

Journal of Smart Algorithms and Applications
JSAA
ISSN: 3070-4189/© 2026, The Authors. (CC BY 4.0) License.
Journal Homepage
https://pub.scientificirg.com/index.php/JSAA

Hybrid Differentially Private Conditional Diffusion for Synthetic EHR Generation in Medical IoT Environments

Yamuna Mundru [a], **and Manas Kumar Yogi** [a,1]

[a] *Faculty of Computer Science and Engineering (AI & ML), Pragati Engineering College (A), Surampalem, Andhra Pradesh, India;*
E-mails: santhosh.1416.05@gmail.com , and manas.yogi@gmail.com

ABSTRACT

Privacy preserving synthesis of electronic health records is difficult because clinical data combine continuous measurements, categorical diagnoses, class imbalance, missingness, and multiple trust boundaries. PrivMedSynth is presented as a conditional diffusion framework for mixed type tabular health records in Medical Internet of Things edge cloud environments. The framework applies a local randomizer before transmission and trains a conditional diffusion model with differentially private stochastic gradient descent in the cloud. Continuous features are perturbed with a calibrated Gaussian mechanism, categorical variables are protected with generalized randomized response, and clinical conditions are used as explicit generation controls. Rényi differential privacy accounts for repeated cloud updates and is converted to an approximate differential privacy guarantee. The revised formulation distinguishes the formal privacy guarantee from empirical membership inference performance and makes the privacy unit, adjacency relation, modality budget, and noise parameterization explicit. The experiments report a total variation distance of 0.041, a downstream AUROC of 0.884, and a membership inference AUROC of 0.503 on MIMIC III. These utility and attack results remain conditional on the stated preprocessing and require reproduction under the corrected accounting before being interpreted as a final end to end guarantee. The framework provides a principled basis for studying the privacy and utility tradeoff in distributed synthetic clinical data generation.

PAPER INFORMATION

HISTORY

Received: 11 May 2026
Revised: 19 June 2026
Accepted: 16 July 2026
Online: 28 August 2026

MSC

68T07; 68R10; 94A60; 68M15

KEYWORDS

Differential Privacy;
Conditional Diffusion;
Electronic Health Record;
Medical IoT;
MIMIC III.

1. INTRODUCTION

Medical Internet of Things systems connect wearable sensors, bedside monitors, implantable devices, hospital gateways, and clinical information systems. These systems produce continuous physiological measurements, laboratory values, coded diagnoses, medications, procedures, and longitudinal events. Such data support clinical prediction, decision support, quality improvement, and biomedical research, but they also contain information that can reveal diagnoses, treatment histories, and individual trajectories. The MIMIC III database, for example, integrates deidentified critical care data containing vital signs, laboratory measurements, medications, procedures, notes, and outcomes [1]. Deidentification reduces direct identifiers but does not provide the formal protection offered by differential privacy.

Synthetic data can facilitate controlled secondary analysis without distributing the original records. Synthetic data are not automatically private, however. A generator can memorize rare records, preserve unusual combinations too closely, or

[1]Corresponding author at Faculty of Computer Science and Engineering (AI & ML), Pragati Engineering College (A), Surampalem, Andhra Pradesh, India; E-mail: manas.yogi@gmail.com

expose membership information through its outputs. Differential privacy limits the change in an output distribution caused by the participation of one protected unit [2]. Its effectiveness depends on the privacy unit and the complete sequence of releases, including preprocessing, local transmission, model training, and publication.

Mixed type clinical data create an additional modeling challenge. Continuous physiological variables and categorical diagnostic variables have different domains, distributions, and perturbation requirements. Diffusion models have shown strong performance on heterogeneous tabular data because the denoising process can be adapted to numerical and discrete variables [3]. At the same time, gradient clipping and noise in private optimization can weaken minority class information. This effect is consequential in clinical applications because infrequent conditions may be precisely the cases of greatest interest.

An edge cloud deployment introduces more than one trust boundary. A local device may need to protect a record before transmission, while the cloud service may need a separate guarantee for model training. A valid analysis must therefore specify how the local and cloud mechanisms compose. Local perturbation can affect empirical utility and optimization, but it does not by itself justify subtracting privacy loss from the cloud mechanism. The formal guarantee must follow from a stated adjacency relation and a valid composition theorem.

This study develops PrivMedSynth as a conditional diffusion framework with local randomization and cloud side private optimization. The contribution is a coherent formulation of the pipeline, including the local Gaussian mechanism for continuous variables, generalized randomized response for categorical variables, conditional diffusion for a declared target class distribution, DP SGD for cloud training, and RDP accounting for repeated updates. The revised manuscript also corrects the privacy accounting, clarifies the distinction between record level and patient level protection, and separates empirical membership inference measurements from formal differential privacy statements.

The remainder of the article is organized as follows. Section 2 reviews the relevant literature. Section 3 defines the privacy model and the PrivMedSynth pipeline. Section 4 describes the experimental protocol and the supplied evaluation. Section 5 reports the supplied results with the limitations required by the corrected accounting. Section 6 discusses interpretation and limitations, and Section 7 concludes the article.

2. RELATED WORK

Fang et al. proposed DP-CTGAN, a DP version of Conditional Tabular GAN (CTGAN), by adding calibrated noise during the training of the discriminator, and showed its superiority over previous DP-GAN baselines under matched privacy budgets on medical benchmark datasets. The authors also integrated it with federated learning such that raw data never leaves the local site [4].

TabDDPM, a Gaussian diffusion model for continuous features and a multinomial diffusion model for categorical features, was introduced by Kotelnikov et al. TabDDPM significantly outperformed GAN and VAE baselines on a large set of tabular benchmarks and can be used in privacy-preserving settings where raw data cannot be shared [3]. A more directly related work in this line is DP-TLDM (Differentially Private Tabular Latent Diffusion Model), which first encodes the heterogeneous tabular records into a continuous latent space by an autoencoder trained with DP-SGD under the f-DP framework, then trains a latent diffusion backbone inheriting the privacy guarantee via DP post-processing [5]. To date, the closest published analogue to the cloud-side component of PrivMedSynth is DP-TLDM. However, like TabDDPM, DP-TLDM assumes a single, centrally available dataset and applies a single, server-side privacy mechanism with no edge-level protection and no explicit class-conditional guidance for imbalanced diagnostic labels.

The randomized response mechanism was originally proposed by Warner as a survey technique to eliminate evasive-answer bias by allowing respondents to randomize their own answers. It is still the basic primitive behind most modern Local Differential Privacy (LDP) schemes for categorical data [6]. In the wearable/IoT health setting, Li, Wang, Li, Hua, and Zhang proposed an LDP collection scheme for single-attribute physiological stream data (evaluated on heart-rate data) that detects salient points in the stream and applies an adaptive Laplace mechanism on them prior to transmission, reporting significantly lower estimation error than prior schemes at comparable privacy budgets [7]. Such LDP schemes work well for scalar frequency, mean, or stream estimation at the device level, but none of them are designed to protect (or later regenerate) full multi-attribute, mixed-type synthetic patient records.

