the above samples, as well as the provided descriptions and your logic and general knowledge, please generate *{b}* new data for data enhancement use in the same format directly. The new samples must not exist in the given data. The samples should also conform to the semantic meaning of each column provided by the column title and the internal logic in the feature combination.

Figure 3. A simplified instruction example.

2.3. Designed Instruction

Traditional statistical and deep learning models are predominantly data-driven, necessitating substantial volumes of samples to accurately model data distributions. This reliance significantly limits their effectiveness when dealing with recommendation tasks with small and sparse datasets, where synthetic data generation becomes crucial. Furthermore, when it comes to textual features, these models typically use ID encodings as input, hindering their ability to capture the semantic associations between features. To address these limitations, we leverage the general knowledge and semantic understanding capabilities of LLMs for data synthesis via few-shot learning. A pivotal component of this approach is the precise framing of the tabular data synthesis task within the prompt, referred to as the instruction, as illustrated in Figure 3. Here, we identify the most effective instruction from a set of manually crafted options tailored to each dataset. In our study, the final instruction for each dataset is selected from five custom-designed instructions. Specifically, for each dataset, 1% synthetic samples are generated and incorporated into the training set. The instruction that delivers the best performance is then established as the standard for that dataset. To further enhance the instruction with relevant knowledge, we input the manually curated instruction, along with official documents O and a few data samples S , into the LLM, which iteratively refines the instruction expression using a Chain-of-Thought process. This procedure is described as follows:

$$I_{t+1} = LLM(I_t, O, S) \quad (4)$$

where the maximum value of iteration times t is fixed at 5. For simplicity, we denote the final output instruction as I .

2.4. Selecting Exemplars

As the key to few-shot learning, exemplars are crucial for providing LLMs with a small number of examples that standardize the format, content, and distribution of LLM-generated samples. However, current studies^[29] often select exemplars randomly from the original dataset. This approach may fail to capture the diverse distribution of the original

data, especially since the number of exemplars is limited by the maximum input tokens of the LLM. Furthermore, the intrinsic distribution differences between LLM’s inherent knowledge and target dataset exacerbate this issue. Consequently, as shown in Figure 1, significant distribution discrepancies are observed between the generated and original tabular samples.

To address these problems, two alignment methods are applied in SampleLLM. The first method aims to maintain sample diversity during exemplar selection to avoid distribution concentration, as detailed below. The second method discussed further in Section 2.5, leverages a novel feature attribution-based importance sampling method to model critical feature relations and address the distribution differences caused by the LLM.

To ensure exemplars reflect the diverse characteristics of the original data and mitigate the sample aggregation issue in LLM output, a simple clustering-based sampling method is first proposed in SampleLLM for exemplar selection. This approach is mathematically represented as follows:

$$\begin{cases} Clu = h(\mathcal{D}_o, a) = \{clu_1, clu_2, \dots, clu_a\} \\ E = \{x_i \mid x_i = g(clu_i); i \in \{1, 2, \dots, a\}\} \end{cases} \quad (5)$$

Here, a clusters $\{clu_1, clu_2, \dots, clu_a\}$ are derived from the original dataset \mathcal{D}_o using a clustering algorithm h —specifically, the K-means method. From each cluster, one exemplar is selected using method g —in this case, random selection—to form an exemplar set E . This approach ensures that each group of exemplars contains samples from diverse regions of the original dataset’s distribution, thereby enhancing diversity and achieving better alignment with the distribution characteristics of the original data.

As illustrated in Figure 2 (a), due to the maximum output token limitations in LLMs, the entire first stage of SampleLLM will iterate Q times, with b samples generated by LLM each time, where Q and b are hyper-parameters. After generation, the generated bQ samples will be formatted in tabular form and aggregated together to form a temporary synthetic tabular dataset \mathcal{D}_s .

2.5. Feature Attribution-based Importance Sampling

Although cluster sampling methods are designed to mitigate discrepancies in data distributions, inherent differences caused by the LLM’s input-output processing still result in a distribution gap and differentiated feature interaction frequency between the generated and the original data. This gap poses significant challenges in recommendation scenarios, where feature associations serve as a foundational modeling basis. To address these issues, SampleLLM employs a feature attribution-based importance sampling method to further refine the feature relation and alignment of distribution in \mathcal{D}_s with that in \mathcal{D}_o .^{[30][31]} Our method is proposed because directly calculating the exact importance weights for tabular samples is challenging due to the high dimensionality of tabular data, which complicates the estimation of the joint distribution of the features and labels in samples:

$$P(x_i) = P(x_i^1, x_i^2, \dots, x_i^K)P(y_i \mid x_i^1, x_i^2, \dots, x_i^K) \quad (6)$$

One straightforward approach to simplify the calculation is to assume independence among all features:

$$P(x_i) = P(x_i^1)P(x_i^2) \dots P(x_i^K)P(y_i \mid x_i^1, x_i^2, \dots, x_i^K) \quad (7)$$

However, this assumption often diverges significantly from reality, as it is generally impractical to assume complete independence among all features in most datasets. Consequently, this approach may lead to inaccuracies in estimating distribution probabilities, thereby complicating the data alignment process.

To mitigate these issues, SampleLLM adopts a semi-independence assumption where feature fields containing significant interactions (e.g., fields 1 and 8) are considered non-independent, while others are treated as independent. This balanced approach reduces computational overhead while minimizing calculation errors, enabling the computation of importance weights based on frequency statistics:

$$P(x_i) = P(x_i^1, x_i^8)P(x_i^2) \dots P(x_i^K)P(y_i | x_i^1, x_i^2, \dots, x_i^K) \quad (8)$$

2.5.1. Feature Interaction Extraction

To determine the distribution probability of a sample within a specific dataset, as specified in Equation (8), it is essential to identify significant feature interactions. This paper introduces a feature attribution-based method for extracting such interactions.

For a dataset \mathcal{D} linked to a particular prediction task and a deep learning-based predictive model f , we could quantify the importance of each feature to the output by approximating model performance degradation after ablating this feature with 1st-order Taylor approximation^[32] as follows:

$$EG_i(x) := \mathbb{E}_{x' \sim \mathcal{D}^b} \left[(x_i - x_i^b) \frac{\delta f(x)}{\delta x_i} \right] \quad (9)$$

where x_i are features in x , x_i^b are non-informative features like zero feature or random feature and \mathcal{D}^b is typically a randomly selected subset of \mathcal{D} or a single sample containing all-zero values^[33]. This approach quantifies the importance of each feature to the output by approximating model performance change after setting an informative feature to non-informative.

