

DOI: https://doi.org/10.48009/1_iis_2026_113

Protecting healthcare machine learning from membership inference: An evaluation of differential privacy

Nick Pierrelouis, *Marymount University*, n0p89926@marymount.eduXiang Liu, *Marymount University*, xliu@marymount.edu

Abstract

Machine-learning models can reveal whether a patient's records were used in training, a membership-inference risk that differential privacy (DP) limits at some cost to accuracy. This study evaluates the aforementioned trade-off in a synthetic cardiac-monitoring setting using classical models. Using a synthetic cardiac-monitoring dataset of 100,000 records and a Gaussian Naïve Bayes classifier, we compare Laplace (ϵ -DP) and Gaussian ((ϵ, δ) -DP, $\delta = 1 \times 10^{-5}$) mechanisms implemented as central DP on the classifier's sufficient statistics across privacy budgets from $\epsilon = 0.1$ to 5.0 under a black-box membership inference attack (MIA). The non-DP baseline achieves 0.75 accuracy. With DP, accuracy stays near the non-private ceiling once the budget reaches $\epsilon = 1.0$ (Laplace 0.755, Gaussian 0.742). Tightening privacy below $\epsilon = 1.0$ trades accuracy for protection, and Laplace retains more accuracy than Gaussian at these strict budgets, identifying $\epsilon \approx 1.0$ as a practical operating point. Across all privacy budgets, attack success remains near the 50% random baseline (approximately 50.0%–50.9%), indicating limited observable membership leakage, with the pattern holding across Random Forest and Logistic Regression. Beyond the empirical results, the paper offers the Investigation-Implementation-Assessment (IIA) framework as a procedural blueprint for deploying DP in applied machine-learning pipelines, with guidance on mechanism choice, noise placement, accounting assumptions, and privacy-budget selection.

Keywords: differential privacy, healthcare machine learning, privacy-utility trade-off, membership inference attack, epsilon selection

Introduction

Healthcare ML models rely on personal health information for research and patient treatment, but are vulnerable to membership inference attacks (MIAs), which can reveal whether a patient's record was used for training. Recent healthcare data breaches affecting approximately 141 million patient records (Vallepu, 2024) have exposed weaknesses in existing privacy frameworks. As AI/ML Software as a Medical Device (SaMD) expands with accompanying cybersecurity guidance, deployable privacy protection becomes critical. Differential privacy (DP) provides mathematical guarantees that individual privacy remains protected when data contributes to model training (Dwork & Roth, 2014). DP mechanisms add calibrated noise to ensure that removing or adding any single record produces negligibly different results. This study applies central DP, privatizing the model's sufficient statistics and parameters rather than the raw feature values, with noise magnitude calibrated to epsilon (ϵ), to address inference attacks that exploit model outputs to reveal sensitive training data information (Shokri et al., 2017). Despite extensive theoretical work, a practical, evaluated procedure tying DP placement and mechanism choice to utility, membership-inference risk, and deployment runtime is lacking. This study provides a procedure, illustrated on a synthetic cardiac-monitoring use case whose features mirror FDA 510(k)-cleared ambulatory monitors. We

evaluate how DP mechanisms affect membership-inference risk and predictive utility and which ϵ values are operationally viable, against a black-box adversary that infers membership from output scores, relative to a 50% random baseline.

Background and Literature Review

DP started as a formal framework for quantifying privacy guarantees in statistical databases. DP bounds how much any single record can change an analysis or training outcome, formalized by parameters ϵ (privacy loss) and, for approximate DP, δ (failure probability). The ϵ value controls the trade-off between privacy and utility, with lower values providing stronger privacy but potentially degrading model performance, and higher values offering less privacy but maintaining model performance. In approximate DP, the parameter δ (Delta) is used to represent the probability that the privacy guarantee may fail and is constrained to be much smaller than $1/n$, where n is the dataset size (Dwork & Roth, 2014). Prior applied studies, such as Federated learning (FL) combined with DP (FL-DP) and privacy-aware federated ECG detection, report high accuracies. However, system designs and guarantees differ from our tabular, central DP configuration, limiting direct comparison (Alghieth, 2025; Bokhari et al., 2025). These studies show that favorable privacy-utility trade-offs are achievable in certain settings; however, systematic ϵ -swept evaluation outside deep learning remains limited, particularly for non-deep-learning approaches suited to resource-constrained clinical environments (Mizani et al., 2025), motivating our focus.

DP in Healthcare ML

DP in healthcare faces unique constraints compared to other industries because medical data exhibits high dimensionality, temporal dependencies, and class imbalance that can amplify utility loss under DP (Badawy et al., 2023; Orabe et al., 2025). Healthcare's regulatory and operational constraints heighten the need for verifiable privacy guarantees and transparent ϵ, δ reporting. Recent systematic reviews of DP in healthcare identify privacy-utility trade-off evaluation as a significant research gap. Most studies focus on algorithm development instead of empirical assessment of privacy-utility outcomes (Demelius et al., 2025). This lack of applied methods suggests that practitioners deploying diagnostic models in clinical settings may still rely on ad hoc experimentation rather than established best practices. The resulting gap supports this study's comprehensive evaluation across multiple ϵ values and mechanisms using IIA (Investigation, Implementation, Assessment). Hybrid approaches combine DP with complementary techniques: FL-DP frameworks, for example, enable distributed training across multiple healthcare institutions while providing formal privacy guarantees (Teo et al., 2024). Resource-adaptive FL emphasizes fair participation across institutions regardless of computing resources, addressing equity concerns in privacy protection (Zhang et al., 2025). Such advances reinforce combining DP with collaborative learning frameworks to provide enhanced privacy-utility trade-offs for multi-institutional healthcare applications.

Membership Inference Attacks

Privacy attacks on ML models exploit the tendency of models to memorize training data, potentially exposing sensitive information. MIAs determine whether specific data points contributed to model training by exploiting differential behavior on training versus testing data through analysis of model confidence scores, loss values, or prediction patterns (Shokri et al., 2017). MIA success varies with model complexity, dataset characteristics, and privacy protections. For healthcare models, MIAs can reveal whether patients participated in a study or were classified with certain conditions. Deep learning models with high capacity are particularly vulnerable due to their ability to memorize training data, while simpler models (e.g., Naïve Bayes, decision trees) show lower observed MIA success yet remain vulnerable without explicit protections (Fang et al., 2024), motivating our evaluation of classical, tabular models with central DP as a practical baseline for resource-constrained deployments.