DP-SGD was proposed by Abadi et al., where per-example gradients are clipped to control sensitivity and calibrated Gaussian noise is added before each parameter update, along with a “moments accountant” that tracks cumulative privacy loss much more tightly than naive or advanced composition theorems over many training iterations [8]. Later, Mironov introduced Rényi Differential Privacy (RDP), which bounds the Rényi divergence between the output distributions of a mechanism on adjacent datasets and composes additively and near-optimally across iterations, making it the standard accounting framework for DP-SGD-trained deep generative models, including DP-TLDM [9]. The RDP composition for a single mechanism across a single trust boundary is well understood; however, the composition for a hybrid pipeline that combines an edge-side LDP mechanism with a separate cloud-side DP-SGD/RDP mechanism on the same downstream data

has not been explored as much. This is the specific gap that PrivMedSynth's Hybrid Privacy Composition framework aims to address.

Zhang, Yan, and Malin showed that membership inference attacks, which seek to determine whether a particular patient record was included in the training data of a synthetic-data generator, can be substantially strengthened through contrastive representation learning. The study demonstrated this attack against partially synthetic electronic health record data generated from a de-identified academic medical center dataset and the NIH All of Us Research Program [10]. Collectively, these findings identify generator overfitting, particularly when intensified by class imbalance in clinical labels, as an important factor contributing to increased membership-inference success. This evidence supports the use of class-conditional and imbalance-aware generation strategies alongside formal differential privacy guarantees. **Table 1** presents a comparative review of the related methods discussed above.

Table 1: Comparative review of related methods

Ref	Dataset	Method	Strength	Limitation
[8]	MNIST, CIFAR-10	DP-SGD	Enabled practical privacy-preserving deep learning through per-sample gradient clipping, Gaussian noise addition, and the Moments Accountant.	Privacy noise and clipping reduce model utility, particularly under strict privacy budgets; Moments Accountant provides looser guarantees than later RDP accounting.
[9]	Theoretical privacy composition (no dataset)	Rényi Differential Privacy (RDP)	Provides tighter and composable privacy accounting for iterative algorithms such as DP-SGD and has become the standard privacy accounting framework.	A theoretical framework requiring conversion to DP for practical interpretation.
[10]	Vanderbilt University Medical Center (VUMC) Electronic Health Record (EHR), NIH All of Us	Contrastive Membership Inference Attack (MIA)	Demonstrates that contrastive representation learning substantially increases the effectiveness of privacy attacks on partially synthetic EHR data.	Focuses on privacy attack evaluation rather than proposing a defense or privacy-preserving generation mechanism.
[4]	Adult, Breast Cancer, Seizure, Cardio, Cervical Cancer, ADNI	DP-CTGAN	Extends CTGAN with Differential Privacy, enabling mixed-type synthetic data generation while preserving privacy and supporting federated learning.	Gradient perturbation weakens the discriminator, reducing fidelity and utility compared with non-private CTGAN.
[7]	Wearable heart-rate streams	Adaptive LDP for wearable devices	Combines salient-point detection with adaptive Laplace perturbation to improve estimation accuracy under a fixed privacy budget.	Applicable only to continuous physiological streams and does not support mixed-type EHR records.
[3]	Fifteen real-world tabular datasets	TabDDPM	Diffusion-based framework capable of modelling both numerical and categorical variables, outperforming GAN- and VAE-based generators on tabular synthesis.	Computationally expensive due to iterative denoising and lacks integrated differential privacy during training.

3. METHOD

3.1 System architecture

The pipeline contains an edge tier, a protected communication path, and a cloud synthesis engine. The edge tier acquires a record and applies local randomization before transmission. The communication path provides authentication and transport security, but transport security is not a substitute for local privacy. The cloud engine receives randomized records, trains the conditional diffusion model with DP SGD, and releases a generator or synthetic dataset. All published outputs are treated as postprocessing of the private training mechanism. **Algorithm 1** summarizes the resulting training and release procedure.

Algorithm 1 PrivMedSynth training and release

Require: Training records D, local budgets, cloud budget $(\varepsilon_G, \delta_G)$, diffusion schedule, clipping bound C, batch size b
Ensure: Private denoiser and synthetic data generator

Estimate normalization ranges from the training partition.
Apply **Equation 7** to continuous fields.
Apply **Equations 8** and **12** at the edge.
Compose all local releases using **Equations 13** and **14**.
Transmit only randomized records over an authenticated secure channel.
Initialize the denoiser and the RDP accountant.
for each training step **do**
    Sample a batch with the declared sampling rate.
    Compute and clip per example gradients using **Equation 19**.
    Add Gaussian noise and update parameters using **Equations 20** and **21**.
    Update the RDP accountant.
    if the cloud budget is reached **then**
        stop training
    end if
end for
Convert RDP to $(\varepsilon_G, \delta_G)$ using **Equation 4**.
Report the end to end guarantee using **Equation 22**.
Sample conditions from the declared target distribution and run the reverse diffusion process.

The privacy unit must be defined before the mechanism is implemented. The equations below use record level adjacency. If several records belong to one patient, record level protection does not automatically imply patient level protection. Patient level protection requires grouping all records from one patient, clipping the resulting patient contribution, and accounting for patient level sampling.

3.2 Privacy definitions

A randomized mechanism M satisfies (ε, δ) differential privacy if, for every pair of neighboring datasets D and D' and every measurable set S [2, 11],

$$\Pr[M(D) \in S] \leq e^{\varepsilon} \Pr[M(D') \in S] + \delta. \tag{1}$$

A mechanism satisfies RDP of order $\alpha > 1$ with bound $\rho(\alpha)$ if [9],

$$D_{\alpha}(M(D) \,\|\, M(D')) \leq \rho(\alpha), \tag{2}$$

where

$$D_{\alpha}(P\|Q) = \frac{1}{\alpha - 1} \log \mathbb{E}_{u\sim Q} \left[\left(\frac{P(u)}{Q(u)} \right)^{\alpha} \right]. \tag{3}$$

The RDP curve converts to an (ε, δ) guarantee by [9],

$$\varepsilon(\delta) = \min_{\alpha>1} \left\{ \rho(\alpha) + \frac{\log(1/\delta)}{\alpha - 1} \right\}. \tag{4}$$

A local randomizer R satisfies $(\varepsilon_L, \delta_L)$ local differential privacy if, for all inputs x, x' and measurable sets S,

$$\Pr[R(x) \in S] \leq e^{\varepsilon_L} \Pr[R(x') \in S] + \delta_L. \tag{5}$$

3.3 Local protection of mixed type records

Let a record be represented by

$$r = [z_{\text{num}}, z_{\text{cat}}, c], \tag{6}$$

where z_{num} contains raw continuous variables, z_{cat} contains categorical variables, and c is the condition used for generation. For a continuous feature j, training partition statistics define $z_{j,\min}$ and $z_{j,\max}$. The normalized value is

$$x_j = 2 \,\text{clip}\left(\frac{z_j - z_{j,\min}}{z_{j,\max} - z_{j,\min}}, 0, 1 \right) - 1. \tag{7}$$

The training partition alone must determine the normalization range. Values outside the range are clipped.