To attribute feature interactions, we extend 1st-order Taylor expansion to 2nd-order, capturing the importance of feature i to the importance of feature j . The resulting expression is:

$$\Gamma_{i,j}(x) = \mathbb{E}_{x \sim \mathcal{D}^b} \left[(x_i - x_i^b) (x_j - x_j^b) \frac{\partial^2 f(x^b)}{\partial x_i \partial x_j} \right] \quad (10)$$

Intuitively, this method measures feature interaction strength as approximated feature importance change when switching interaction from informative to non-informative.

By applying this interaction attribution method to sample samples from the training set, we could capture the complex interactions between feature fields across the entire dataset. Therefore, summing the absolute values of interaction maps $\Gamma(x)$ for all samples yields a comprehensive matrix that highlights the most significant feature interactions in the whole dataset, irrespective of whether their impacts on the output are positive or negative:

$$\Gamma^{\mathcal{D}_o} = \sum_{x \sim \mathcal{D}_o} |\Gamma(x)| \quad (11)$$

Interaction values above a threshold, defined by a hyper-parameter γ , are then clustered to delineate independent and non-independent feature groups, further refining data alignment:

$$\begin{cases} F_\gamma = \gamma \cdot \max(\Gamma^{\mathcal{D}_o}) \\ Pairs = \text{argwhere}_{\Gamma^{\mathcal{D}_o}}(\Gamma^{\mathcal{D}_o} > F_\gamma) \\ Gp = \text{Merge}(Pairs) \end{cases} \quad (12)$$

where the *Merge* function aggregates all pairs with overlapping feature fields to form a feature group z^t (e.g., $z^1 = \{x^1, x^3, x^8\}$) with several feature fields, resulting in T groups $Gp = \{z^1, \dots, z^t, \dots, z^T\}$ where T is determined by the aggregation process, ensuring $T \leq K$.

In this study, our objective is to align the distribution of the dataset $\mathcal{D}_{s'}$ with that of the original dataset \mathcal{D}_o . Consequently, the feature groups are derived from \mathcal{D}_o and are utilized for calculating the distribution probability of samples in both \mathcal{D}_o and $\mathcal{D}_{s'}$. This methodology facilitates a refined understanding of the underlying feature interactions, thus enhancing model interpretability and improving alignment between datasets.

2.5.2. Obtaining Discriminant Probability

As a dataset, each tabular sample contains not only a set of features but also a label. Therefore, another critical challenge in calculating the probability of a sample is obtaining the discriminant probability $P(y_i | x_i^1, x_i^2, \dots, x_i^K)$ for each sample $x_i = \{x_i^1, x_i^2, \dots, x_i^K, y_i\}$. In SampleLLM, this is achieved by training a task-specific predictive model \mathcal{M} on the corresponding dataset:

$$\mathcal{M}(x_i) = P(y_i | x_i^1, x_i^2, \dots, x_i^K) \quad (13)$$

For a specific task (e.g., binary classification), two predictive models \mathcal{M}_o and $\mathcal{M}_{s'}$ are trained to predict the sample discriminant probabilities on \mathcal{D}_o and $\mathcal{D}_{s'}$, respectively.

2.5.3. Calculating Sample Weights

After obtaining the feature groups Gp and discriminant models \mathcal{M}_o and $\mathcal{M}_{s'}$, the importance weights of samples can be derived using the following formula:

$$\begin{cases} P_d(x, y) = P_d(z^1)P_d(z^2) \quad (d \in \{o, s'\}) \\ \dots P_d(z^T)\mathcal{M}_d(x_i) \\ w = P_o(x, y)/P_{s'}(x, y) \end{cases}, \quad (14)$$

where the first term is an alternative expression of Equation (8), and the second term is used to calculate the importance weights based on the sample's distribution probability in both the original and the synthetic datasets. $P_d()$ is the frequency of related features in dataset d .

Finally, after obtaining all synthetic samples' weights, the final synthetic tabular dataset \mathcal{D}_s is created through an importance sampling on $\mathcal{D}_{s'}$. The overall pseudo-code is also provided in Appendix A to further clarify the overall flow of SampleLLM.

3. Experiments

In this section, we conduct experiments on three recommendation datasets to address the following research questions:

- **RQ1:** How does SampleLLM perform in comparison to other tabular data synthesis methods?
- **RQ2:** How effective is the generated tabular data when used as a replacement for the original tabular data?

- **RQ3:** What impact do cluster sampling and feature attribution-based importance sampling in SampleLLM have in the process of generating synthetic tabular data?

Moreover, two extensive general datasets are also applied to validate SampleLLM’s effectiveness in general tabular data-related tasks. In the following sections, we first describe the experimental setup used in this study in the following subsections. Subsequently, detailed analyses are provided for each research question based on our experimental results.

3.1. Experimental Setup

3.1.1. Dataset

We conducted experiments on three commonly used recommendation datasets with Click-Through Rate (CTR) prediction tasks^{[34][35]} to verify the validity of SampleLLM, i.e., ML-1M¹, Amazon² and Douban³. In addition, a binary classification dataset HELOC⁴ and a multi-classification dataset Coverttype⁵ are further applied for extended experiments to verify the potential of SampleLLM in general tasks. In accordance with previous studies^{[36][37]}, each dataset is divided into training/validation/test sets in an 8:1:1 ratio. Data statistics could be found in Table 1.

Dataset	Train	Validation	Test
ML-1M	800,167	100,020	100,022
Amazon	658,827	82,353	82,354
Douban	1,348,399	168,549	168,551
HELOC	8,367	1,045	1,047
Coverttype	464,809	58,101	58,102

Table 1. Data Statistics.

3.1.2. Baselines

To verify the effectiveness of SampleLLM, a comprehensive comparison is conducted between SampleLLM and various baseline approaches:

- **CTGAN**^[26] utilizes a conditional generator and a mode-specific normalization method in synthetic tabular data generation to address the complexities of tabular data, including mixed data types and imbalanced categories.
- **TVAE**^[26], a Variational AutoEncoder (VAE) adapted for synthetic tabular data generation, leverages specialized preprocessing and a modified loss function to handle mixed data types.
- **TabDDPM**^[38] leverages a denoising diffusion probabilistic model to handle the inherent heterogeneity of tabular data, providing superior generative performance across benchmarks.

- **PATE-GAN**^[40] leverages the Private Aggregation of Teacher Ensembles (PATE) framework to introduce differential privacy guarantees into the Generative Adversarial Network (GAN) setting. By modifying the GAN discriminator with the PATE mechanism, PATE-GAN ensures that the synthetic data generated maintains privacy while preserving data utility.
- **ADS-GAN**^[40] is designed to generate synthetic data that closely approximate the joint distribution of variables in original datasets. Utilizing a conditional Generative Adversarial Network (GAN) framework, ADS-GAN ensures data anonymization by minimizing identifiability based on the probability of re-identification.
- **GReaT**^[8] leverages an auto-regressive generative LLM to synthesize realistic tabular data. This method efficiently models tabular data distributions by conditioning on any subset of features, allowing for flexible and authentic data generation.