Synthetic Data for Privacy Protection

Synthetic data reduces disclosure risk by approximating real data distributions. CTGAN (Conditional Tabular Generative Adversarial Network) learns conditional distributions over mixed data types and is used to generate the synthetic cardiac dataset in this study (Xu et al., 2019). Applying DP to synthetic data generation provides formal privacy guarantees; for instance, PATE-GAN combines multiple teacher models with DP principles to generate synthetic healthcare data using ensemble techniques to limit individual influence (Jordon et al., 2019). Synthetic healthcare data can be medically reasonable while complying with strong privacy guidelines, though its quality requires careful validation to ensure synthesis introduces no artifacts or biases.

Software as a Medical Device (SaMD)

As of 2025, the FDA had authorized more than 1,000 AI-enabled medical devices and issued lifecycle and cybersecurity guidance (Babic et al., 2025). Ambulatory cardiac monitoring systems constantly monitor heart rhythm for arrhythmia and atrial fibrillation detection, demanding both strict privacy and high utility (Steinberg et al., 2017); unlike imaging or genomic data, cardiac bio-signals follow well-characterized waveform and heart rate variability patterns, well-suited to privacy-preserving analysis (de Marvao et al., 2020). The privacy-utility tension is acute here because small accuracy losses can carry serious clinical consequences, and the real-time, low-power constraints of wearables leave little room for DP's computational overhead. Cardiac monitoring is also underrepresented in DP-health research, where even a moderate budget ($\epsilon = 10$) substantially degrades high-capacity models, especially for underrepresented populations (Tayebi Arasteh et al., 2023). Calibrating ϵ without losing clinical effectiveness remains difficult (Noh et al., 2022), though attack-aware noise calibration may ease the trade-off (Kulynych et al., 2024); the budget's cumulative depletion forces periodic retraining (Ziller et al., 2024), and ϵ 's abstractness invites misinterpretation (Nanayakkara et al., 2023). This study addresses these gaps in a cardiac use case, applying the IIA framework for structured parameter selection, attack simulation, and statistical validation.

Gaps in Current Literature

Most DP studies in healthcare focus on imaging or genomic applications, with bio-signal applications comprising few published studies (Mohammadi et al., 2026). The limited number of peer-reviewed sources offering practical guidance on ϵ and δ parameter selection for cardiac healthcare ML models represents an important gap (Saifuzzaman et al., 2022). Existing literature prioritizes either strong privacy guarantees at the expense of model performance or maintains utility with inadequate privacy protections (Zeng et al., 2022). Few studies systematically evaluate the full spectrum of privacy-utility trade-offs across multiple ϵ , mechanisms, and model architectures (Demelius et al., 2025).

Recent DP-ML healthcare studies span diverse domains, including ICU records (Batchala et al., 2025), medical imaging (Adnan et al., 2022; Tayebi Arasteh et al., 2023), biobank data (Qian et al., 2024), and genomic datasets (Chen et al., 2021). These studies explore various DP implementations: training-time DP-SGD (Adnan et al., 2022; Tayebi Arasteh et al., 2023), multiple noise mechanisms (Batchala et al., 2025), differentially private GANs (Qian et al., 2024), and federated learning approaches (Adnan et al., 2022; Shukla et al., 2025). However, we found no studies that evaluate privacy-utility trade-offs across comprehensive ϵ ranges, multiple mechanisms, and multiple architectures simultaneously. Only two studies we identified (Chen et al., 2021; Qian et al., 2024) evaluate MIA vulnerability, and none provide systematic frameworks for ϵ -selection based on empirical attack resistance. NIST SP 800-226 (Near et al., 2025) provides an authoritative framework for evaluating DP guarantees across parameters, algorithmic assumptions, and trust model. It is evaluative, not operational: it specifies what to assess, not how to choose a budget or noise placement for a task. IIA is complementary, giving a selection procedure that ties ϵ and placement to measured attack resistance. In summary, prior work rarely (i) specifies where to add noise

within applied pipelines, (ii) reports ϵ alongside an explicit selection procedure, and (iii) evaluates MIA success and utility across mechanisms and ϵ values in non-DL, tabular settings.

Main Contribution of the Paper

This paper makes three main contributions. First, it introduces the IIA framework as a procedural blueprint for deploying DP in applied ML pipelines, clarifying mechanism choice (Laplace vs. Gaussian), noise placement (central DP), accounting assumptions (including δ), and ϵ selection (Figure 1). Second, it provides an empirical evaluation on a synthetic cardiac-monitoring task that quantifies membership-inference risk and predictive utility across $\epsilon \in \{0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 5.0\}$ under a black-box adversary. Third, it identifies $\epsilon \approx 1.0$ as a practical operating point where utility approaches the non-private ceiling while membership-inference attack success remains near chance.




Figure 1. IIA Framework

Research Methodology

We formalize a deployment framework for DP in biomedical ML pipelines that specifies where perturbation occurs on the model's aggregated statistics and parameters, i.e., central DP, not on individual records, which mechanism is selected (Laplace ϵ -DP or Gaussian (ϵ, δ)-DP), and how the privacy budget (ϵ, δ) is calibrated

against target utility and attack resilience. Assumptions, parameter bounds, and evaluation criteria are made explicit to support reuse in related tabular clinical tasks. The methodology consists of three sequential phases: Investigation, Implementation, and Assessment, each with specific objectives, activities, and expected outcomes.

Data & Synthetic Generation

This study used CTGAN (300 training epochs) to generate 100,000 synthetic cardiac monitoring records (80/20 train-test split) with seven features: patient age, prior AF diagnosis, ECG readings, PPG readings, signal quality, thirty-minute context windows, and confidence scores. The generator uses a shared-latent design in which every feature derives from a hidden atrial-fibrillation propensity with independent measurement noise, so that no feature is a near-copy of the label. The synthetic features were checked for redundancy with the outcome: No single feature is strongly tied to the AF label (highest correlation ≈ 0.46 ; ECG-AF ≈ 0.44), and the best single-feature classifier reaches only ≈ 0.72 accuracy, slightly above the 0.625 majority-class baseline. Using all features, the non-private accuracy ceiling is ≈ 0.75 , with a positive prevalence of ≈ 0.372 .