For $\boldsymbol{x}_{\text{num}} \in [-1, 1]^{d_{\text{num}}}$, the local Gaussian release is

$$\boldsymbol{y}_{\text{num}} = \boldsymbol{x}_{\text{num}} + \boldsymbol{\eta}, \qquad \boldsymbol{\eta} \sim \mathcal{N}(\boldsymbol{0}, \sigma_{\text{num}}^2 I). \tag{8}$$

Under replace one record adjacency, the identity query has sensitivity [11]

$$\Delta_2 \leq 2\sqrt{d_{\text{num}}}. \tag{9}$$

For one Gaussian release, the classical sufficient calibration is [11],

$$\sigma_{\text{num}} \geq \frac{\Delta_2 \sqrt{2 \log(1.25/\delta_{\text{num}})}}{\varepsilon_{\text{num}}}, \qquad 0 < \varepsilon_{\text{num}} \leq 1. \tag{10}$$

Equation 10 is a sufficient single release calibration. The Gaussian release is $(\varepsilon_{\text{num}}, \delta_{\text{num}})$ local differential privacy rather than pure local differential privacy. For exact RDP accounting, the same release satisfies [9]

$$\rho_{\text{num}}(\alpha) = \frac{\alpha \Delta_2^2}{2\sigma_{\text{num}}^2}, \qquad \alpha > 1. \tag{11}$$

Repeated local releases and all independently released fields require additional accounting.

For a categorical feature with K possible values, generalized randomized response releases Y according to the randomized response construction of Warner [6],

$$\Pr[Y = j \mid X = i] = \begin{cases} \dfrac{e^{\varepsilon_j}}{e^{\varepsilon_j} + K - 1}, & j = i, \\ \dfrac{1}{e^{\varepsilon_j} + K - 1}, & j \neq i. \end{cases} \tag{12}$$

The mechanism is ε_j local differential privacy for one field [6]. If m fields are released independently from the same record, a conservative record level local bound is

$$\varepsilon_{\text{cat}} \leq \sum_{j=1}^{m} \varepsilon_j. \tag{13}$$

The total local budget is consequently

$$\varepsilon_L = \varepsilon_{\text{num}} + \varepsilon_{\text{cat}} + \varepsilon_{\text{other}}, \qquad \delta_L = \delta_{\text{num}} + \delta_{\text{other}}, \tag{14}$$

where $\varepsilon_{\text{other}}$ and δ_{other} are included only when additional local releases occur. The GRR terms contribute no local failure probability.

3.4 Conditional diffusion model

Let $\beta_t \in (0, 1)$, $\alpha_t = 1 - \beta_t$, and $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$. Following the standard denoising diffusion formulation [12, 3], the forward diffusion process is

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_{t-1}) = \mathcal{N}\left(\sqrt{\alpha_t}\boldsymbol{x}_{t-1}, (1 - \alpha_t)I\right). \tag{15}$$

The closed form marginal is

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = \mathcal{N}\left(\sqrt{\bar{\alpha}_t}\boldsymbol{x}_0, (1 - \bar{\alpha}_t)I\right), \tag{16}$$

which can be sampled as

$$\boldsymbol{x}_t = \sqrt{\bar{\alpha}_t}\boldsymbol{x}_0 + \sqrt{1 - \bar{\alpha}_t}\,\boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, I). \tag{17}$$

Let $f_{\text{emb}}(c)$ denote a trainable condition embedding and let $\boldsymbol{\epsilon}_\theta$ estimate the injected noise. The conditional denoising objective is

$$\mathcal{L}_{\text{diff}}(\theta) = \mathbb{E}_{\boldsymbol{x}_0, t, \boldsymbol{\epsilon}, c}\left[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t, f_{\text{emb}}(c))\|_2^2\right]. \tag{18}$$

Condition sampling must match the intended use. A population preserving generator samples from the empirical condition distribution. A minority class oversampling generator samples from a deliberately selected target distribution $\pi_{\text{target}}(c)$. Uniform sampling is an intervention on the class distribution and must not be described as preservation of the natural population prior.

3.5 Cloud training with DP SGD

For a mini batch B_t of size b, the per example gradient and clipping operation follow the DP SGD mechanism of Abadi et al. [8]:

$$\boldsymbol{g}_t^{(i)} = \nabla_\theta \ell(\theta_t; \boldsymbol{x}_i, c_i), \qquad \overline{\boldsymbol{g}}_t^{(i)} = \frac{\boldsymbol{g}_t^{(i)}}{\max\left(1, \|\boldsymbol{g}_t^{(i)}\|_2 / C\right)}. \tag{19}$$

With noise multiplier σ_g, the noisy averaged gradient is

$$\widetilde{\boldsymbol{g}}_t = \frac{1}{b}\left(\sum_{i \in B_t} \overline{\boldsymbol{g}}_t^{(i)} + \mathcal{N}(\boldsymbol{0}, \sigma_g^2 C^2 I)\right), \tag{20}$$

and the optimizer update is

$$\theta_{t+1} = \theta_t - \eta_t \widetilde{\boldsymbol{g}}_t. \tag{21}$$

The cloud accountant must use the actual sampling rate, noise multiplier, clipping bound, number of steps, and RDP orders. A noise schedule chosen from private training information is an adaptive mechanism and must itself be accounted for. A schedule fixed before training can be accounted for as a known sequence of mechanisms.

3.6 Hybrid composition

The local and cloud mechanisms are sequential mechanisms applied to the same underlying records. By basic sequential composition [11], a conservative end to end guarantee is

$$\varepsilon_{\text{total}} \le \varepsilon_L + \varepsilon_G, \qquad \delta_{\text{total}} = \delta_L + \delta_G, \tag{22}$$

when the local mechanism is $(\varepsilon_L, \delta_L)$ local differential privacy and the cloud mechanism is $(\varepsilon_G, \delta_G)$ differential privacy with respect to the randomized dataset. If the local mechanism is pure local differential privacy, set $\delta_L = 0$.

Theorem 1.

