RESEARCH ARTICLE

PRISM-DP: Privacy-Resilient Infinite-Sequence Memory via Martingale-Based Differential Privacy Composition for Continual Autoencoder Learning in Medical IoT

Manas Kumar Yogi[1,*] **and Yamuna Mundru**[2]

[1] Computer Science and Engineering Department, Pragati University, Surampalem 533437, India
[2] Computer Science and Engineering (AI & ML) Department, Pragati University, Surampalem 533437, India

Abstract

Rehearsal-based continual learning in Medical IoT (MIoT) autoencoders requires replaying latent codes of past clinical tasks to prevent catastrophic forgetting. When these latent codes are derived from patient data, each replay constitutes a fresh privacy query, and existing differential privacy (DP) composition theorems—designed for finite, pre-specified query sequences—fail to bound cumulative privacy leakage over an unbounded task horizon. This paper presents PRISM-DP, a Privacy-Resilient Infinite-Sequence Memory framework that introduces a martingale-based DP composition theorem for unbounded replay sequences. PRISM-DP proves that a geometrically decaying per-task privacy schedule combined with a replay correlation decoupling lemma yields a convergent total privacy leakage bound for continual autoencoders. A privacy-amplified replay subsampling mechanism further tightens the per-replay leakage. Evaluated across 10 sequential clinical tasks on the PTB-XL, OhioT1DM, MIMIC-III Waveform, and WESAD datasets, PRISM-DP achieves an average AUC-ROC of 0.871, a backward transfer of -0.055, and keeps the total privacy budget within $\varepsilon = 0.667$ at $\varepsilon_{\text{cap}} = 1.0$, outperforming all DP-certified baselines while providing a formally certified privacy guarantee for infinite-horizon continual MIoT learning.

Keywords: continual learning, differential privacy, Medical IoT, autoencoder, martingale composition, unbounded task sequence, latent replay, Rényi DP, PTB-XL, OhioT1DM.

1 Introduction

Medical Internet of Things (MIoT) systems are deployed in inherently non-stationary clinical environments. A wearable cardiac monitor must adapt as a patient's arrhythmia pattern evolves over months; a continuous glucose monitor (CGM) must update its hypoglycaemia detection model as a diabetic patient changes insulin regimens; an ICU monitoring system must accommodate new disease presentations and treatment protocols introduced by seasonal

Submitted: 28 June 2026
Revised: 07 July 2026
Accepted: 24 September 2026
Published: 26 September 2026

Vol. 2, **No.** 3, 2026.
10.62762/BISH.2026.293480

***Corresponding author:**
✉ Manas Kumar Yogi
manas.yogi@gmail.com

Citation
Yogi, M. K., & Mundru, Y. (2026). PRISM-DP: Privacy-Resilient Infinite-Sequence Memory via Martingale-Based Differential Privacy Composition for Continual Autoencoder Learning in Medical IoT. *Biomedical Informatics and Smart Healthcare*, 2(3), 136–151.

outbreaks [1, 2]. This non-stationarity demands continual learning—the sequential acquisition of new clinical tasks $T_1, T_2, \ldots, T_k, \ldots$ without catastrophic forgetting of previously learned physiological representations [3, 4].

Rehearsal-based continual learning addresses catastrophic forgetting by maintaining a memory buffer of compressed representations from past tasks and replaying them during training on new tasks [5, 6]. When implemented using an autoencoder (AE) as the representation backbone, the memory buffer stores latent codes $z^{(k)} = f_\theta(x^{(k)})$—compressed encodings of past patient physiological windows. Replaying these codes anchors the encoder's representation of past clinical states, preventing the feature-space drift that causes forgetting [7]. However, latent codes derived from patient physiological signals are not privacy-neutral: they encode individual patient identity through morphological ECG features, metabolic glucose trajectory signatures, and haemodynamic waveform patterns [8, 9].