Threat Model & Privacy Requirements

The Investigation phase identifies vulnerabilities and privacy risks in the ML training pipeline. For this experiment, it was tailored for ambulatory cardiac monitoring systems. Systematic assessment identifies potential attack vectors in the training pipeline. In this study, the primary vulnerability relevant to the evaluated black-box MIA arises from model prediction scores accessible through inference APIs. Other attack surfaces, such as gradient leakage in federated learning or exposed model checkpoints, represent broader deployment risks but were not evaluated here. The vulnerability assessment considers both direct disclosure (e.g., revealing AF diagnosis) and indirect inference through correlated features. Risk assessment evaluates potential harm if patient privacy is compromised, considering data sensitivity levels and attack likelihood. Risk quantification accounts for attack sophistication, adversary knowledge, and disclosure impact. The assessment prioritized membership inference (revealing patient participation in training) as the primary threat requiring mitigation. Based on identified vulnerabilities and assessed risks, privacy requirements establish that MIA success must fall below 52%, within 2 percentage points of the 50% random guessing baseline. Evaluation employed 5 independent runs with stratified 80/20 train-test splits (80,000/20,000 samples) to ensure reproducibility. These thresholds balance privacy protection with clinical utility requirements. The Investigation phase outputs a vulnerability assessment and privacy requirements that guide mechanism selection and parameter calibration in subsequent phases.

DP Mechanisms & Calibration

The Implementation phase focuses on executing operational protection strategies through DP mechanisms targeting the identified vulnerable points in the training pipeline. Key activities include model selection, mechanism selection, (ϵ, δ) -DP calibration, and noise addition implementation. Three model architectures were evaluated: Naïve Bayes (primary model), Random Forest (RF), and Logistic Regression (LR). Naïve Bayes (NB) was selected as the primary model due to computational efficiency, interpretability, and suitability for healthcare applications requiring transparent decision-making (Singh & Pal, 2024). All implementations used scikit-learn 1.8, NumPy 2.4.6, and pandas 3.0.3, with IBM Differential Privacy Library (Diffprivlib) 0.6.6 used for the Laplace-NB cross-check and a direct output-perturbation routine for LR (Chaudhuri et al., 2011) on Python 3.12.5. Feature bounds were calculated from training data (min-max per feature) to constrain noise addition to ensure privacy mechanisms operate within valid feature ranges without explicit normalization. Features were min-max scaled to $[0,1]$ using public training-set bounds, which bounds the sensitivity of the released statistics.

Both mechanisms used the same five seeds (42–46), so each seed produces a matched Laplace and Gaussian run on the same test set and budget. The mechanisms are compared by a paired t-test over these five seed-matched pairs, which rests on each mechanism's mean and standard deviation over the five runs. Selection of appropriate DP mechanisms depends on data types, attack vectors, and utility requirements. The Laplace mechanism achieves ϵ -DP by adding noise from a Laplace distribution, providing pure DP guarantees without probabilistic failure. The Gaussian mechanism achieves (ϵ, δ) -DP by adding noise from a Gaussian distribution, allowing a small probability δ of privacy guarantee failure in exchange for potentially improved utility at equivalent privacy levels. NB is made private by adding calibrated noise to its sufficient statistics, class counts, and per-class, per-feature sums and sums-of-squares (Vaidya et al., 2013); the Laplace and Gaussian variants differ only in the noise distribution drawn on those statistics. diffprivlib's GaussianNB, the same method, validates the Laplace NB. Random Forest uses diffprivlib's central-DP forest, and Logistic Regression uses output perturbation (Chaudhuri et al., 2011). The Laplace-versus-Gaussian comparison is therefore made on NB, where the noise distribution is the only thing that varies. During mechanism selection, classical ML algorithms like NB, RF, and LR indicated effective protection in this setting, while maintaining compatibility with existing implementations.

Careful tuning of (ϵ, δ) -DP achieves the desired balance between privacy protection and model utility; ϵ was selected by sweeping a fixed grid and choosing the point that best balances privacy loss against utility loss. For the Gaussian mechanism, the δ parameter is constrained by: $\delta \ll \frac{1}{n}$, where n represents the dataset size, ensuring the probability of privacy guarantee failure remains negligible. This study evaluates $\epsilon \in \{0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 5.0\}$ for Laplace and $\delta = 10^{-5}$ for Gaussian mechanisms based on dataset size ($n = 100,000$). Privacy noise is applied during model fitting rather than to the raw records. For the central-DP implementation, the released quantities are the model's aggregated sufficient statistics:

Each statistic is the class count $n(c)$, the feature sum $S(c,j) = \sum x(i,j)$, and the sum of squares $Q(c,j) = \sum x(i,j)^2$, giving the privatized set $\tilde{S} = \{ \tilde{n}(c), \tilde{S}(c,j), \tilde{Q}(c,j) \}$, in which every statistic t is released as $\tilde{t} = t + \eta$. With k classes and d features, there are $k(1 + 2d)$ statistics, and the privacy budget is divided across them by sequential composition, so each receives $\epsilon_q = \epsilon / [k(1 + 2d)]$ (and $\delta_q = \delta / [k(1 + 2d)]$ for the Gaussian mechanism). Features are min-max scaled to $[0,1]$ using public training-set bounds, so each statistic has L_1 -sensitivity $\Delta f = 1$. The Laplace scale and Gaussian standard deviation below are applied with this per-statistic budget, so ϵ and δ in those equations denote ϵ_q and δ_q . Each released statistic receives noise calibrated to its sensitivity and privacy allocation, the scale parameter for Laplace noise is: $b = \frac{\Delta f}{\epsilon}$ where Δf is the L_1 -sensitivity of the function f . For Gaussian noise, the standard deviation is:

$$\sigma = \frac{\Delta f \cdot \sqrt{2 \ln(1.25/\delta)}}{\epsilon}$$

The privatized statistics give the class priors and, for each class c and feature j , the mean $\tilde{\mu}(c,j) = \tilde{S}(c,j) / \tilde{n}(c)$ and variance $\tilde{\sigma}^2(c,j) = \tilde{Q}(c,j) / \tilde{n}(c) - \tilde{\mu}(c,j)^2$ (variances floored at a small constant) used at prediction time (Vaidya et al., 2013). DP can be applied centrally, to these aggregated statistics, or feature-wise, to individual records before training. This work uses central placement: a provider that controls its own training data operates under the central-DP trust model, and perturbing the aggregated statistics preserves the feature-label structure that per-record noise erodes at usable privacy budgets. Appendix A reports this comparison: at matched budgets, feature-wise perturbation collapses to the majority class while central DP retains utility.