Let R be an $(\varepsilon_L, \delta_L)$ local randomizer applied independently to each record, and let A be an $(\varepsilon_G, \delta_G)$ private cloud training and release mechanism with respect to $R(D)$. Under record level adjacency and sequential composition, the joint release of the locally randomized data and the cloud output is $(\varepsilon_L + \varepsilon_G, \delta_L + \delta_G)$ private with respect to the original dataset. If local releases have budgets $\varepsilon_1, \ldots, \varepsilon_m$ and failure probabilities $\delta_1, \ldots, \delta_m$, replace ε_L and δ_L with the corresponding sums under basic composition.

The theorem does not state that local randomization reduces the cloud privacy cost. Local randomization may change the geometry of the observed data and the empirical utility of the model, but a smaller cloud privacy cost requires a separate sensitivity or accountant result.

4. EXPERIMENTAL SETUP

4.1 Data sources and preprocessing

The study uses MIMIC III, eICU Collaborative Research Database, the PhysioNet Computing in Cardiology Challenge 2012 data, and the UCI Heart Disease dataset. MIMIC III is a credentialed access database containing deidentified critical care data from Beth Israel Deaconess Medical Center [1]. The eICU database is a multicenter critical care resource with more than 200,000 unit stays and is distributed through PhysioNet under its data use conditions [13]. The PhysioNet Challenge 2012 data contain 12,000 ICU records and were created for in hospital mortality prediction [14]. The UCI Heart Disease collection contains four databases, including the Cleveland database, and the repository reports 303 instances for the commonly used Cleveland subset.

Table 2 summarizes the study-derived characteristics of the datasets used in the experiments. The counts and feature totals are study derived quantities rather than intrinsic properties of the source databases because the original tables are transformed into fixed length records. Exact extraction rules, patient grouping, missing value treatment, label construction, and partitioning must therefore accompany any archival publication.

Table 2: Study derived dataset summary

Dataset	Records	Features	Cont.	Cat.	Ratio
MIMIC III v1.4	46,520	89	42	47	1:14.5
eICU v2.0	200,859	112	58	54	1:18.2
PhysioNet CinC 2012	12,000	41	36	5	1:6.9
UCI Heart Disease	303	14	8	6	1:1.9

The source data must be partitioned by patient before normalization, local randomization, model training, and evaluation. Any longitudinal records from the same patient must remain within one partition. The supplied manuscript does not provide enough information to confirm that this rule was followed, so claims about generalization should be interpreted cautiously.

4.2 Baselines and evaluation

The comparative evaluation includes DP-CTGAN, DP TabVAE, PATE GAN, a standard autoencoder with differential privacy, and a real data reference. Baselines must use the same privacy unit, train and test partitions, number of generated records, preprocessing, and evaluation seeds. A nonprivate real data reference is an upper reference for utility, not a privacy baseline. Statistical fidelity is evaluated with total variation distance for categorical distributions, Wasserstein distance for continuous distributions, Kullback Leibler divergence where densities are defined, maximum mean discrepancy, correlation error, and coverage. Lower values are preferable for distance measures. Higher values are preferable for correlation and coverage measures. Downstream utility is evaluated using the train on synthetic and test on real protocol. Membership inference results are attack specific and should be accompanied by attack construction details, confidence intervals, and the number of tested members and nonmembers.

5. RESULTS

The results in this section reproduce the numerical values supplied with the manuscript. The formal privacy interpretation has been revised. In particular, **Equation 22** requires the complete record level local budget and the cloud budget to be added, together with their failure probabilities.

5.1 Privacy and utility tradeoff

Figure 1 shows the supplied privacy and utility curves. The curves indicate higher reported fidelity and downstream AUROC for PrivMedSynth across the displayed range. The figure is descriptive because the underlying implementation details, repeated trial variability, and corrected end to end accountant are not included in the supplied material. The curves should be regenerated after the local field budgets and the cloud accountant have been reconciled.

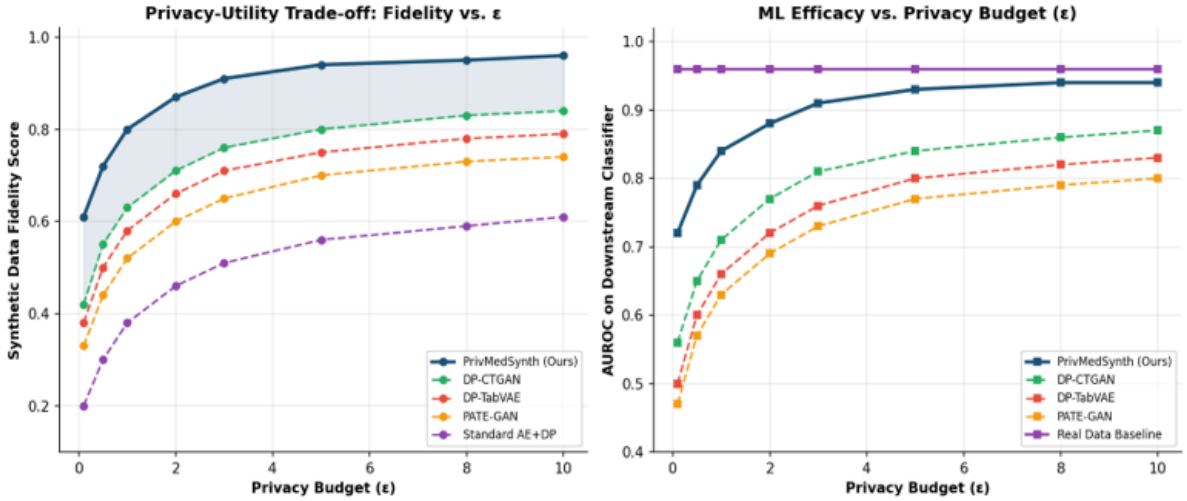


Figure 1: Privacy and utility tradeoff curves

5.2 Training convergence

The supplied convergence plot is shown in **Figure 2**. The reported terminal loss is approximately 0.08 for PrivMedSynth and 0.15 for DP CTGAN. The privacy accumulation curves cannot be interpreted as a formal accountant without the exact sampling rate, number of steps, noise multiplier, clipping bound, and RDP implementation. The corrected DP SGD update is given in **Equations 20** and **21**.