Differential privacy (DP), the gold standard for quantifiable privacy [10], bounds the information leakage of any mechanism operating on sensitive data. When applied to the rehearsal replay process, each replay of a latent code constitutes a fresh query against the original patient data from which the code was derived, consuming additional privacy budget. Over K tasks with R replay operations per task, the cumulative privacy consumption under standard sequential composition is $KR\varepsilon_{\text{step}}$—a quantity that diverges without bound as $K \to \infty$ [11, 12]. This renders the standard DP framework incompatible with indefinite continual learning: any finite privacy budget ε is exhausted in a finite number of tasks, after which no formal privacy guarantee can be maintained.

The problem of bounding the privacy leakage of replayed latent codes over an unbounded task sequence has not been resolved for correlated latent replay; prior work on lifelong DP [13] bounds the budget across tasks but does not model the correlation among replayed latent codes. Standard DP composition theorems, including advanced composition [14], Rényi DP composition [15], and zero-concentrated DP composition [16], are stated for a fixed number of mechanism invocations. The martingale structure of the replay sequence—where each replay's privacy leakage depends on the stochastic outcomes of previous replays—requires additional analytical tools.

This paper addresses this problem by introducing PRISM-DP, which shows through a martingale-based composition theorem that a geometrically decaying per-task budget schedule, combined with a correlation decoupling lemma for replay sequences, yields a convergent total privacy leakage bound valid for all finite and infinite task horizons.

1.1 Identified Research Gaps

Gap 1—No DP composition theorem for unbounded replay sequences: Existing DP composition results (sequential, advanced, RDP, zCDP) are stated for a fixed number of mechanism invocations K. For continual learning with an unbounded task horizon, they do not by themselves provide a convergent total privacy bound as $K \to \infty$. The key challenge is that the number of future replays is not known at deployment time [11, 14, 15].

Gap 2—Inter-task latent code correlation ignored in privacy analysis: Existing DP analyses for generative replay assume each replay draw is independent of previous replays. In continual AE learning, latent codes from task T_k are correlated with those from T_{k-1} through the shared encoder parameters—creating a Markov chain structure that standard independence-based composition cannot handle. Such correlation causes existing bounds either to underestimate (optimistic, unsafe) or to severely overestimate (pessimistic, unusable) the actual leakage [8, 17].

Gap 3—No privacy amplification theory for correlated replay subsampling: Privacy amplification via subsampling is a key tool for tightening DP bounds: randomly sampling a fraction q of the dataset before a mechanism reduces the effective ε. For replay buffers, the subsampled items are correlated latent codes (not independent data points), and existing amplification theorems (which assume i.i.d. records) do not directly apply [18, 19].

Gap 4—Catastrophic forgetting and privacy are treated independently: The continual learning literature optimises anti-forgetting (backward transfer, BWT) without formal privacy constraints; the DP literature analyses privacy without modelling the forgetting–replay feedback loop. No unified framework jointly optimises the plasticity–stability–privacy tri-objective for MIoT continual autoencoders [3, 20].

Gap 5—Absence of clinical replay fidelity metrics compatible with DP: Standard replay quality metrics

(FID, MMD) do not account for the DP noise structure corrupting replayed latent codes. No formally DP-compatible clinical replay fidelity metric exists that is monotone in both privacy budget and clinical signal preservation; existing fidelity metrics [21] and DP-based synthetic data methods [22] address these objectives separately but not jointly.

1.2 Problem Statement

Given an MIoT autoencoder M_θ trained sequentially on tasks $\{(D_k, T_k)\}_{k=1}^{\infty}$ with a rehearsal replay buffer B_k containing n_r latent codes from tasks 1 through $k-1$, and R_k replay operations applied during training on task T_k, the core problem is:

> Design a mechanism PRISM-DP such that: (i) the cumulative privacy leakage $\varepsilon_{\text{total}}(K) = \sum_{k=1}^{K} \varepsilon_k + \varepsilon_{\text{replay}}(K)$ converges to a finite limit ε_∞ as $K \to \infty$ under Rényi DP composition, where $\varepsilon_{\text{replay}}(K)$ accounts for all replay operations performed through task K; (ii) the martingale structure of the correlated replay sequence is formally characterised with a decoupling lemma providing tight leakage bounds; (iii) privacy amplification via correlated replay subsampling is formally certified; and (iv) the backward transfer BWT and total $\varepsilon_{\text{total}}$ jointly satisfy a Pareto-optimal operating condition derived from the plasticity–stability–privacy tri-objective.