REaLTabFormer^[27] generates high-quality synthetic data for both non-relational and relational datasets. It employs an autoregressive model for parent tables and a sequence-to-sequence model for related child tables with GPT-2 LLM as backbone modules, ensuring realistic data relationships and preventing data copying through target masking and statistical bootstrapping.

3.1.3. Implementation Details

For ML-1M, Amazon, Douban and HELOC datasets, the widely used metrics of AUC and Logloss are deployed for evaluation^{[41][42]}. For Covertypes whose task is multi-class classification, the widely used metrics of weighted Precision, Recall, and F1 score^[43] are deployed for evaluation. For a fair comparison and clearly illustrate the effectiveness of each module in SampleLLM. A simple Deep Neural Network (DNN) is selected as the predictive model for all tasks. Specifically, for different datasets and tasks, we first carried out a grid search for network parameters to ensure prediction performance, in which the search range for the number of layers is {1-4} and the search range for the number of neurons per layer is {16,32,64,128}. The simple grid search is also applied in selecting other hyper-parameters such as learning rate and dropout rate in training. Meanwhile, the INT4 quantized version of Llama-3-70b-instruct-awq^[44] is applied as the LLM backbone for the first stage of SampleLLM. In addition, we ran each experiment 10 times and reported the average performance.

3.2. Overall Performance (RQ1, 2)

This section gives overall performance comparisons between SampleLLM and various baselines on the augmentation utility and MLE utility as illustrated in Section 2.1 to answer the RQ1 and RQ2.

3.2.1. Augmentation Utility

As illustrated in Equation (3), the utility of synthetic data could be evaluated by the data enhancement effect after injecting synthetic data into the original data^{[10][28]}. Specifically, we generate 10% synthetic data for the training set using different baselines, and integrate them with the original training set for training. The performance of the final

predictive model could, therefore, serve as evidence to measure the utility of the synthetic data. The experimental results are presented in Table 2. A case study is also provided in Appendix C for visualization analysis.

Approach	ML-1M		Amazon		Douban		HELOC		CoverType		
	AUC ↑	Logloss ↓	AUC ↑	Logloss ↓	AUC ↑	Logloss ↓	AUC ↑	Logloss ↓	Precision ↑	Recall ↑	F1 ↑
Original	0.8173	0.5163	0.7075	0.4602	0.8016	0.5158	0.7638	0.6445	0.7395	0.7339	0.7250
PATE-GAN	0.8154	<u>0.5153</u>	0.7022	0.4620	0.8010	0.5162	0.7648	0.6461	0.7388	0.7350	0.7258
ADS-GAN	0.8147	0.5157	0.7026	0.4619	0.8016	0.5151	0.7672	0.6451	0.7414	0.7413	0.7312
CTGAN	0.8148	0.5156	0.7027	0.4631	0.8016	0.5150	0.7677	0.6419	0.7417	0.7366	0.7295
TVAE	0.8143	0.5162	0.7031	0.4616	0.8019	0.5146	0.7701	0.6455	0.7422	0.7378	0.7305
TabDDPM	0.8141	0.5159	0.7036	0.4619	0.8020	0.5150	0.7704	0.6424	0.7418	0.7388	0.7299
GReaT	0.8153	0.5198	0.7040	0.4616	0.8021	0.5147	0.7703	0.6444	0.7423	0.7440	<u>0.7317</u>
REaLTabFormer	<u>0.8156</u>	0.5182	<u>0.7041</u>	<u>0.4611</u>	<u>0.8022</u>	<u>0.5145</u>	<u>0.7707</u>	<u>0.6417</u>	<u>0.7428</u>	0.7351	0.7264
SampleLLM	0.8180*	0.5140*	0.7082*	0.4601*	0.8027*	0.5139*	0.7732*	0.6403*	0.7445*	<u>0.7437</u>	0.7364*

Table 2. Augmentation utility. **The boldface denotes the best score. The underline indicates the second-best score.** The “Original” method indicates training without synthetic data and is not considered in scoring. ↑: higher is better; ↓: lower is better. “*” indicates the statistically significant improvements (i.e., two-sided t-test with $p < 0.05$) over the best baseline.

From Table 2 we can conclude that:

- For baselines without LLM in modeling, TabDDPM and TVAE outperform GAN-based models in most cases, verifying the effectiveness of more detailed modeling for distribution alignment in diffusion probabilistic models and VAE models.
- For the ML-1M, Amazon, and Douban datasets, traditional models generally struggle to match the performance seen on the original data. However, they tend to perform better on the HELOC and Covertype datasets. This disparity highlights the importance of complex feature associations in recommendation modeling, which traditional models fail to capture due to their inability to understand semantic relationships. GReaT and REaLTabFormer outperform other baselines, showcasing LLMs’ strength in modeling feature relations through semantic comprehension. Additionally, by utilizing alignment methods, SampleLLM achieves superior AUC and Logloss across all datasets.
- SampleLLM outperforms all baselines in most cases, illustrating the effectiveness of LLM’s semantic understanding in modeling feature relations and data distributions and the effectiveness of the feature attribution-based importance sampling method in further aligning the feature correlations and data distribution of the generated data with the original data.

3.2.2. MLE Utility

As illustrated in Equation (2), the utility of synthetic data could also be evaluated by supplanting the real data in training a predictive model. The corresponding evaluation method is referred to as Machine Learning Efficacy (MLE) utility^[26] [27]. Specifically, to reduce the cost of training time, we generate 10% synthetic data for the training set using different baselines and directly use them as the training sets for model training. The performance of the final predictive model could, therefore, serve as evidence to measure the utility of the synthetic data. The experimental results are presented in Table 3.

Approach	ML-1M		Amazon		Douban		HELOC		Covertypes		
	AUC ↑	Logloss ↓	AUC ↑	Logloss ↓	AUC ↑	Logloss ↓	AUC ↑	Logloss ↓	Precision ↑	Recall ↑	F1 ↑
Original	0.8173	0.5163	0.7075	0.4602	0.8016	0.5158	0.7638	0.6445	0.7395	0.7339	0.7250
PATE-GAN	0.4944	0.9450	0.5800	0.5001	0.7560	<u>0.5600</u>	0.6215	0.6906	0.7126	<u>0.7152</u>	0.7006
ADS-GAN	0.5590	<u>0.6771</u>	0.5802	0.5819	0.7548	0.5624	0.7211	0.9640	0.7069	0.7115	0.6953
CTGAN	0.5602	0.6774	0.5775	0.6092	0.7551	0.5603	0.6289	0.6817	0.7055	0.7132	0.6991
TVAE	<u>0.5622</u>	0.6776	0.5828	0.5981	0.7552	0.5610	0.7627	0.5856	0.7070	0.7091	0.6906
TabDDPM	0.5334	0.6806	0.5853	0.5005	0.7556	0.5603	0.7388	0.6028	0.7105	0.7132	0.6993
GReaT	0.5596	0.8181	0.5881	<u>0.4986</u>	0.7562	0.5615	0.7444	0.5987	0.7111	0.7134	0.7010
REaLTabFormer	0.5574	0.7718	<u>0.5900</u>	0.4992	<u>0.7567</u>	0.5604	<u>0.7628</u>	<u>0.5826</u>	<u>0.7175</u>	0.7103	<u>0.7023</u>
SampleLLM	0.5668*	0.6758*	0.5922*	0.4977*	0.7590*	0.5571*	0.7652*	0.5799*	0.7196*	0.7159*	0.7086*

Table 3. MLE utility. The **boldface** denotes the best score, and the underline indicates the second-best score. The “Original” method indicates training with the original training set and is not considered in scoring. ↑: higher is better; ↓: lower is better. “*” indicates the statistically significant improvements (i.e., two-sided t-test with $p < 0.05$) over the best baseline.