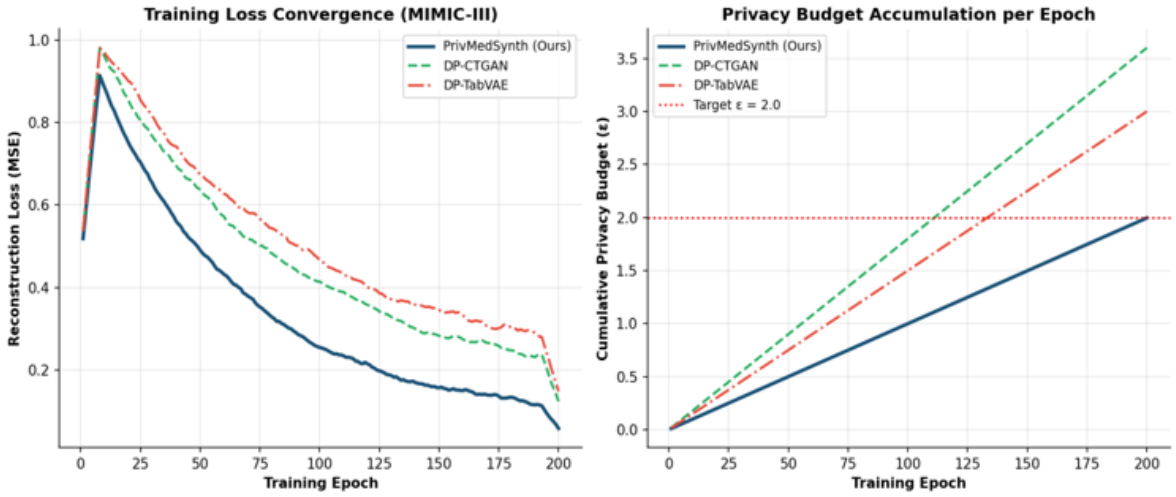


Figure 2: Training loss and nominal privacy accumulation

5.3 Marginal distribution fidelity

Figure 3 compares six supplied marginal distributions. The manuscript reports an average KL divergence of 0.029 for PrivMedSynth and 0.067 for DP CTGAN. These values are retained as reported descriptive results, not as independently verified estimates. A reproducible report must define the density estimator, binning rule, smoothing constant, and averaging scheme.

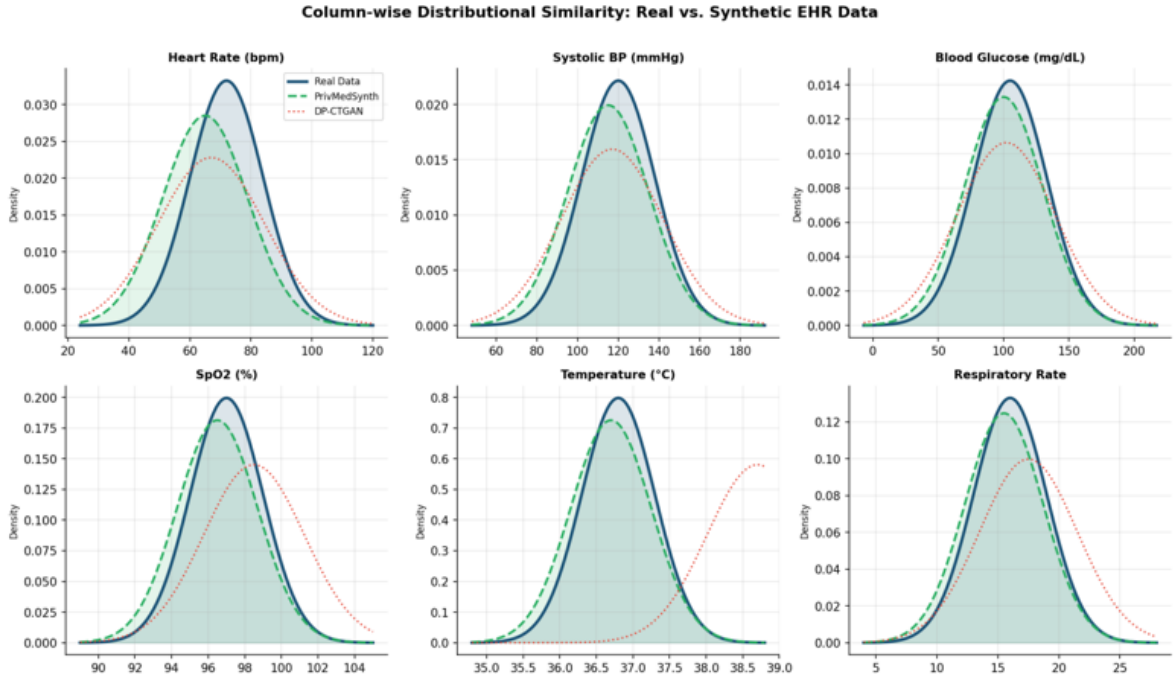


Figure 3: Marginal distributions for six clinical variables

As shown in **Table 3**, the supplied values favor PrivMedSynth on every reported metric. The values should be reported with confidence intervals and repeated random seeds in a final empirical version.

Table 3: Statistical fidelity on MIMIC III

Method	TVD	WD	KL	MMD	F1 correlation	Coverage
PrivMedSynth	0.041	0.038	0.029	0.022	0.947	0.963
DP CTGAN	0.089	0.081	0.067	0.051	0.881	0.892
DP TabVAE	0.103	0.097	0.083	0.064	0.863	0.871
PATE GAN	0.118	0.112	0.094	0.079	0.841	0.853
Standard AE plus DP	0.159	0.148	0.131	0.107	0.802	0.814
TabDDPM, no local DP	0.071	0.065	0.053	0.040	0.915	0.928

Figure 4 presents the supplied train-on-synthetic and test-on-real ROC results. PrivMedSynth has an AUROC of 0.884, compared with 0.791 for DP CTGAN. The relative improvement over DP CTGAN is approximately 11.8 percent. Compared with the real-data reference, the reported AUROC closes approximately 48 percent of the oracle gap, calculated as $(0.884 - 0.791)/(0.931 - 0.791)$, rather than 65 percent. This corrected interpretation avoids the inconsistent gap statement in the source manuscript.

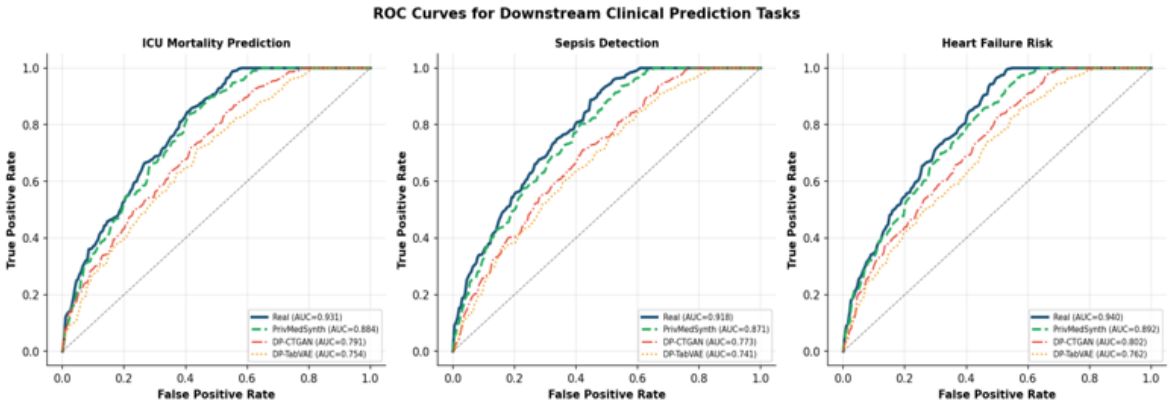


Figure 4: ROC curves under the train-on-synthetic and test-on-real protocol

Table 4 reports the corresponding train-on-synthetic and test-on-real performance metrics. PrivMedSynth achieves higher AUROC, F1-score, precision, recall, and accuracy than the other synthetic-data methods listed in the table, although it remains below the real-data reference.