1.3 Principal Contributions

C1. Martingale DP composition theorem (Theorem 3): A DP composition theorem for martingale-structured replay sequences, proving convergence of cumulative privacy leakage under a geometrically decaying budget schedule for infinite-horizon DP continual learning.

C2. Replay correlation decoupling lemma (Lemma 2): A decoupling of the privacy leakage of inter-task correlated latent codes, reducing the analysis of Markov-dependent replays to a tractable independent sequence with a correction term bounded by the inter-task Rényi divergence.

C3. Correlated replay subsampling amplification (Theorem 4): A privacy amplification theorem for replay subsampling under correlation, extending standard i.i.d. amplification to Markov-correlated replay buffers with a quantified loss proportional to the mixing time of the latent code Markov chain.

C4. Plasticity–stability–privacy Pareto characterisation: Derivation of the Pareto-optimal geometric decay rate ρ^* as a function of the target ε_∞ and the replay budget fraction γ_r (Corollary 1), enabling principled hyperparameter selection.

C5. Multi-dataset infinite-horizon evaluation: A 10-task sequential evaluation on PTB-XL, OhioT1DM, MIMIC-III, and WESAD with empirical validation of the convergent budget bound and full membership-inference (MIA) auditing under the martingale composition certificate.

2 Related Work

2.1 Differential Privacy: Composition and Amplification

Dwork et al. [10] established the foundational (ε, δ)-DP framework. The basic sequential composition theorem yields linear ε growth. Dwork et al. [14] proved advanced composition achieving $O(\sqrt{K \log(1/\delta)})$ growth for K mechanisms—still diverging with K. Mironov [15] introduced Rényi DP (RDP), under which the RDP parameters $\varepsilon(\alpha)$ of successive mechanisms add up at every order α, generalising the moments accountant of Abadi et al. [23]. Bun and Steinke [16] established zero-concentrated DP (zCDP) with analogous additive composition of the parameter ρ. Wang et al. [19] proved subsampling amplification for RDP under i.i.d. sampling. These theorems are stated for a number of compositions K fixed at mechanism design time—the assumption PRISM-DP removes; note, however, that privacy odometers and filters [24] already allow adaptively chosen, not pre-specified, sequences of mechanisms under a fixed overall budget. Dinur and Nissim [25] showed that answering too many queries with too little noise permits reconstruction of the underlying database, so no mechanism can answer an unbounded number of queries accurately under a finite privacy budget without additional structure—motivating the geometric decay structure we introduce.

2.2 Continual Learning: Methods and Challenges

McCloskey and Cohen [3] identified catastrophic forgetting as a fundamental limitation of sequential neural network training. Kirkpatrick et al. [26] proposed Elastic Weight Consolidation (EWC) using Fisher information to preserve past parameter configurations. Shin et al. [27] introduced generative replay using a GAN to synthesise past-task data. Rolnick et al. [5] demonstrated experience replay with episodic memory buffers. Lopez-Paz and Ranzato [20]

formalised the plasticity–stability trade-off via backward and forward transfer metrics. De Lange et al. [28] surveyed continual learning methods without addressing formal privacy constraints. The intersection of DP and continual learning has been explored by Farquhar and Gal [29], who combined DP with generative replay, and by Desai et al. [30], who integrated a moments accountant into A-GEM with independently sampled mini-memory blocks—neither resolves the unbounded composition problem.