From Table 3 we can conclude that:

- Almost all methods exhibit poorer performance compared to directly using the original training set (“Original”). This result is intuitive, as synthetic data cannot fully replicate or substitute the nuances of the original data. Interestingly, SampleLLM surpasses “Original” on the HELOC dataset, likely due to the dataset’s data scarcity. The model’s convergence rate may remain similar on HELOC whether trained on the original training set or synthetic data.
- The performance degradation across all methods is most pronounced on the ML-1M dataset. This is likely due to the discrete, sparse, and strongly correlated nature of the features in ML-1M. Notably, SampleLLM outperforms other baselines across all datasets on both augmentation and MLE utility measures, demonstrating its suitability for tabular

data synthesis on recommendations and even more various downstream tasks such as binary classification and multi-class classification.

3.3. Ablation Study (RQ3)

This section verifies the effectiveness of each module in SampleLLM through ablation experiments to answer RQ3. Specifically, we compare the augmentation utility of SampleLLM on the Douban dataset with the following substitutions:

- **w/o LLM (-LLM)**: with TabDDPM, a non-LLM baseline as the alternative of the LLM in the first stage of SampleLLM;
- **w/o cluster sampling (-CS)**: without cluster sampling in the first stage of SampleLLM. The exemplars are selected randomly.
- **w/o importance sampling (-IS)**: without the feature attribution-based importance sampling in the second stage of SampleLLM.

We also calculate their SDV similarity^[45], a score measuring distribution similarity considering column shapes and pair trends. The result is shown in Table 4.




Figure 4. Ablation study of augmentation utility on Douban.

Models	SampleLLM	-LLM	-CS	-IS
SDV Similarity	94.83%	92.32%	94.13%	93.26%

Table 4. SDV Similarity Score with the original dataset.

Based on Figure 4 and Table 4, we could conclude that:

- Replacing the LLM backbone in the first stage of SampleLLM with the TabDDPM results in a significant performance decline for SampleLLM. This decline may be attributed to LLM’s superior ability to model feature relations and joint distributions at the semantic level, which allows it to better approximate the sample distribution of the original data

compared to baseline methods. This characteristic is crucial for ensuring the effectiveness of feature attribution-based importance sampling in the second stage of SampleLLM, as the distribution similarity between the two datasets positively impacts the effectiveness of the importance sampling method^[46] when the sample amount is fixed.

- Without the cluster sampling method in selecting exemplars at the first stage, SampleLLM also suffers from performance decline, indicating that selecting diverse exemplars is essential in LLM's few-shot learning process.
- Without the feature attribution-based importance sampling in the second stage, SampleLLM's performance degenerates significantly, indicating that refining sample distribution and feature relation is vital in LLM-based tabular data synthesis.

In order to further intuitively illustrate the effectiveness of different modules, a visualization analysis is conducted through the T-Distributed Stochastic Neighbor Embedding (TSNE) method for the synthetic tabular data generated by the above alternatives, as shown in Figure 5. From these subfigures, we could observe that:

- As depicted in Figure 5(b), substituting the TabDDPM baseline model for the LLM in the first stage fundamentally results in a worse synthetic data distribution (especially in the middle part), despite the integration of feature attribution-based importance sampling in the subsequent stage. This observation substantiates our initial analysis.
- As shown in Figure 5(c), the selection of random exemplars over those derived from cluster sampling leads to an overly concentrated synthetic data distribution in comparison with Figure 5(a) when generated through the LLM.
- Although Figure 5(d) shows a relatively similar shape, the point density in different regions differs greatly from the original dataset. This demonstrates that, in the absence of feature attribution-based importance sampling, achieving a comprehensive alignment with the original dataset is challenging.

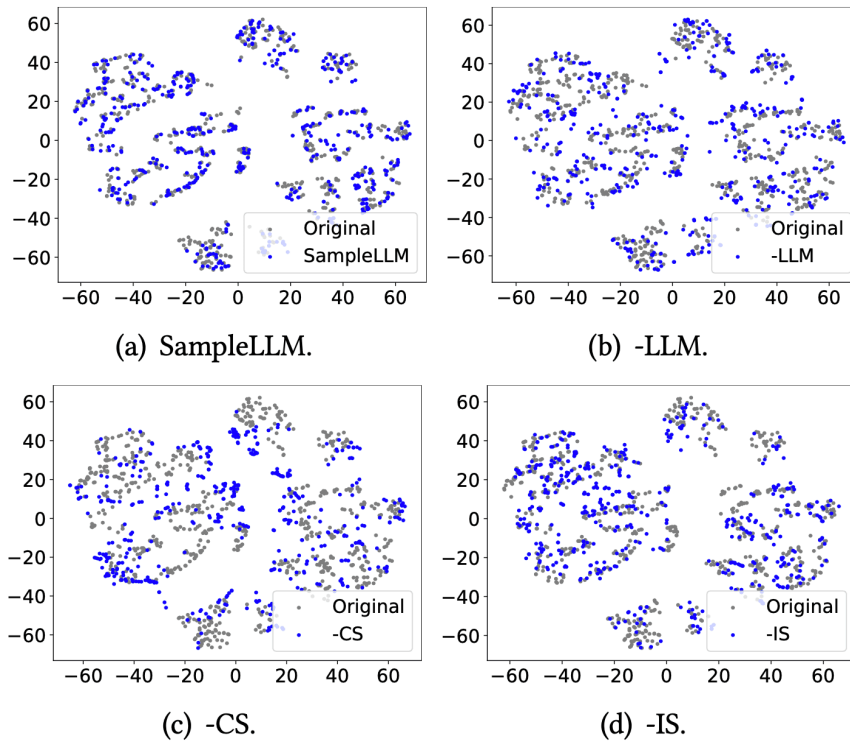


Figure 5. Visualization analysis of the ablation study.