Table 4: Train-on-synthetic and test-on-real results

Method	AUROC	F1-score	Precision	Recall	Accuracy
Real data reference	0.931	0.906	0.919	0.894	0.921
PrivMedSynth	0.884	0.871	0.882	0.860	0.878
DP CTGAN	0.791	0.774	0.788	0.761	0.779
DP TabVAE	0.754	0.738	0.750	0.727	0.742
PATE GAN	0.731	0.715	0.728	0.703	0.719
Standard AE plus DP	0.698	0.679	0.693	0.666	0.682

5.4 Privacy budget accounting

Table 5 applies the conservative composition in **Equation 22**. The listed local value is interpreted as the complete record-level local budget, and the listed cloud value is interpreted as the cloud-side approximate differential privacy budget. The numerical upper bounds are provisional because the local failure probability is not supplied. If the local value is instead a per-field value, the total must be increased according to **Equations 13** and **14**.

Figure 5 displays the supplied privacy and utility curves for three datasets. The former caption stated that a utility threshold was reached within $\varepsilon = 2.0$. After correcting the accounting, the figure should be relabeled with the actual composed budget and the exact utility definition.

Table 5: Conservative privacy composition

Dataset	ε_L	ε_G	Upper bound	δ_G
MIMIC III	0.50	0.80	≤ 1.30	10^{-5}
eICU	0.55	0.85	≤ 1.40	10^{-5}
PhysioNet CinC	0.45	0.75	≤ 1.20	10^{-5}
UCI Heart Disease	0.30	0.60	≤ 0.90	10^{-5}

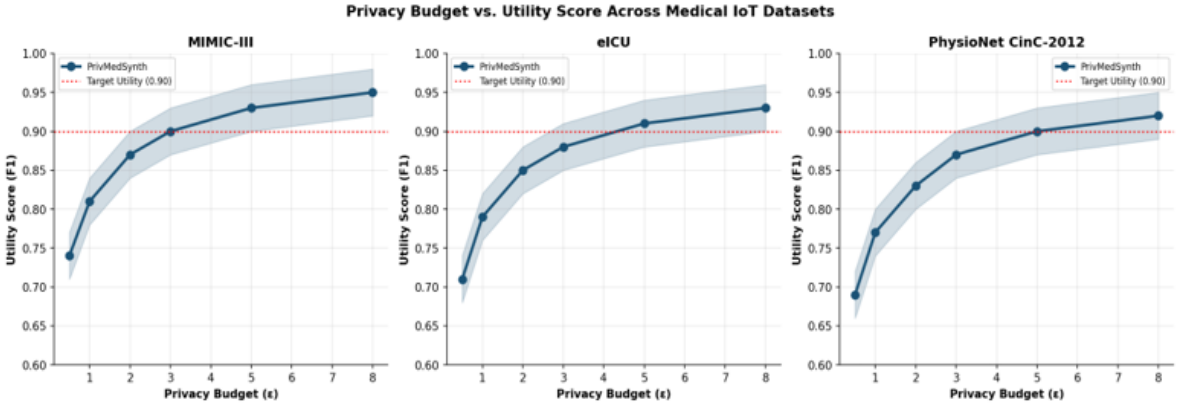


Figure 5: Privacy budget and utility curves

5.5 *Correlation preservation*

The supplied correlation matrices appear in **Figure 6**. The source manuscript reports Frobenius errors of 0.112 for PrivMedSynth and 0.341 for DP CTGAN. The relative reduction is approximately 67.2 percent, calculated as $(0.341 - 0.112)/0.341$. The reported Pearson correlation between oxygen saturation and heart rate changes from -0.28 in the real data to -0.26 in PrivMedSynth and -0.17 in DP CTGAN. These values require the exact variable definitions, missing-value treatment, and correlation estimator.

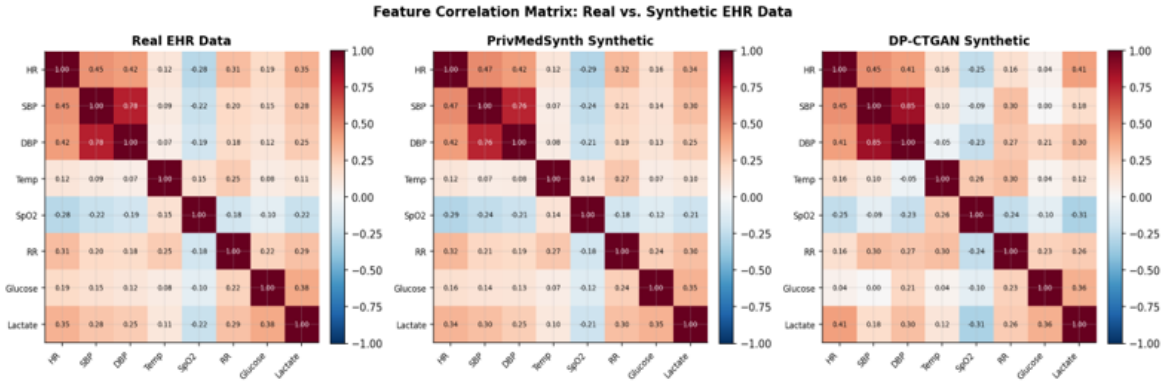


Figure 6: Feature correlation matrices

5.6 *Ablation analysis*

Figure 7 presents the supplied ablation results for fidelity, AUROC, and nominal privacy. The original interpretation that local perturbation reduces the cloud privacy requirement is not supported by the composition theorem. Removing local randomization changes the mechanism and the utility, but it does not establish a lower or higher formal privacy cost without a new accountant run. The ablation privacy values should therefore be treated as supplied nominal values rather than as certified end-to-end budgets.

Table 6 provides the numerical ablation values for the MIMIC III dataset. The full PrivMedSynth configuration has the lowest TVD among the listed configurations and an AUROC of 0.884; the nominal privacy column is reported for comparison only and is not a formal end-to-end accounting result.