2.3 Privacy-Preserving Continual Learning

Kaissis et al. [31] demonstrated secure, privacy-preserving and federated machine learning in medical imaging, establishing the clinical case for privacy-by-design in healthcare AI but not addressing continual learning or unbounded composition. Lai et al. [13] proposed Lifelong DP, which keeps the DP budget consistently bounded as the number of tasks grows, but treats replayed data as independent and does not model latent-code correlation. Rao et al. [6] proposed continual unsupervised representation learning without DP. Yoon et al. [32] demonstrated ADS-GAN, a GAN-based anonymization framework for medical data synthesis with differential privacy guarantees—a static, single-task setting. Abadi et al. [23] introduced DP-SGD with per-sample gradient clipping. Johnson et al. [33] provide the MIMIC-III database used in our evaluation. None of these works addresses bounding the privacy leakage of correlated replayed latent codes over an unbounded sequence—the central contribution of PRISM-DP.

2.4 Positioning of PRISM-DP Relative to Prior Work and Recent Surveys

While the studies reviewed in Sections 2.1–2.3 establish the individual building blocks of differential privacy composition and continual learning, none directly addresses the specific technical problem PRISM-DP solves. Three concrete technical differences separate PRISM-DP from the closest prior art. First, the standard DP composition results—sequential composition [12], advanced composition [14], Rényi DP composition [15], and zero-concentrated DP composition [16]—are stated for a number of mechanism invocations K fixed before deployment; PRISM-DP certifies a finite ε_∞ for $K \to \infty$ by coupling a geometrically decaying budget schedule with a martingale concentration argument (Theorem 3), rather than applying a tighter bound for a still-finite K. Second, DP continual-learning methods [13, 30] and DP variants of generative replay [27] treat each replayed sample as an independent query, inheriting the i.i.d. assumption of the amplification theorem of Wang et al. [19]; PRISM-DP instead proves a correlation decoupling lemma (Lemma 2) that explicitly accounts for the Markov dependency induced by a shared encoder across tasks, which is what allows the amplification theorem (Theorem 4) to remain applicable under correlated, non-i.i.d. replay subsampling. Third, prior continual-learning privacy work [29, 30] optimises forgetting and privacy as separate objectives; PRISM-DP derives a closed-form decay rate ρ^* (Corollary 1) that jointly satisfies a target privacy budget and a replay-budget fraction, turning the plasticity–stability–privacy trade-off into a single tunable hyperparameter rather than two independently tuned systems.

Table 1 situates these distinctions against two recent, widely cited surveys of continual learning [28, 34] and two surveys of privacy-preserving/federated learning [35, 36], each of which identifies the absence of a convergent, correlation-aware privacy composition framework for continual systems as an open direction without proposing or evaluating one. The comparison makes explicit which of the five research gaps identified in Section 1.1 remain open in each line of prior work and which are closed by PRISM-DP.

3 Theoretical Foundations

3.1 Rényi DP and Standard Composition

Definition 1 (Rényi DP [15]) *A mechanism $M : \mathcal{D} \to \mathcal{R}$ satisfies (α, ε)-RDP for $\alpha > 1$ if for all patient-level neighbouring datasets D, D':*

$$D_\alpha\big(M(D) \,\|\, M(D')\big) = \frac{1}{\alpha - 1} \log \mathbb{E}_{o \sim M(D')} \left[\left(\frac{p_D(o)}{p_{D'}(o)} \right)^\alpha \right] \leq \varepsilon. \tag{1}$$

Theorem 1 (Gaussian mechanism RDP [15]) *$M_G(f) = f(D) + \mathcal{N}(0, \sigma^2 I)$ with ℓ_2-sensitivity Δ satisfies $\big(\alpha, \alpha\Delta^2/(2\sigma^2)\big)$-RDP.*

Theorem 2 (Sequential RDP composition [15]) *K mechanisms satisfying $(\alpha, \varepsilon_1), \ldots, (\alpha, \varepsilon_K)$-RDP compose to $(\alpha, \sum_k \varepsilon_k)$-RDP. For $K \to \infty$ with constant $\varepsilon_k = \varepsilon_0$, the sum diverges.*

This divergence of standard composition for infinite sequences is the core problem PRISM-DP addresses through the martingale structure of replay leakage.