4. Online A/B Test

The Huawei app advertising platform contends with the challenge of low recall rates for cold-start apps, primarily due to the sparse user interaction data resulting from limited exposure. This scarcity hinders existing models from achieving global optimality. To address this challenge, we implemented SampleLLM to generate synthetic interactions between users and cold-start apps. In the user profile generation phase, clustering sampling was employed to iteratively select user profile seeds, which were then utilized by the LLM to synthesize diverse user profiles, simulating app exposure across a wide range of user groups. The LLM further integrated these profiles with the cold-start app's features and descriptions to predict user click behaviors, preserving the result as a raw dataset for distribution alignment.

Subsequently, we decomposed the complete feature sets into several strong interaction feature clusters using Hessian-based methods on a previously trained recommendation model. The distributions of these feature clusters were calculated based on occurrence frequency. This allows us to determine the importance weights of each synthetic sample, which serves as augmenting data to improve training for cold-start apps.

For deployment, we downsampled one day's user data and clustered it into 1,000 groups using Spark. From each cluster, samples were selected as seeds according to cluster size for subsequent data synthesis. The selected user profiles, combined with app descriptions, served as seeds for Qwen2.5 to simulate user click behaviors on the apps. The synthesized samples were then assigned importance weights based on statistics extracted from the current day's

recommendation model, ensuring alignment with real user behavior for non-cold-start apps. These synthetic samples are utilized as augmentation for the following week without further updates.

We integrated the model trained with synthetic data as a new pipeline in the user-item recalling module. During the experimental phase, real cold-start samples were updated daily, while synthetic data received incremental updates weekly. The experiment lasted a total of one week and exhibited a 2% improvement in RPM (Revenue Per Mille), a 1.86% improvement in ECPM (Effective Cost Per Mille) in our cold-start scenario, thereby validating the method's effectiveness.

5. Related Work

Tabular data synthesis has evolved to address diverse needs across domains, initially through probabilistic models like the Synthetic Data Vault (SDV)^[45] that focused on privacy and statistical integrity. The introduction of deep learning approaches, including GANs^{[26][40][39]} and VAEs^[26], has furthered this field, offering solutions to issues such as data scarcity and privacy. Models like CTGAN and TVAE^[26] effectively handle specific challenges, including mixed data types, significantly enhancing the quality of synthetic datasets. Specialized techniques, such as TabDDPM^[38], leverage models like denoising diffusion probabilistic models for better generative performance across benchmarks^{[30][47]}. However, these methods often require large datasets, limiting their effectiveness in data-scarce contexts, especially in recommendation scenarios^{[48][49][50][51]}. Moreover, they fail to capture the semantic associations between features^[52]^[53], which are increasingly important in recommendation modeling.

The rise of transformer^[54] architectures^[55] and LLMs^{[56][57][58][59]} has introduced new possibilities for tabular data synthesis, thanks to LLMs' few-shot learning^{[16][60][61]} and semantic understanding capabilities^[62]. Methods like GReaT^[8] and ReaLTabFormer^[27] have explored converting tabular data into natural language to leverage LLM strengths. Despite promising outcomes, these approaches often fail to address alignment issues between LLMs and target datasets, leading to distribution divergence. Instead, our proposed method, SampleLLM, addresses these shortcomings by combining few-shot learning with alignment techniques to better match distributions and feature relationships of generated and original datasets, thereby enhancing data utility and downstream task performance in recommendations.

6. Conclusion

In this paper, a two-stage LLM-based framework SampleLLM is proposed to integrate LLM with sampling methods for optimizing tabular data synthesis in recommendations. Specifically, a manually designed instruction, together with a group of exemplars generated through a cluster sampling method serves as the input prompt for few-shot learning via LLM. Then a novel feature attribution-based importance sampling method is proposed to serve as a second stage for further feature relation modeling and distribution alignment. By doing so, SampleLLM is able to generate synthetic tabular data with semantic understanding, higher utility, and distribution similarity. Experiments on three public

recommendation datasets, two general datasets, and an online application demonstrate the effectiveness of the proposed SampleLLM.

Appendix A. Pseudo-code of SampleLLM

To elucidate the overall process of SampleLLM, we present the pseudo-code in Algorithm 1. Specifically, SampleLLM operates in two distinct stages. The first stage (i.e., lines 1-7) employs a cluster sampling technique for exemplar selection in LLM’s few-shot learning, aimed at generating an initial synthetic tabular dataset, denoted as $\mathcal{D}_{s'}$. In the second stage (i.e., lines 8-12), a novel feature attribution-based importance sampling method is proposed to enhance the modeling of feature relationships and achieve distribution alignment between the synthetic dataset $\mathcal{D}_{s'}$ and the original tabular dataset \mathcal{D}_o . This process results in the final synthetic tabular dataset, \mathcal{D}_s .

Algorithm 1 The whole process of SampleLLM

Input: The original tabular dataset \mathcal{D}_o ; An LLM model LLM ;

Hyper-parameters a, b, Q, γ

Output: The final synthetic tabular dataset \mathcal{D}_s

- 1: $\mathcal{D}_{s'} = []$
 - 2: **for** Iteration in $1, \dots, Q$ **do**
 - 3: Obtain instruction I for \mathcal{D}_o through Section 2.3
 - 4: Obtain exemplar set E with a exemplars via Equation (5)
 - 5: Obtain output synthetic data \mathcal{D}_{tmp} with b samples via $LLM(I + E)$
 - 6: $\mathcal{D}_{s'} = \mathcal{D}_{s'} + \mathcal{D}_{tmp}$
 - 7: **end for**
 - 8: Obtain $\Gamma^{\mathcal{D}_o}$ via Equation (10), (11)
 - 9: Obtain feature groups Gp via Equation (12)
 - 10: Obtain discriminant probability \mathcal{M}_o and $\mathcal{M}_{s'}$ via Equation (13)
 - 11: Obtain sample weight w for samples in $\mathcal{D}_{s'}$ via Equation (14)
 - 12: Obtain \mathcal{D}_s by processing importance sampling on $\mathcal{D}_{s'}$ with sample weight w .
-

Appendix B. Overall Interaction Map and Feature Groups

In this section, we present the overall interaction maps and feature groups for all five experimental datasets, as illustrated in Figure 6, and Table 5. It is important to note that the interaction maps and feature groups may exhibit slight variations from the provided figures and tables due to fluctuations in the performance of the predictive model trained on each dataset.