Where Δf is the L_2 -sensitivity. With features min-max scaled to $[0,1]$ using public training-set bounds, each released statistic has unit sensitivity, with bounds = $(X_train.min(axis=0), X_train.max(axis=0))$ with

the privacy budget split across the $k(1+2d)$ released statistics by sequential composition. For the cardiac dataset, this resulted in L_1 -sensitivity of $\Delta f = \max(\text{feature_max} - \text{feature_min})$ across all features, which under the $[0,1]$ scaling gives each released statistic unit sensitivity.

Experimental Setup & Evaluation

The Assessment phase evaluates the effectiveness of implemented privacy-preserving strategies and their impact on model performance. This phase, where simulation-based testing may occur, includes privacy guarantee verification, utility impact, computational, and statistical analysis. Privacy protection effectiveness requires validation beyond theoretical guarantees through adversarial testing. This study evaluates the membership attack vector to empirically assess the resilience of DP implementations: MIAs (determining training set inclusion). The Adversarial Robustness Toolbox (ART) provides standardized attack implementations for reproducible evaluation: Membership Inference exploits loss value differences $L(x, y; \theta)$ between training and test data. This is a black-box attack, the scenario most relevant to a deployed service in which the adversary has only query access.

Measurement of privacy protection levels ensures that implemented DP mechanisms align with theoretical guarantees. Empirical verification confirmed that privacy budgets were tracked correctly by the accountant and that noise addition behaved as expected across epsilon values, without claiming to verify the formal DP guarantee:

$$\max_{D, D'} \left| \log \frac{\Pr[M(D) = o]}{\Pr[M(D') = o]} \right| \leq \epsilon$$

where D and D' are neighboring datasets differing by one record, and o represents any possible output. For approximate DP with Gaussian mechanisms:

$$\Pr \left[\max_{D, D'} \left| \log \frac{\Pr[M(D) = o]}{\Pr[M(D') = o]} \right| > \epsilon \right] \leq \delta$$

Systematic evaluation of privacy mechanism impact on model performance assesses accuracy, precision, recall, and F1-score across privacy configurations. Utility degradation is measured through comparative analysis:

$$\text{Utility Loss} = \frac{\text{Metric}_{\text{baseline}} - \text{Metric}_{\text{DP}}}{\text{Metric}_{\text{baseline}}} \times 100\%$$

where $\text{Metric}_{\text{baseline}}$ represents performance without DP and $\text{Metric}_{\text{DP}}$ represents performance with DP applied. Utility loss analysis identifies optimal privacy configurations providing acceptable utility degradation while achieving privacy objectives. Simulated attacks verify the effectiveness of privacy-preserving implementations against real-world threats. ART modeled MIAs to assess model resilience as the attacks attempt to exploit the tendency of models to produce lower loss values for training data:

$$A(x; \varphi): X \rightarrow \{0 \text{ (non-member)}, 1 \text{ (member)}\}$$

where the adversary A uses parameters φ to map input x to binary membership decisions. If $L(x, y; \theta) < \tau$, the point is classified as a member, where L is the model's loss function, and τ is the attack threshold.

The Assessment phase produces an output detailing privacy-utility balance and optimization paths, informing deployment decisions by quantifying trade-offs between privacy protection strength, model

utility, and computational efficiency. Models were evaluated using standard classification metrics (accuracy, precision, recall, F1-score), with privacy protection assessed through MIA success rates. The non-DP model achieved 49.8% attack success at random level, indicating no substantial privacy leakage before DP protection was applied. Computational efficiency metrics (inference time in milliseconds per sample and memory usage in megabytes) were evaluated to assess deployment feasibility for resource-constrained environments like wearable cardiac monitors and mobile health applications. Statistical validation employed paired t -tests comparing Laplace and Gaussian mechanisms with matching epsilon values across 5 independent iterations (significance threshold $\alpha=0.05$). Preprocessing order: feature-specific bounds were computed from training data as $[\min(x_j), \max(x_j)]$ for each feature j , then models were trained with DP, where calibrated noise (Laplace for ϵ -DP or Gaussian for (ϵ, δ) -DP) was added during training using these bounds. All experiments used conventional hardware (Apple M4 Pro, 12-core CPU, 24 GB unified memory) to demonstrate practical feasibility. Core libraries supporting DP implementation were Python, IBM's Diffprivlib, ART, and Jupyter notebook. NB was the primary classifier, with RF and LR used for comparative analysis.

Findings

The Laplace mechanism achieved a mean accuracy ranging from 0.6877 at $\epsilon = 0.1$ to 0.7557 at $\epsilon = 5.0$. The Gaussian mechanism followed the same trend but recovered more slowly, ranging from 0.6341 at $\epsilon = 0.1$ to 0.7560 at $\epsilon = 5.0$. The non-private NB ceiling was 0.7555, so both mechanisms approach it once $\epsilon \geq 1.0$. Five independent runs per mechanism (seeds 42–46 for both mechanisms) show that variance is largest at strict budgets and shrinks as ϵ grows: the Laplace standard deviation falls from 0.0184 at $\epsilon = 0.1$ to below 0.001 by $\epsilon = 1.0$, and the Gaussian from 0.0803 to 0.0170 over the same range (Table 1). Variance is largest at strict budgets and shrinks as ϵ grows, a privacy-utility trade-off at small ϵ that closes as the budget increases.

Table 1. Laplace and Gaussian Mechanism Performance

Epsilon (ϵ)	Laplace Accuracy	Laplace F1	Gaussian Accuracy	Gaussian F1
0.1	0.6877	0.4559	0.6341	0.4816
0.5	0.7520	0.6744	0.7065	0.5370
1.0 ^a	0.7554	0.6847	0.7421	0.6471
2.0	0.7555	0.6849	0.7543	0.6814
5.0	0.7557	0.6851	0.7560	0.6859

Note. Values represent the mean across five independent runs with consistent random seeds (42–46 for Laplace and Gaussian). ^a $\epsilon = 1.0$ indicates the recommended optimal privacy budget based on privacy-utility analysis. Baseline (Non-DP): Accuracy = 0.7555, F1 = 0.6847. Laplace recovers utility faster than Gaussian at strict budgets, with the two converging by $\epsilon \approx 2.0$; MIA remains near the 0.50 chance level throughout.