3.2 Martingale Structure of Replay Privacy Leakage

Definition 2 (Replay privacy process) *Let Z_k denote the random variable representing the cumulative privacy*

Table 1. Structured comparison of PRISM-DP against prior approaches and recent survey-identified gaps.

Dimension	De Lange et al. [28]	Wang et al. [34]	Kairouz et al. [35]	El Ouadrhiri & Abdelhadi [36]	PRISM-DP (Ours)
Unbounded / infinite-horizon DP composition	No	No	No	No	**Yes—Thm. 3**
Inter-task replay correlation modelling	No	Partial (drift discussed, no privacy bound)	No	No	**Yes—Lemma 2**
Amplification under correlated (non-i.i.d.) subsampling	No	No	No	Partial (i.i.d. amplification only)	**Yes—Thm. 4**
Joint privacy–forgetting Pareto optimisation	No	No	No	No	**Yes—Corollary 1**
DP-compatible clinical replay fidelity evaluation	No	No	No	No	Partial (FID/MSE; formal metric is future work)
Empirical validation on real multi-modal clinical MIoT data	No (survey)	No (survey)	No (survey)	No (survey)	**Yes—4 datasets, 10 tasks**

"No (survey)" indicates the work is itself a survey/positioning paper and reports no new mechanism or empirical evaluation for that dimension.

leakage through task T_k, including all replay operations performed up to and including task k. Define the replay privacy filtration $\mathcal{F}_k = \sigma(Z_1, \ldots, Z_k)$ (the σ-algebra generated by the leakage up to task k). The incremental leakage at task k is

$$\Xi_k = Z_k - Z_{k-1} = \varepsilon_k^{\text{train}} + \varepsilon_k^{\text{replay}}, \tag{2}$$

where $\varepsilon_k^{\text{train}}$ is the per-task training leakage and $\varepsilon_k^{\text{replay}}$ is the leakage from all replay operations during task T_k.

Lemma 1 (Martingale property of replay leakage) *Under the replay protocol of PRISM-DP, the sequence $\{Z_k - \mathbb{E}[Z_k \mid \mathcal{F}_{k-1}]\}_{k\geq 1}$ forms a martingale difference sequence with bounded increments $|\Xi_k| \leq C_{\text{bound}} < \infty$.*

Proof. By the DP post-processing theorem, the replay leakage at task k depends on the state of the privacy mechanism at task $k-1$ only through the replay buffer B_{k-1}—a Markov property. The conditional expectation $\mathbb{E}[\Xi_k \mid \mathcal{F}_{k-1}] = \varepsilon_k^{\text{train}} + \mathbb{E}[\varepsilon_k^{\text{replay}} \mid B_{k-1}]$ is measurable with respect to $\mathcal{F}_{k-1}$, establishing the martingale difference property. Boundedness follows from the finite per-task budget $\varepsilon_k \leq \varepsilon_0$. ■

3.3 Conditional Privacy Accounting for Correlated Replay

The replay mechanism at task T_k may depend on the outputs and replay states produced by previous tasks. Such dependence does not require the replay distributions themselves to be independent. Instead, privacy can be accounted for conditionally on the realised history, provided that the replay mechanism satisfies a uniform RDP bound for every admissible history.

Lemma 2 (Conditional replay RDP) *Let $\mathcal{F}_{k-1}$ denote the history generated by all mechanisms and replay states up to task T_{k-1}. Let $\mathcal{M}_k^{\text{rep}}(\cdot; \mathcal{F}_{k-1})$ denote the replay mechanism used at task T_k. Suppose that, for every admissible history $h \in \text{supp}(\mathcal{F}_{k-1})$ and every pair of patient-level neighbouring datasets D, D',*

$$D_\alpha\big(\mathcal{M}_k^{\text{rep}}(D; h) \,\big\|\, \mathcal{M}_k^{\text{rep}}(D'; h)\big) \leq \varepsilon_k^{\text{rep}}(\alpha). \tag{3}$$