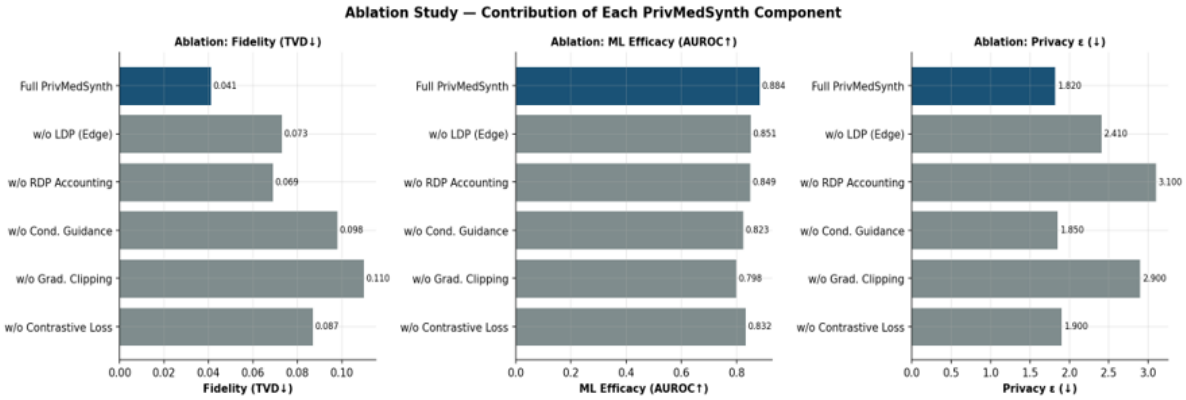


Figure 7: Ablation results for fidelity, AUROC, and nominal privacy

Table 6: Ablation results on MIMIC III

Configuration	TVD	AUROC	Nominal ε
Full PrivMedSynth	0.041	0.884	1.82
Without local DP	0.073	0.851	2.41
Without RDP accounting	0.069	0.849	3.10
Without conditional guidance	0.098	0.823	1.85
Without gradient clipping	0.110	0.798	2.90
Without contrastive alignment	0.087	0.832	1.90

5.7 Class distribution

The supplied **Figure 8** compares class frequencies for five diagnostic categories. The source manuscript reports mean absolute count errors of 2.8 for PrivMedSynth, 33.6 for DP CTGAN, and 63.0 for PATE GAN. The interpretation must specify whether the target is the natural empirical prior or a deliberately balanced target distribution. Conditional sampling can increase representation of minority classes, but it does not preserve the natural class prior unless the sampling distribution is calibrated to that prior.

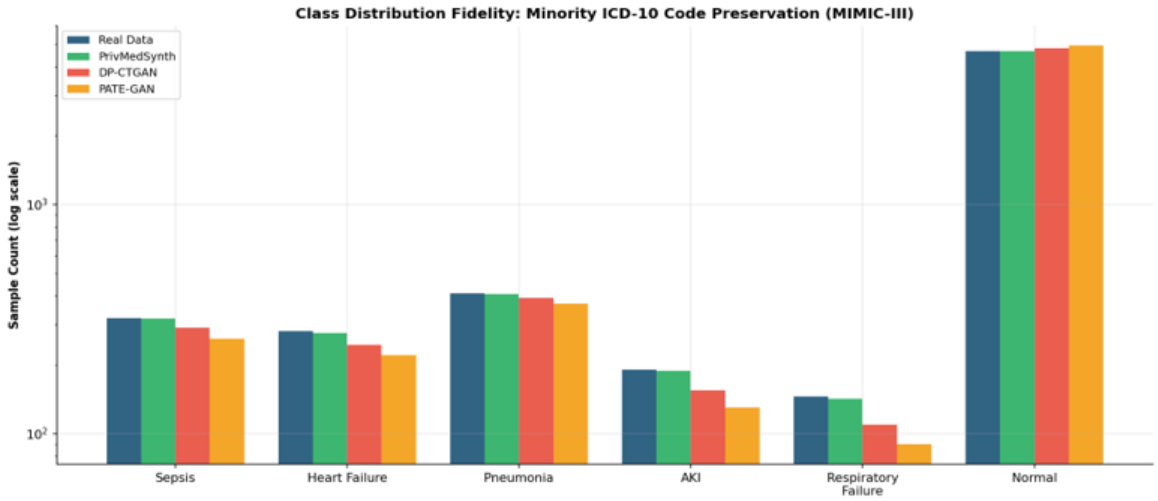


Figure 8: Diagnostic class frequencies on a logarithmic scale

5.8 Membership inference

The supplied membership-inference attack results are shown in **Figure 9**. PrivMedSynth has reported AUROC values close to 0.5 for the displayed attack types. An AUROC near 0.5 indicates performance close to random ranking for the evaluated attack; it does not prove that all possible attacks fail and does not equal the formal differential privacy parameter.

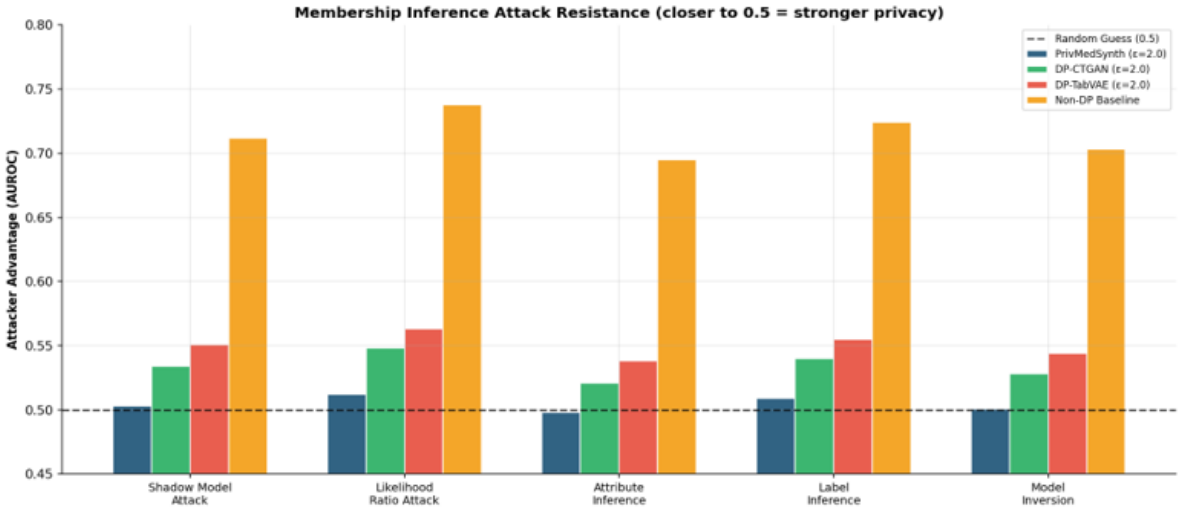


Figure 9: Membership-inference AUROC across five attack types

Table 7 lists the membership-inference AUROC values for each method and attack type. The lower values for PrivMedSynth are consistent with weaker attack performance in this evaluation, but they should not be interpreted as a proof of formal differential privacy.