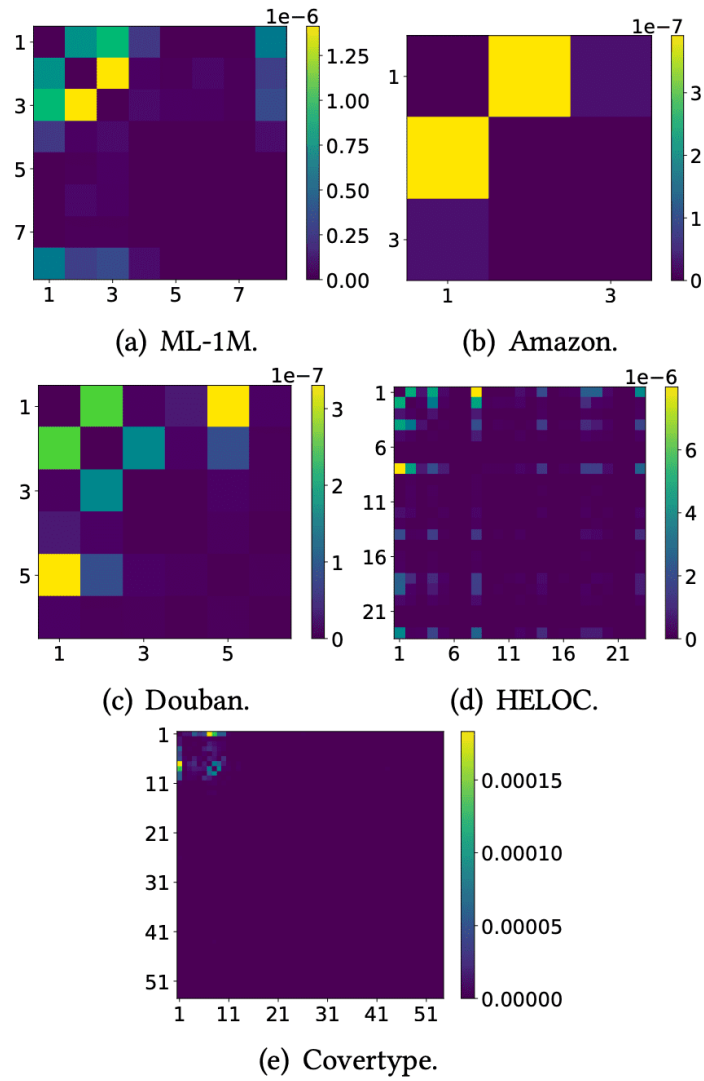


Figure 6. Interaction maps for different datasets. The numbers on the axes represent the indices of the feature fields.

Dataset	Feature Group
ML-1M	[1, 2, 3]
Amazon	[1, 2]
Douban	[1, 2, 5]
HELOC	[1, 8]
Coverttype	[1, 7, 8]

Table 5. Feature groups. Note that only groups with more than one feature field are shown in the table for simplicity.

Appendix C. Case Study

This section presents a case study using TSNE visualizations and SDV similarity^[45], a score measuring distribution similarity considering column shapes and pair trends, to analyze the synthetic tabular data generated by high-performance baselines such as TVAE, TabDDPM, and REaLTabFormer, thereby highlighting the advantages of SampleLLM. As shown in Figure 7 and Table 6, TVAE and TabDDPM struggle to effectively replicate the original data distribution. Even with the incorporation of LLM, REaLTabFormer’s distribution remains concentrated in several central regions, as noted in Figure 1. In contrast, SampleLLM demonstrates significant improvements in aligning with the original sample distribution by leveraging the few-shot learning capabilities of LLM and employing a novel two-stage alignment strategy.

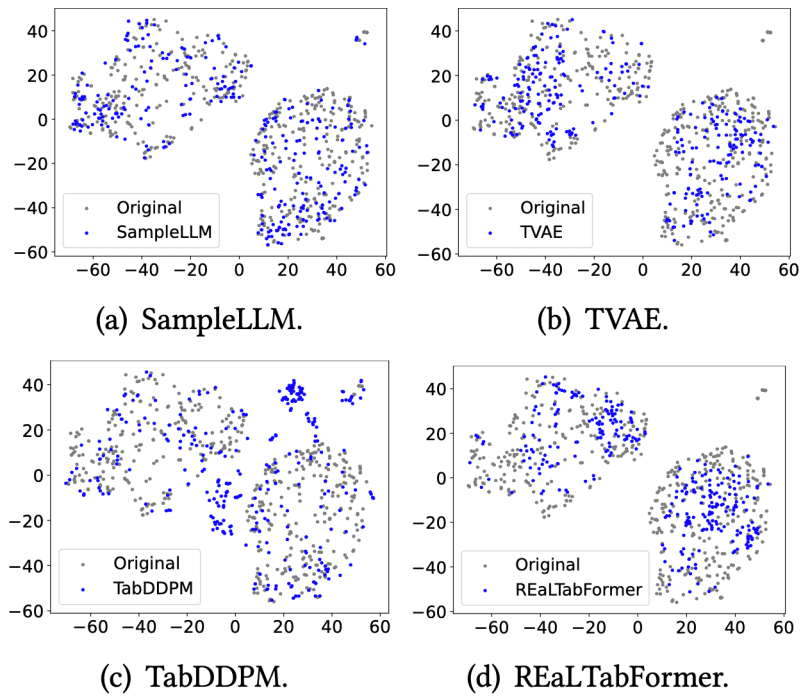


Figure 7. Case study on HELOC dataset.

Models	SampleLLM	TVAE	TabDDPM	REaLTabFormer
SDV Similarity	90.10%	85.96%	84.84%	87.23%

Table 6. SDV Similarity Score with the original dataset.

Acknowledgements

This research was partially supported by Research Impact Fund (No.R1015-23), Collaborative Research Fund (No.C1043-24GF), APRC - CityU New Research Initiatives (No.9610565, Start-up Grant for New Faculty of CityU), Hong Kong ITC Innovation and Technology Fund Midstream Research Programme for Universities Project (No.ITS/034/22MS), and Huawei (Huawei Innovation Research Program).

Footnotes

¹ <https://grouplens.org/datasets/movielens/1m/>

² <https://jmcauley.ucsd.edu/data/amazon/>

³ <https://www.kaggle.com/datasets/fengzhujoey/douban-datasetratingreviewside-information>