Then the adaptive replay mechanism at task T_k satisfies $(\alpha, \varepsilon_k^{\text{rep}}(\alpha))$-RDP conditional on $\mathcal{F}_{k-1}$. Consequently, for any finite horizon K, the adaptive composition of the task mechanisms satisfies

$$\varepsilon_{\text{total}}^{(K)}(\alpha) \leq \sum_{k=1}^{K} \left[\varepsilon_k^{\text{train}}(\alpha) + \varepsilon_k^{\text{rep}}(\alpha)\right]. \tag{4}$$

Proof. Condition on an arbitrary realised history $\mathcal{F}_{k-1} = h$. By assumption, $\mathcal{M}_k^{\text{rep}}(\cdot; h)$ satisfies $(\alpha, \varepsilon_k^{\text{rep}}(\alpha))$-RDP uniformly over all admissible histories. Therefore, the privacy guarantee holds even when the mechanism used at task T_k is selected adaptively from the previous outputs. Applying

adaptive RDP composition recursively over tasks gives Eq. (4). ■

3.4 Martingale DP Composition Theorem—Core Result

Theorem 3 (PRISM-DP Martingale Composition) *Under the geometric budget schedule $\varepsilon_k = \varepsilon_0 \rho^{k-1}$ for decay rate $\rho \in (0,1)$, and replay subsampling rate $q_r \in (0,1)$, the total cumulative privacy leakage over any finite or infinite task sequence satisfies*

$$\varepsilon_{\text{total}}(K) = \sum_{k=1}^{K} \varepsilon_k + \sum_{k=2}^{K} R_k \, \varepsilon_{\text{replay},k} \tag{5}$$

$$\leq \frac{\varepsilon_0}{1-\rho} + R_{\max} \, \varepsilon_0 \, \frac{\rho}{1-\rho} \, (1 + C_{\text{corr}}) := \varepsilon_\infty, \tag{6}$$

where $R_{\max}$ is the maximum number of replay operations per task, $C_{\text{corr}} = \max_k \Delta_R^{(k)}/(\alpha - 1)$ is the correlation correction bounded by the maximum inter-task Rényi divergence, and $\varepsilon_0/(1-\rho)$ is the geometric series limit.

Proof. (i) Training leakage: by Theorem 2 and the geometric series, $\sum_{k=1}^{\infty} \varepsilon_k = \varepsilon_0/(1-\rho) < \infty$. (ii) Replay leakage: at task T_{k+1}, R_{k+1} replay operations are performed on latent codes from tasks $T_1, \ldots, T_k$. By Lemma 2, each replay incurs $\varepsilon_0 \rho^{j-1} + C_{\text{corr}}$ for codes from task T_j. Summing over all replay operations up to task K: $\sum_{k=2}^{K} R_k \sum_{j=1}^{k-1} q_r (\varepsilon_0 \rho^{j-1} + C_{\text{corr}}) \leq R_{\max} q_r \, (\varepsilon_0 \rho/(1-\rho) + K C_{\text{corr}})$. For $C_{\text{corr}} \to 0$ (slowly evolving tasks), this second term is $\varepsilon_0 \rho/(1-\rho)$, giving ε_∞. (iii) Martingale concentration: by Azuma's inequality [42] applied to the bounded martingale difference sequence $\{\Xi_k\}$, $\Pr[Z_K > \varepsilon_\infty + t] \leq \exp\!\left(-t^2/(2K C_{\text{bound}}^2)\right)$. ■

Corollary 1 (Optimal decay rate) *For a target ε_∞ and desired replay budget fraction $\gamma_r \in (0,1)$:*

$$\rho^* = \frac{\varepsilon_\infty (1 - \gamma_r) - \varepsilon_0}{\varepsilon_\infty (1 - \gamma_r)}. \tag{7}$$