Mechanism Comparison at $\epsilon = 1.0$

Laplace and Gaussian mechanisms behaved differently across ϵ . At $\epsilon = 1.0$, Laplace reached 0.7554 ± 0.0006 and Gaussian 0.7421 ± 0.0170 ; Laplace was higher, though the paired t -test was not significant at this budget ($t = 1.598$, $p = 0.185$), reflecting Gaussian's larger run-to-run variance as it nears the ceiling. At strict budgets Laplace led by about 5 points at $\epsilon = 0.1$, though the difference was not significant ($t = 1.28$, $p = 0.271$) given Gaussian's high variance; the advantage was significant at $\epsilon = 0.5$ (about 4.5 points, $t = 3.18$, $p = 0.033$), and the two converge by $\epsilon \approx 2.0$. The pattern follows the mechanisms' noise scales: with the budget split across the released statistics, the Gaussian standard deviation carries the $\sqrt{(2 \ln(1.25/\delta))}$

factor, roughly five times the Laplace scale at the same per-statistic budget, so Gaussian pays more when ϵ is small (Table 2).

Table 2. Mechanism Statistical Summary at $\epsilon = 1.0$

Metric	Laplace	Gaussian	Abs. Difference	Favors
Accuracy	0.7554 ± 0.0006	0.7421 ± 0.0170	0.0133	Laplace
F1-Score	0.6847	0.6471	0.0376	Laplace
Utility Loss (%) ^a	0.01	1.77	1.76	Laplace
MIA Success Rate ^b	0.509	0.501	0.008	~chance
Paired t-test (t) ^c	—	—	1.598	—
p-value	—	—	0.185	—
Significant at $\alpha=0.05$?	—	—	No (ns at $\epsilon = 1.0$)	—

Note. NB classifier, 5 runs at $\epsilon = 1.0$. ^a Utility Loss = (Baseline – DP Accuracy) / Baseline $\times$ 100%. ^b MIA accuracy vs. 0.500 chance; ~0.500 = no membership advantage. ^c Paired t-test on accuracy; ns at $\epsilon = 1.0$, significant in Laplace's favor at $\epsilon = 0.5$.

Membership-Inference Results

MIA evaluation showed attack success near the random baseline across all privacy budgets for both mechanisms (approximately 50.0-50.9%, against a 50.0% baseline). This holds for a model that learned the task: the rule-based attack still cannot separate members from non-members. This result is consistent with research showing that well-generalized models resist MIAs because there is no distinguishable signal between training and test samples. DP therefore provides formal mathematical guarantees on top of inherent privacy from good generalization. This experimental result supports a layered defense: practical privacy, as attacks fail even on non-DP models (approximately 49% success); formal guarantees, as DP provides mathematical proof of privacy bounds; high utility, with accuracy near the 0.7555 ceiling at $\epsilon \geq 1.0$; and regulatory compliance, as formal privacy definitions satisfy governance requirements. Because the ART membership attack is a black-box test, it reflects the realistic deployment threat in which an adversary has only query access, so this result bounds membership risk under that attack class; stronger white-box or trained-shadow attacks, such as LiRA, which assume greater access, are left to future work.

Model Architecture Comparison

RF preserved utility under central DP (0.41% loss at $\epsilon = 1.0$) across all ϵ values. This resilience stems from the ensemble structure, where noise added to individual trees averages out during aggregation, preserving overall prediction accuracy (Liu et al., 2023). NB had the smallest loss (0.01% Laplace, 1.77% Gaussian at $\epsilon=1.0$), while LR stayed within 3% utility loss (2.84% at $\epsilon = 1.0$), under central DP. LR carries the largest DP cost because its output-perturbation routine must also regularize, a difficulty that grows with dimensionality (Hu et al., 2022). These differences reflect each model's sensitivity to noise on its parameters or statistics, not feature-level noise. Inference cost is noise-dominated and negligible across models; differences reflect model class, with RF the slowest (Table 3).

Table 3. Model Performance Comparison at $\epsilon = 1.0$

Model	Mechanism	Acc	F1	Utility Loss	MIA
NB	Laplace	0.7554	0.6847	0.01%	0.509
NB	Gaussian	0.7421	0.6471	1.77%	0.501
RF	diffprivlib (central)	0.7536	0.6383	0.41%	0.502
LR	output perturbation	0.7429	0.6526	2.84%	0.499

Note. Inference time measured across 50 trials on 20,000 test samples (prediction only; training time excluded). Utility loss is measured against each model's own non-private accuracy (NB 0.7555, RF 0.7568, LR 0.7647); timing is noise-dominated and reported only with that caveat.

Discussion of Findings

This research demonstrated that DP can be effectively integrated into cardiac healthcare ML models while maintaining clinical utility. $\epsilon = 1.0$ emerged as a practical operating point for healthcare organizations implementing privacy-preserving analytics. At $\epsilon = 1.0$, accuracy stays near the non-private ceiling (Laplace 0.7554, Gaussian 0.7421) and attack success near the 0.50 baseline. The non-private model also sits at chance (0.498), so this does not show DP defeating the attack: a well-generalized classifier resists black-box membership inference on this synthetic tabular task with or without DP, and the finding bounds risk under one attack rather than proving robustness. DP's value is its formal worst-case guarantee, which holds regardless of any single attack and matters most outside this study: higher-capacity or memorization-prone models, smaller or imbalanced datasets, rare subpopulations, and real patient data.

Evaluation across $\epsilon \in \{0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 5.0\}$ revealed a clear inflection point near $\epsilon = 1.0$, where utility loss is minimal at a still-meaningful privacy budget, providing evidence-based guidance for parameter selection and addressing a well-documented gap in DP literature regarding ϵ interpretability (Cummings et al., 2024). Statistical comparison between mechanisms (paired t-test at $\epsilon = 1.0$: $t = 1.598$, $p = 0.185$) showed Laplace ahead, though not significantly at this budget; the advantage is significant at $\epsilon = 0.5$. Model architecture comparison at $\epsilon = 1.0$ reveals distinct performance profiles. NB demonstrates optimal balance: minimal utility loss (0.01% for Laplace), negligible inference cost, and strong privacy protection, making it the primary recommendation for cardiac monitoring applications. RF exhibits resilience to DP noise (0.41% utility loss), consistent with prior research on ensemble methods' robustness (Liu et al., 2023), but shows higher inference cost than NB, which is noise-dominated and small in absolute terms. Under central DP, all three model types preserve utility (loss under 3%); NB is the lightest and the primary recommendation.