⁴ <https://huggingface.co/datasets/mstz/heloc>

⁵ <https://archive.ics.uci.edu/dataset/31/covertime>

References

- ¹ ^ΔHabibi O, Chemmakha M, Lazaar M (2023). "Imbalanced tabular data modelization using CTGAN and machine learning to improve IoT Botnet attacks detection". *Engineering Applications of Artificial Intelligence*. 2023: 105669.
- ² ^ΔMcElfresh D, Khandagale S, Valverde J, Prasad C V, Ramakrishnan G, Goldblum M, White C (2024). "When do neural nets outperform boosted trees on tabular data?" *Proc. of NeurIPS* (2024).
- ³ ^ΔGoel A, Goel AK, Kumar A (2023). "The role of artificial neural network and machine learning in utilizing spatial information". *Spatial Information Research*. 2023: 275–285.
- ⁴ ^Δ [♠] [♠]Bourou S, El Saer A, Velivassaki T-H, Voulkidis A, Zahariadis T (2021). "A review of tabular data synthesis using GANs on an IDS dataset". *Information*. 2021: 375.
- ⁵ ^ΔFayaz SA, Zaman M, Kaul S, Butt MA (2022). "Is deep learning on tabular data enough? An assessment". *International Journal of Advanced Computer Science and Applications*. (2022).
- ⁶ ^ΔHernandez M, Epelde G, Alberdi A, Cilla R, Rankin D (2023). "Synthetic tabular data evaluation in the health domain covering resemblance, utility, and privacy dimensions". *Methods of information in medicine*. 2023: e19–e38.
- ⁷ ^ΔQian Z, Cebere BC, van der Schaar M (2023). "Synthcity: facilitating innovative use cases of synthetic data in different data modalities". *arXiv preprint arXiv:2301.07573*.
- ⁸ ^Δ [♠] [♠] [♠]Borisov V, Se{\ss}ler K, Leemann T, Pawelczyk M, Kasneci G (2022). "Language models are realistic tabular data generators". *arXiv preprint arXiv:2210.06280*.
- ⁹ ^ΔBansal MA, Sharma DR, Kathuria DM (2022). "A systematic review on data scarcity problem in deep learning: solution and applications". *ACM Computing Surveys (Csur)*. 2022: 1–29.
- ¹⁰ ^Δ [♠] [♠] [♠]Figueira A, Vaz B (2022). "Survey on synthetic data generation, evaluation methods and GANs". *Mathematics*. 2022: 2733.
- ¹¹ ^ΔSufi F (2024). "Addressing Data Scarcity in the Medical Domain: A GPT-Based Approach for Synthetic Data Generation and Feature Extraction". *Information*. 2024: 264.
- ¹² ^ΔFonseca J, Bacao F (2023). "Tabular and latent space synthetic data generation: a literature review". *Journal of Big Data*. 2023: 115.
- ¹³ ^ΔLiu Q, Hu J, Xiao Y, Zhao X, Gao J, Wang W, Li Q, Tang J (2024). "Multimodal recommender systems: A survey". *ACM Computing Surveys*. (2024): 1–17.
- ¹⁴ ^ΔLi J, Wang M, Li J, Fu J, Shen X, Shang J, McAuley J (2023). "Text is all you need: Learning language representations for sequential recommendation". In: *Proc. of KDD*. pp. 1258–1267.
- ¹⁵ ^ΔXu D, Chen W, Peng W, Zhang C, Xu T, Zhao X, Wu X, Zheng Y, Wang Y, Chen E (2024). "Large language models for generative information extraction: A survey". *Frontiers of Computer Science*. 2024: 186357.

16. ^a ^b Cahyawijaya S, Lovenia H, Fung P (2024). "LLMs Are Few-Shot In-Context Low-Resource Language Learners". arXiv preprint arXiv:2403.16512. Available from: <https://arxiv.org/abs/2403.16512>.
17. ^Δ Dong Q, Li L, Dai D, Zheng C, Wu Z, Chang B, Sun X, Xu J, Sui Z (2022). "A survey on in-context learning". arXiv preprint arXiv:2301.00234.
18. ^Δ Liu Q, Wu X, Zhao X, Zhu Y, Xu D, Tian F, Zheng Y (2024). "When moe meets llms: Parameter efficient fine-tuning for multi-task medical applications". In: Proc. of SIGIR. p. 1104–1114.
19. ^Δ Gao J, Zhao X, Chen B, Yan F, Guo H, Tang R (2023). "AutoTransfer: instance transfer for cross-domain recommendations". In: Proc. of SIGIR. pp. 1478–1487.
20. ^Δ Liu D, Yang C, Tang X, Wang Y, Lyu F, Luo W, He X, Ming Z, Zhao X (2024). "MultiFS: Automated Multi-Scenario Feature Selection in Deep Recommender Systems". In Proc. of WSDM. pp. 434–442.
21. ^Δ Song F, Chen B, Zhao X, Guo H, Tang R (2022). "Autoassign: Automatic shared embedding assignment in streaming recommendation". In: Proc. of ICDM. pp. 458–467.
22. ^Δ Zhao X, Xia L, Zhang L, Ding Z, Yin D, Tang J (2018). "Deep reinforcement learning for page-wise recommendations". In: Proc. of RecSys. pp. 95–103.
23. ^Δ Zhao X, Zhang L, Ding Z, Xia L, Tang J, Yin D (2018). "Recommendations with negative feedback via pairwise deep reinforcement learning". In: Proc. of KDD. pp. 1040–1048.
24. ^Δ Wang Y, Zhao X, Chen B, Liu Q, Guo H, Liu H, Wang Y, Zhang R, Tang R (2023). "PLATE: A prompt-enhanced paradigm for multi-scenario recommendations". In: Proc. of SIGIR. pp. 1498–1507.
25. ^Δ Lin W, Zhao X, Wang Y, Zhu Y, Wang W (2023). "Autodenoise: Automatic data instance denoising for recommendations". In Proc. of WWW. pp. 1003–1011.
26. ^a ^b ^c ^d ^e ^f Xu L, Skoularidou M, Cuesta-Infante A, Veeramachaneni K (2019). "Modeling Tabular data using Conditional GAN". In: Proc. of NeurIPS.
27. ^a ^b ^c Solorio AV, Dupriez O (2023). "Realtabformer: Generating realistic relational and tabular data using transformer". arXiv preprint arXiv:2302.02041.
28. ^a ^b Chatterjee S, Hazra D, Byun YC, Kim YW (2022). "Enhancement of image classification using transfer learning and GAN-based synthetic data augmentation". Mathematics. 2022: 1541.
29. ^Δ Borisov V, Leemann T, Seßler K, Haug J, Pawelczyk M, Kasneci G (2022). "Deep neural networks and tabular data: A survey". IEEE transactions on neural networks and learning systems. (2022).
30. ^a ^b Qian Z, Cebere BC, van der Schaar M (2023). "Synthcity: facilitating innovative use cases of synthetic data in different data modalities".
31. ^Δ Janizek JD, Sturmfels P, Lee S-I (2021). "Explaining Explanations: Axiomatic Feature Interactions for Deep Networks". Journal of Machine Learning Research. 2021: 1–54.
32. ^Δ Sundararajan M, Taly A, Yan Q (2016). "Gradients of counterfactuals". arXiv preprint arXiv:1611.02639. [arXiv:1611.02639](https://arxiv.org/abs/1611.02639).
33. ^Δ Janizek JD, Sturmfels P, Lee S-I (2021). "Explaining explanations: Axiomatic feature interactions for deep networks". Journal of Machine Learning Research. 2021: 1–54.