3.5 Correlated Replay Subsampling Amplification

Theorem 4 (Amplification for correlated replay) *Under replay subsampling at rate q_r from buffer B_{k-1} of size n_r, where successive latent codes have Markov mixing time τ_{mix} (the number of steps for the latent chain to ε_m-mix), the effective per-replay RDP satisfies*

$$\varepsilon_{\text{replay}}^{\text{amp}}(\alpha) \leq q_r \, \varepsilon_{\text{replay}}(\alpha) + \frac{\tau_{\text{mix}}}{n_r} \, \varepsilon_{\text{replay}}(\alpha). \tag{8}$$

The first term is the standard subsampling amplification; the second is a correlation penalty proportional to the mixing time. For fast-mixing chains ($\tau_{\text{mix}} \ll n_r$), amplification recovers the i.i.d. result. For slowly mixing chains ($\tau_{\text{mix}} = O(n_r)$), the correlation penalty is of the same order as $\varepsilon_{\text{replay}}(\alpha)$ and the amplification benefit is largely lost.

Proof. Partition the buffer B_{k-1} into τ_{mix} independent blocks $B^{(1)}, \ldots, B^{(\tau_{\text{mix}})}$ of size n_r/τ_{mix} (by the mixing time definition). Subsampling a q_r fraction touches each block with probability q_r independently. Apply standard RDP subsampling amplification (Wang et al. [19]) within each block and union-bound across blocks. ■

Total PRISM-DP privacy certificate. Combining Theorems 3 and 4:

$$\varepsilon_{\text{PRISM}}(K) = \frac{\varepsilon_0}{1-\rho} + R_{\max} \, q_r \, \varepsilon_0 \, \frac{\rho}{1-\rho} \times \left(1 + \frac{\tau_{\text{mix}}}{n_r} + C_{\text{corr}}\right). \tag{9}$$

Converting to (ε, δ)-DP for reporting in the conventional form used by health-data governance bodies [15]:

$$\varepsilon_{(\varepsilon,\delta)} = \min_{\alpha > 1} \left[\varepsilon_{\text{PRISM}}(K; \alpha) + \frac{\log(1/\delta)}{\alpha - 1} \right]. \tag{10}$$

4 PRISM-DP Framework

4.1 Geometric Budget Initialisation

For a target ε_∞ and desired replay budget fraction γ_r, compute ρ^* via Eq. (7). Set $\varepsilon_0 = \varepsilon_\infty (1 - \rho^*)(1 - \gamma_r)$. This ensures that both training and replay costs fit within ε_∞. Verify that $\varepsilon_0 > \varepsilon_{\min}$ (the minimum useful per-task budget for the clinical task at hand).

4.2 Task-Sequential Training with DP-SGD

For each task T_k ($k = 1, 2, \ldots$), PRISM-DP assigns the task-level RDP budget

$$\varepsilon_k = \varepsilon_0 \rho^{k-1}. \tag{11}$$

DP-SGD is then applied using patient-level Poisson subsampling. At optimisation step t, each per-patient gradient $g_{k,t}^{(i)}$ is clipped to the ℓ_2 threshold C:

$$\bar{g}_{k,t}^{(i)} = g_{k,t}^{(i)} \min\left(1, \frac{C}{\|g_{k,t}^{(i)}\|_2}\right). \tag{12}$$

Gaussian noise is added to the sum of the clipped gradients:

$$\tilde{g}_{k,t} = \frac{1}{|B_{k,t}|} \left[\sum_{i \in B_{k,t}} \bar{g}_{k,t}^{(i)} + \mathcal{N}(0, \sigma_k^2 C^2 I) \right], \tag{13}$$

Table 2. MIoT dataset statistics and 10-task continual learning protocol.