Implications for Research/Practice

Utility loss under central DP stays under 3 percent at $\epsilon = 1.0$, so privacy protection does not meaningfully degrade predictions. Minimal computational overhead at $\epsilon = 1.0$ (Table 3), negligible across all three models, suggests that DP-protected models may be feasible for resource-constrained environments, including wearable devices and mobile applications. More broadly, the findings suggest that DP can be incorporated into healthcare ML workflows without substantial loss of predictive performance in this setting. For maximum privacy applications (rare disease screening, pediatric populations), $\epsilon \leq 0.1$ is recommended despite utility costs, though accuracy drops substantially at this budget (about 0.69 for Laplace and 0.63 for Gaussian), so it suits only applications willing to trade accuracy for the strongest formal guarantee. For general clinical applications (routine cardiac monitoring, chronic disease management), $\epsilon = 0.5$ -1.0 provides a reasonable balance, with Laplace already near the ceiling at $\epsilon = 0.5$ (0.7520). For rapid response applications (emergency cardiac monitoring, acute event detection), $\epsilon = 1.0$ reaches full utility (Laplace 0.7554) with chance-level membership inference. DP in healthcare ML systems has significant social implications, promoting autonomy, building public confidence, and ensuring equitable access to health technology.

Public willingness to share health data correlates with trust in healthcare systems, making it essential to demonstrate that clinical accuracy can be maintained with strong privacy protections. However, if DP measures are resource-intensive, privacy protections may only be available at well-resourced institutions, compounding inequity. Resource-adaptive frameworks can increase both privacy and equity (Zhang et al., 2025). This study's demonstration of computational feasibility through Laplace mechanisms and efficient architectures makes DP scalable and implementable in constrained infrastructure, lowering adoption barriers and democratizing privacy protections. Regulatory bodies increasingly require demonstrable privacy controls, yet adequate privacy remains loosely defined, and algorithmic-bias concerns call for diverse datasets, fairness-aware methods, and transparent AI (Ueda et al., 2024). The IIA framework offers practitioners and policymakers a replicable process for operationalizing responsible AI that balances privacy, performance, equity, and regulation, supporting compliance with emerging data-protection standards.

Limitations

This study used synthetic cardiac monitoring data rather than real patient data, limiting generalizability to actual clinical deployments. The task has a non-private ceiling near 0.75, so the reported accuracies reflect a model that learned the underlying signal. While synthetic data maintained physiologically plausible relationships and statistical properties, it may not capture the full complexity of real-world cardiac monitoring systems. Rare arrhythmia variations or borderline disease features may be underrepresented in synthetic datasets. Validation with real clinical data remains necessary to confirm these findings in operational settings. The study focused on three model architectures (NB, RF, and LR), excluding deep learning models that may exhibit different privacy-utility trade-offs and could benefit from alternative DP implementations such as DP-SGD. Only Laplace and Gaussian mechanisms were evaluated; alternative mechanisms, such as attack-aware noise calibration, were not analyzed (Kulynych et al., 2024). The attack evaluation covered black-box MIA only; the attack was a black-box test, the realistic case of query-only access, so stronger attacks that assume greater adversary access, including trained-shadow attacks such as LiRA, white-box, and adaptive attacks, remain untested. The observed chance-level MIA results may not extend to datasets with greater class imbalance, lower quality, or rare event detection scenarios, where stricter DP guarantees may be required.

Future work should validate the IIA framework on real clinical datasets and additional model classes, including deep temporal models trained with DP-SGD, to test whether the $\epsilon \approx 1.0$ operating point generalizes. Evaluation could also expand beyond accuracy to include calibration, AUROC, PR-AUC, and subgroup fairness analyses reported with confidence intervals across independent DP draws. Additional research should examine stronger attack settings, advanced DP mechanisms, automated parameter selection for edge deployments, and real-world adoption barriers in clinical environments.

Conclusions

Across seven privacy budgets ($\epsilon = 0.1\text{--}5.0$) and two noise mechanisms, $\epsilon = 1.0$ emerged as a practical operating point that balances formal privacy guarantees with diagnostic accuracy in cardiac monitoring. A critical finding is that a model that learned the task (accuracy ≈ 0.75) still resisted the black-box MIA, the realistic query-only deployment threat, so this bounds risk under that attack class rather than proving robustness. The attack stayed at chance (approximately 50.0-50.9%) across all privacy budgets, indicating that, in this synthetic tabular setting, appropriate model design already limited observable black-box membership leakage, a protection that DP mechanisms complement with formal mathematical guarantees.

At $\varepsilon \approx 1.0$, both mechanisms maintained diagnostic performance within ≈ 0.01 to 1.8% of the non-private ceiling (0.7555), demonstrating minimal utility loss combined with formal privacy guarantees and empirical attack resistance. The IIA framework provides a procedural blueprint specifying threat modeling, mechanism selection, parameterization, and evaluation protocols, coupling MIA resilience with utility and computational feasibility. This methodology can be adapted to various healthcare ML applications, providing replicable guidance for balancing privacy protection, model utility, and clinical deployment feasibility. The low computational overhead demonstrates that privacy and performance need not be competing priorities, supporting broader adoption of privacy-preserving analytics in healthcare as regulatory requirements evolve and systems continue to digitize.