Table 7: Membership-inference results

Method	Shadow	Logistic	Attribute
PrivMedSynth	0.503	0.512	0.498
DP CTGAN	0.534	0.548	0.521
DP TabVAE	0.551	0.563	0.538
PATE GAN	0.572	0.581	0.558
Nonprivate baseline	0.712	0.738	0.695

6. DISCUSSION

The utility results suggest that PrivMedSynth can model mixed type clinical variables and retain useful predictive structure under the tested configuration. The reported total variation distance, downstream AUROC, and marginal comparisons are encouraging, but their evidential strength depends on the validity of the preprocessing, the independence of the evaluation partition, and the reproducibility of the training procedure. The corrected privacy analysis changes the principal interpretation. A local mechanism and a cloud mechanism compose, and their privacy costs do not cancel. Under record level adjacency, the conservative guarantee is the sum of the complete local and cloud privacy parameters, with the local and cloud failure probabilities also added. The Gaussian local release therefore contributes both an epsilon term and a delta term. If multiple categorical fields are randomized independently, their per field privacy values must be added. This issue is especially important for high cardinality diagnostic data, where a single per field value can otherwise be mistaken for a complete record level guarantee. The conditional mechanism also requires a precise statement of the target distribution. Uniform conditioning can improve minority class coverage, which may help a downstream classifier, but the resulting synthetic data are no longer a direct sample from the natural class prior. Both objectives are legitimate. They should be reported separately, with the target condition distribution and calibration procedure made explicit. The supplied ablation plots indicate that conditional guidance and local randomization affect utility. The plots do not establish causality beyond the tested configurations, and they do not show that local randomization lowers the cloud privacy cost. A complete ablation must keep the privacy unit, total composed budget, data partition, training steps, and model capacity comparable across configurations. The present formulation targets fixed length tabular records. It does not explicitly model irregular longitudinal events, waveform data, images, genomic data, or clinical text. The source data transformations are not described in sufficient detail to verify the reported feature counts and class ratios. In addition, several supplied numerical results lack confidence intervals, repeated seeds, and implementation details. The privacy budget table contained arithmetic inconsistencies and omitted the local Gaussian failure probability, so a final end to end privacy claim requires a complete accountant output. Future work should implement patient level rather than record level protection for datasets with repeated admissions, develop temporal diffusion models for irregular event sequences, evaluate multimodal records,

and compare exact analytic Gaussian accounting with subsampled RDP accounting. Stronger privacy audits should include attribute inference, reconstruction, linkage, and model inversion attacks. Resource measurements on edge devices and cloud hardware are also needed to assess deployment feasibility.

7. CONCLUSION

PrivMedSynth provides a framework for conditional diffusion based synthesis of mixed type EHR data in an edge cloud setting. The method combines local Gaussian perturbation for continuous variables, generalized randomized response for categorical variables, conditional diffusion, cloud DP SGD, and RDP accounting. The revised equations define the mechanisms and their composition without attributing an unsupported privacy reduction to local perturbation or treating membership inference AUROC as a formal DP bound. The supplied experiments report favorable utility and empirical attack results for PrivMedSynth. Those results are informative but should be treated as provisional until the preprocessing, privacy unit, local field budget, cloud accountant, and repeated trial protocol are fully documented and rerun. With those corrections, the framework can provide a sound basis for evaluating privacy preserving clinical data synthesis across distributed Medical IoT environments.

AUTHOR CONTRIBUTION STATEMENT

All authors contributed equally to the study conception and design. Material preparation, data collection, and analysis were performed by the authors. The first draft of the manuscript was written by the authors, and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

ETHICS APPROVAL AND CONSENT TO PARTICIPATE

Ethics declaration: not applicable. This study did not involve human participants or animals. Therefore, ethical approval and consent to participate are not applicable.

CONSENT FOR PUBLICATION

Consent to Publish declaration: not applicable.

DATA AVAILABILITY

The datasets analyzed during this study are publicly available from the PhysioNet repository and the UCI Machine Learning Repository. The MIMIC-III and eICU datasets are available through PhysioNet subject to registration, credentialing, and applicable data-use requirements. The PhysioNet CinC-2012 dataset is publicly available through PhysioNet, and the UCI Heart Disease dataset is available from the UCI Machine Learning Repository.

ACKNOWLEDGMENTS

The authors acknowledge the use of ChatGPT-5 for assistance in refining the manuscript. The authors reviewed and edited all AI-generated output and take full responsibility for the final content.

FUNDING

No funding.

DISCLOSURE STATEMENT

The authors declare no conflict of interest.

REFERENCES

[1] A. E. Johnson, T. J. Pollard, L. Shen, L.-w. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Anthony Celi, and R. G. Mark, "Mimic-iii, a freely accessible critical care database," *Scientific data*, vol. 3, no. 1, p. 160035, 2016.

[2] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*, pp. 265–284, Springer, 2006.

[3] A. Kotelnikov, D. Baranchuk, I. Rubachev, and A. Babenko, "Tabddpm: Modelling tabular data with diffusion models," in *International conference on machine learning*, pp. 17564–17579, PMLR, 2023.

[4] M. L. Fang, D. S. Dhami, and K. Kersting, "Dp-ctgan: Differentially private medical data generation using ctgans," in *International conference on artificial intelligence in medicine*, pp. 178–188, Springer, 2022.

[5] C. Zhu, J. Tang, J. F. Pérez, M. van Dijk, and L. Y. Chen, "Dp-tldm: Differentially private tabular latent diffusion model," in *International Conference on Availability, Reliability and Security*, pp. 337–357, Springer, 2025.

[6] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American statistical association*, vol. 60, no. 309, pp. 63–69, 1965.

[7] Z. Li, B. Wang, J. Li, Y. Hua, and S. Zhang, "Local differential privacy protection for wearable device data," *Plos one*, vol. 17, no. 8, p. e0272766, 2022.

[8] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.

[9] I. Mironov, "Rényi differential privacy," in *2017 IEEE 30th computer security foundations symposium (CSF)*, pp. 263–275, IEEE, 2017.

[10] Z. Zhang, C. Yan, and B. A. Malin, "Membership inference attacks against synthetic health data," *Journal of biomedical informatics*, vol. 125, p. 103977, 2022.

[11] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and trends® in theoretical computer science*, vol. 9, no. 3-4, pp. 211–487, 2014.

[12] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.

[13] T. J. Pollard, A. E. Johnson, J. D. Raffa, L. A. Celi, R. G. Mark, and O. Badawi, "The eicu collaborative research database, a freely available multi-center database for critical care research," *Scientific data*, vol. 5, no. 1, p. 180178, 2018.

[14] I. Silva and G. Moody, "An open-source toolbox for analysing and processing physionet databases in matlab and octave," *Journal of open research software*, vol. 2, no. 1, 2014.