34. [△]Fu Z, Li X, Wu C, Wang Y, Dong K, Zhao X, Zhao M, Guo H, Tang R (2023). "A unified framework for multi-domain ctr prediction via large language models". *ACM Transactions on Information Systems*. (2023).
35. [△]Liang J, Zhao X, Li M, Zhang Z, Wang W, Liu H, Liu Z (2023). "Mmmlp: Multi-modal multilayer perceptron for sequential recommendations". In: *Proc. of WWW*. pp. 1109–1117.
36. [△]Yuan G, Yuan F, Li Y, Kong B, Li S, Chen L, Yang M, Yu C, Hu B, Li Z, et al. (2022). "Tenrec: A large-scale multipurpose benchmark dataset for recommender systems". *Proc. of NeurIPS*. 2022: 11480–11493.
37. [△]Lian J, Zhou X, Zhang F, Chen Z, Xie X, Sun G (2018). "xdeepfm: Combining explicit and implicit feature interactions for recommender systems". In: *Proc. of KDD*. pp. 1754–1763.
38. [△][‡]Kotelnikov A, Baranchuk D, Rubachev I, Babenko A (2023). "Tabddpm: Modelling tabular data with diffusion models". In *Proc. of ICML*. pp. 17564–17579.
39. [△][‡]Jordon J, Yoon J, Van Der Schaar M (2018). "PATE-GAN: Generating synthetic data with differential privacy guarantees". In *Proc. of ICLR*.
40. [△][‡]Yoon J, Drumright LN, Van Der Schaar M (2020). "Anonymization through data synthesis using generative adversarial networks (ads-gan)". *IEEE journal of biomedical and health informatics*. 2020: 2378–2388.
41. [△]Kirasich K, Smith T, Sadler B (2018). "Random forest vs logistic regression: binary classification for heterogeneous datasets". *SMU Data Science Review*. 2018: 9.
42. [△]Guo H, Tang R, Ye Y, Li Z, He X (2017). "DeepFM: a factorization-machine based neural network for CTR prediction". *arXiv preprint arXiv:1703.04247*.
43. [△]Grandini M, Bagli E, Visani G (2020). "Metrics for multi-class classification: an overview". *arXiv preprint arXiv:2008.05756*. [arXiv:2008.05756](https://arxiv.org/abs/2008.05756).
44. [△]AI@Meta (2024). "Llama 3 Model Card".
45. [△][‡]Patki N, Wedge R, Veeramachaneni K (2016). "The synthetic data vault". In: *2016 IEEE international conference on data science and advanced analytics (DSAA)*. pp. 399–410.
46. [△]Elvira V, Martino L (2021). "Advances in importance sampling". *arXiv preprint arXiv:2102.05407*.
47. [△]Qian Z, Davis R, van der Schaar M (2024). "Synthcity: a benchmark framework for diverse use cases of tabular synthetic data". *Proc. of NeurIPS*.
48. [△]Zhao X, Zheng X, Yang X, Liu X, Tang J (2020). "Jointly learning to recommend and advertise". In: *Proc. of KDD*. pp. 3319–3327.
49. [△]Liu Z, Tian J, Cai Q, Zhao X, Gao J, Liu S, Chen D, He T, Zheng D, Jiang P, et al. (2023). "Multi-task recommendations with reinforcement learning". In *Proc. of WWW*. pp. 1273–1282.
50. [△]Lin W, Zhao X, Wang Y, Xu T, Wu X (2022). "AdaFS: Adaptive feature selection in deep recommender system". In: *Proc. of KDD*. pp. 3309–3317.
51. [△]Li X, Qiu Z, Zhao X, Wang Z, Zhang Y, Xing C, Wu X (2022). "Gromov-wasserstein guided representation learning for cross-domain recommendation". In: *Proc. of CIKM*. 2022:1199–1208.
52. [△]Zheng Z, Qiu Z, Xiong H, Wu X, Xu T, Chen E, Zhao X (2022). "Ddr: Dialogue based doctor recommendation for online medical service". In *Proc. of KDD*. pp. 4592–4600.

53. [△]Xu D, Zhang Z, Lin Z, Wu X, Zhu Z, Xu T, Zhao X, Zheng Y, Chen E (2024). "Multi-perspective Improvement of Knowledge Graph Completion with Large Language Models". In *Proc. of COLING*. pp. 11956–11968.
54. [△]Zhao K, Zou L, Zhao X, Wang M, Yin D (2023). "User retention-oriented recommendation with decision transformer". In: *Proc. of WWW*. pp. 1141–1149.
55. [△]Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017). "Attention is all you need". *Proc. of NeurIPS* (2017).
56. [△]Zhao WX, Zhou K, Li J, Tang T, Wang X, Hou Y, Min Y, Zhang B, Zhang J, Dong Z, et al. (2023). "A survey of large language models". *arXiv preprint arXiv:2303.18223*.
57. [△]Minaee S, Mikolov T, Nikzad N, Chenaghlu M, Socher R, Amatriain X, Gao J (2024). "Large language models: A survey". *arXiv preprint arXiv:2402.06196*.
58. [△]Chang Y, Wang X, Wang J, Wu Y, Yang L, Zhu K, Chen H, Yi X, Wang C, Wang Y, et al. (2024). "A survey on evaluation of large language models". *ACM Transactions on Intelligent Systems and Technology*. 2024: 1–45.
59. [△]Liu Q, Wu X, Zhao X, Zhu Y, Zhang Z, Tian F, Zheng Y (2024). "Large Language Model Distilling Medication Recommendation Model". *CoRR*. (2024).
60. [△]Liu Q, Wu X, Zhao X, Wang Y, Zhang Z, Tian F, Zheng Y (2024). "Large Language Models Enhanced Sequential Recommendation for Long-tail User and Item". *arXiv preprint arXiv:2405.20646* (2024).
61. [△]Zhang C, Zhang H, Wu S, Wu D, Xu T, Gao Y, Hu Y, Chen E (2024). "NoteLLM-2: Multimodal Large Representation Models for Recommendation". *arXiv preprint arXiv:2405.16789* (2024).
62. [△]Li Y, Wen H, Wang W, Li X, Yuan Y, Liu G, Liu J, Xu W, Wang X, Sun Y, et al. (2024). "Personal llm agents: Insights and survey about the capability, efficiency and security". *arXiv preprint arXiv:2401.05459*.

Declarations

Funding: This research was partially supported by Research Impact Fund (No.R1015-23), Collaborative Research Fund (No.C1043-24GF), APRC - CityU New Research Initiatives (No.9610565, Start-up Grant for New Faculty of CityU), Hong Kong ITC Innovation and Technology Fund Midstream Research Programme for Universities Project (No.ITS/034/22MS), and Huawei (Huawei Innovation Research Program).

Potential competing interests: No potential competing interests to declare.