References

- Adnan, M., Kalra, S., Cresswell, J. C., Taylor, G. W., & Tizhoosh, H. R. (2022). Federated learning and differential privacy for medical image analysis. *Scientific Reports*, 12(1), 1953. <https://doi.org/10.1038/s41598-022-05539-7>
- Alghieth, M. (2025). DeepECG-Net: A hybrid transformer-based deep learning model for real-time ECG anomaly detection. *Scientific Reports*, 15(1), 20714. <https://doi.org/10.1038/s41598-025-07781-1>
- Babic, B., Glenn Cohen, I., Stern, A. D., Li, Y., & Ouellet, M. (2025). A general framework for governing marketed AI/ML medical devices. *Npj Digital Medicine*, 8(1), 328. <https://doi.org/10.1038/s41746-025-01717-9>
- Badawy, M., Ramadan, N., & Hefny, H. A. (2023). Healthcare predictive analytics using machine learning and deep learning techniques: A survey. *Journal of Electrical Systems and Information Technology*, 10(1), 40. <https://doi.org/10.1186/s43067-023-00108-y>
- Batchala, R., Jain, P., Gyanchandani, M., Shukla, S., & Wadhvani, R. (2025). Enhance differential privacy mechanisms for clinical data analysis using CNNs and reinforcement learning. *Journal of Big Data*, 12(1), 178. <https://doi.org/10.1186/s40537-025-01215-5>
- Bokhari, S. M., Sohaib, S., & Shafi, M. (2025). Fusion of Personalized Federated Learning (PFL) with Differential Privacy (DP) Learning for Diagnosis of Arrhythmia Disease. *PloS One*, 20(7), e0327108. <https://doi.org/10.1371/journal.pone.0327108>
- Chaudhuri, K., Monteleoni, C., & Sarwate, A. D. (2011). Differentially Private Empirical Risk Minimization. *Journal of Machine Learning Research: JMLR*, 12, 1069–1109.
- Chen, J., Wang, W. H., & Shi, X. (2021). Differential Privacy Protection Against Membership Inference Attack on Machine Learning for Genomic Data. *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, 26, 26–37.
- Cummings, R., Desfontaines, D., Evans, D., Geambasu, R., Huang, Y., Jagielski, M., Kairouz, P., Kamath, G., Oh, S., Ohrimenko, O., Papernot, N., Rogers, R., Shen, M., Song, S., Su, W., Terzis, A., Thakurta, A., Vassilvitskii, S., Wang, Y.-X., ... Zhang, W. (2024). Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment. *Harvard Data Science Review*, 6(1). <https://doi.org/10.1162/99608f92.d3197524>

- de Marvao, A., Dawes, T. J. W., & O'Regan, D. P. (2020). Artificial Intelligence for Cardiac Imaging-Genetics Research. *Frontiers in Cardiovascular Medicine*, 6, 195. <https://doi.org/10.3389/fcvm.2019.00195>
- Demelius, L., Kern, R., & Trügler, A. (2025). Recent Advances of Differential Privacy in Centralized Deep Learning: A Systematic Survey. *ACM Comput. Surv.*, 57(6), 158:1-158:28. <https://doi.org/10.1145/3712000>
- Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2016). Calibrating Noise to Sensitivity in Private Data Analysis. *Journal of Privacy and Confidentiality*, 7(3), 17–51. <https://doi.org/10.29012/jpc.v7i3.405>
- Dwork, C., & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4), 211–407. <https://doi.org/10.1561/04000000042>
- Fang, C., Dziedzic, A., Zhang, L., Oliva, L., Verma, A., Razak, F., Papernot, N., & Wang, B. (2024). Decentralised, collaborative, and privacy-preserving machine learning for multi-hospital data. *eBioMedicine*, 101, 105006. <https://doi.org/10.1016/j.ebiom.2024.105006>
- Hu, L., Ni, S., Xiao, H., & Wang, D. (2022). High Dimensional Differentially Private Stochastic Optimization with Heavy-tailed Data. *Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, 227–236. SIGMOD/PODS '22: International Conference on Management of Data. <https://doi.org/10.1145/3517804.3524144>
- Jordon, J., Yoon, J., & van der Schaar, M. (2019). PATE-GAN: Generating Synthetic Data with Differential Privacy Guarantees. *International Conference on Learning Representations*.
- Kulynych, B., Gomez, J. F., Kaissis, G., du Pin Calmon, F., & Troncoso, C. (2024). Attack-Aware Noise Calibration for Differential Privacy. *Advances in Neural Information Processing Systems*, 37, 134868–134901.
- Liu, J., Li, X., Wei, Q., Liu, S., Liu, Z., & Wang, J. (2023). A two-phase random forest with differential privacy. *Applied Intelligence*, 53(10), 13037–13051. <https://doi.org/10.1007/s10489-022-04119-6>
- Mizani, M. A., Sheikh, A., & Banerjee, A. (2025). An empirical assessment of differential privacy in real-world observational data: A case-control study of asthma exacerbation in UK Biobank linked with electronic health records. *Journal of the American Medical Informatics Association*, 32(8), 1328–1339. <https://doi.org/10.1093/jamia/ocaf090>
- Mohammadi, M., Vejdanihemmat, M., Lotfinia, M., Rusu, M., Truhn, D., Maier, A., & Tayebi Arasteh, S. (2026). Differential privacy for medical deep learning: Methods, tradeoffs, and deployment implications. *Npj Digital Medicine*, 9(1), 93. <https://doi.org/10.1038/s41746-025-02280-z>
- Nanayakkara, P., Smart, M. A., Cummings, R., Kaptchuk, G., & Redmiles, E. M. (2023). What are the chances? Explaining the epsilon parameter in differential privacy. *Proceedings of the 32nd USENIX Conference on Security Symposium, SEC '23*, 1613–1630.
- Near, J. P., Darais, D., Lefkovitz, N., & Howarth, G. S. (2025). *Guidelines for Evaluating Differential Privacy Guarantees* (NIST Special Publication 800-226). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-226>

- Orabe, Z., Vasankari, A., Pahikkala, T., Kaisti, M., & Airola, A. (2025). *Securing Deep Learning Models with Differential Privacy for Cardiovascular Disease Prediction* (SSRN Scholarly Paper No. 5375201). Social Science Research Network. <https://doi.org/10.2139/ssrn.5375201>
- Qian, Z., Callender, T., Cebere, B., Janes, S. M., Navani, N., & van der Schaar, M. (2024). Synthetic data for privacy-preserving clinical risk prediction. *Scientific Reports*, 14(1), 25676. <https://doi.org/10.1038/s41598-024-72894-y>
- Saifuzzaman, M., Ananna, T. N., Chowdhury, M. J. M., Ferdous, M. S., & Chowdhury, F. (2022). A systematic literature review on wearable health data publishing under differential privacy. *International Journal of Information Security*, 21(4), 847–872. <https://doi.org/10.1007/s10207-021-00576-1>
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership Inference Attacks Against Machine Learning Models. *2017 IEEE Symposium on Security and Privacy (SP)*, 3–18. <https://doi.org/10.1109/SP.2017.41>
- Shukla, S., Rajkumar, S., Sinha, A., Esha, M., Elango, K., & Sampath, V. (2025). Federated learning with differential privacy for breast cancer diagnosis enabling secure data sharing and model integrity. *Scientific Reports*, 15(1), 13061. <https://doi.org/10.1038/s41598-025-95858-2>
- Singh, G., & Pal, A. (2024). Exploring Explainable Machine Learning in Healthcare: Closing the Predictive Accuracy and Clinical Interpretability Gap. In Y. Singh, P. K. Singh, P. J. S. Gonçalves, & A. K. Kar (Eds.), *Proceedings of International Conference on Recent Innovations in Computing* (pp. 167–182). Springer Nature. https://doi.org/10.1007/978-981-97-2839-8_12
- Steinberg, J. S., Varma, N., Cygankiewicz, I., Aziz, P., Balsam, P., Baranchuk, A., Cantillon, D. J., Dilaveris, P., Dubner, S. J., El-Sherif, N., Krol, J., Kurpesa, M., Rovere, M. T. L., Lobodzinski, S. S., Locati, E. T., Mittal, S., Olshansky, B., Piotrowicz, E., Saxon, L., ... Piotrowicz, R. (2017). 2017 ISHNE-HRS expert consensus statement on ambulatory ECG and external cardiac monitoring/telemetry. *Heart Rhythm*, 14(7), e55–e96. <https://doi.org/10.1016/j.hrthm.2017.03.038>
- Tayebi Arasteh, S., Lotfinia, M., Nolte, T., Sähn, M.-J., Isfort, P., Kuhl, C., Nebelung, S., Kaissis, G., & Truhn, D. (2023). Securing Collaborative Medical AI by Using Differential Privacy: Domain Transfer for Classification of Chest Radiographs. *Radiology: Artificial Intelligence*, 6(1), e230212. <https://doi.org/10.1148/ryai.230212>
- Teo, Z. L., Jin, L., Li, S., Miao, D., Zhang, X., Ng, W. Y., Tan, T. F., Lee, D. M., Chua, K. J., Heng, J., Liu, Y., Goh, R. S. M., & Ting, D. S. W. (2024). Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. *Cell Reports Medicine*, 5(2), 101419. <https://doi.org/10.1016/j.xcrm.2024.101419>
- Ueda, D., Kakinuma, T., Fujita, S., Kamagata, K., Fushimi, Y., Ito, R., Matsui, Y., Nozaki, T., Nakaura, T., Fujima, N., Tatsugami, F., Yanagawa, M., Hirata, K., Yamada, A., Tsuboyama, T., Kawamura, M., Fujioka, T., & Naganawa, S. (2024). Fairness of artificial intelligence in healthcare: Review and recommendations. *Japanese Journal of Radiology*, 42(1), 3–15. <https://doi.org/10.1007/s11604-023-01474-3>

- Vaidya, J., Basu, A., Shafiq, B., & Hong, Y. (2013). Differentially private naïve bayes classification. *Proceedings - 2013 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2013*, 571–576. <https://doi.org/10.1109/WI-IAT.2013.80>
- Vallepu, R. (2024). Exploring Data Security and Privacy Challenges in Master Data Governance Systems. *International Journal of Computer Trends and Technology*, 72(11), 126–134. <https://doi.org/10.14445/22312803/IJCTT-V72I11P113>
- Xu, L., Skoularidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling Tabular data using Conditional GAN. *Advances in Neural Information Processing Systems*, 32. <https://proceedings.neurips.cc/paper/2019/hash/254ed7d2de3b23ab10936522dd547b78-Abstract.html>
- Zeng, X., Yang, C., & Dai, B. (2022). Utility–Privacy Trade-Off in Distributed Machine Learning Systems. *Entropy*, 24(9), 1299. <https://doi.org/10.3390/e24091299>
- Zhang, F., Zhai, D., Bai, G., Jiang, J., Ye, Q., Ji, X., & Liu, X. (2025). Towards fairness-aware and privacy-preserving enhanced collaborative learning for healthcare. *Nature Communications*, 16(1), 2852. <https://doi.org/10.1038/s41467-025-58055-3>
- Ziller, A., Mueller, T. T., Stieger, S., Feiner, L. F., Brandt, J., Braren, R., Rueckert, D., & Kaissis, G. (2024). Reconciling privacy and accuracy in AI for medical imaging. *Nature Machine Intelligence*, 6(7), 764–774. <https://doi.org/10.1038/s42256-024-00858-y>

Appendix A. Central versus Feature-wise Differential Privacy

DP placement determines the utility cost. Central DP adds noise to aggregated statistics (class counts, feature sums, sums of squares); because each aggregate summarizes the full dataset, one record's influence is small, and little utility is lost. Feature-wise DP adds noise to each record's features before training. With the budget split across the d features, the per-record Laplace scale is approximately $\Delta f/(\epsilon/d) \approx 7/\epsilon$ for seven features in $[0,1]$, which swamps the class-conditional signal until ϵ is large. This collapses the classifier to the majority-class rate ($F1 \approx 0$) across the swept budgets (Table 4). At matched budgets, placement alone decides whether the model is usable: central recovers the non-private ceiling by $\epsilon \approx 1$, feature-wise does not. The feature-wise attack reading near 0.51 is not a privacy result, since a single-class predictor has nothing to leak. The recommendation is to privatize aggregates, not raw features, unless an untrusted-aggregator setting forces local DP.

Table 4. Noise placement: test accuracy by ϵ (same model and data; only placement differs).

Epsilon (ϵ)	Central Laplace	Central Gaussian	Feature-wise Laplace	Feature-wise Gaussian
0.1	0.6877	0.6341	0.6251	0.6251
0.5	0.7520	0.7065	0.6251	0.6251
1.0	0.7554	0.7421	0.6251	0.6251
1.5	0.7554	0.7463	0.6251	0.6251
2.0	0.7555	0.7543	0.6251	0.6251
3.0	0.7554	0.7561	0.6251	0.6251
5.0	0.7557	0.7560	0.6382	0.6251