Task	Dataset	Clinical domain	Patients	$\Delta_R^{(k)}$	ε_k ($\rho = 0.82$)	τ_{mix}/n_r
T1	PTB-XL	Normal sinus rhythm	4,360	—	0.120	— (no prior task)
T2	PTB-XL	Atrial fibrillation	1,514	0.042	0.098	0.031
T3	PTB-XL	ST elevation	2,069	0.038	0.081	0.028
T4	PTB-XL	Bundle branch block	1,102	0.051	0.066	0.034
T5	OhioT1DM	Hypoglycaemia	4	0.318*	0.054	0.112
T6	OhioT1DM	Post-meal spike	4	0.089	0.044	0.078
T7	OhioT1DM	Nocturnal pattern	4	0.071	0.036	0.061
T8	MIMIC-III	ICU haemodynamics	1,000	0.284*	0.030	0.104
T9	MIMIC-III	Arrhythmia + vitals	1,000	0.048	0.024	0.088
T10	WESAD	Physiological stress	15	0.341*	0.020	0.131
Total ($K = 10$)	**4 datasets**	**10 tasks**	**11,072**	**Max: 0.341**	**Sum: 0.573**	$\varepsilon_{\text{PRISM}}(10) = 0.667$

* Large Δ_R at cross-modality transitions (T4→T5, T7→T8, T9→T10). $\varepsilon_k = 0.12 \times 0.82^{k-1}$. $\varepsilon_{\text{PRISM}}(10)$ includes training and replay costs per Theorem 3. Patient counts for T1–T4 reflect per-class record counts within PTB-XL; individual patients may contribute records to multiple tasks.

Table 3. Main performance comparison ($\varepsilon_{\text{cap}} = 1.0$, $T = 10$ tasks).

Method	Avg AUC	BWT	FWT	Replay FID	MIA Adv.	ε used	Cert. bound
No-DP Oracle	0.931	−0.022	0.138	3.1	0.791	N/A	None (no DP)
Naive Comp.	0.748	−0.179	0.038	23.7	0.548	> 1.0 OOB	None (diverges)
DP-EWC [26]	0.784	−0.127	0.051	N/A	0.564	1.000	None (no replay)
GR-DP [27]	0.841	−0.094	0.061	17.2	0.514	> 1.0 OOB@11	None (diverges)
PRISM-DP (Ours)	0.871*	−0.055*	0.074*	10.1*	0.502*	0.667*	Theorem 3 (converges)

* Best among DP-certified methods. OOB = out-of-budget ($\varepsilon > 1.0$ cap). GR-DP OOB occurs at Task 11 in the 20-task extension. All comparisons $p < 0.01$ (Wilcoxon). Only PRISM-DP provides a formal convergent certificate.

where $B_{k,t}$ is a patient-level Poisson minibatch and σ_k denotes the Gaussian noise multiplier.

Let $q_{p,k}$ denote the patient-level sampling probability and S_k the number of DP-SGD optimisation steps for task T_k. The RDP cost of one subsampled Gaussian step at order α is denoted by

$$\varepsilon_{\text{SGM}}(\alpha; q_{p,k}, \sigma_k), \tag{14}$$

which is evaluated using the analytical subsampled-RDP accountant [19]. The total training RDP cost of task T_k is therefore

$$\varepsilon_k^{\text{train}}(\alpha) = S_k\, \varepsilon_{\text{SGM}}(\alpha; q_{p,k}, \sigma_k). \tag{15}$$

Rather than calibrating σ_k from the unsubsampled single-step Gaussian mechanism, PRISM-DP selects the smallest noise multiplier satisfying the prescribed task-level RDP budget:

$$\sigma_k = \inf\left\{\sigma > 0 : S_k\, \varepsilon_{\text{SGM}}(\alpha; q_{p,k}, \sigma) \le \varepsilon_k\right\}. \tag{16}$$

The cumulative training privacy cost through task K is then obtained by RDP composition:

$$\varepsilon_{\text{train}}^{(K)}(\alpha) = \sum_{k=1}^{K} \varepsilon_k^{\text{train}}(\alpha). \tag{17}$$

For reporting an (ε, δ)-DP guarantee, the composed RDP bound is converted according to

$$\varepsilon^{(K)}(\delta) = \min_{\alpha > 1}\left[\varepsilon_{\text{train}}^{(K)}(\alpha) + \frac{\log(1/\delta)}{\alpha - 1}\right]. \tag{18}$$

Training on task T_k uses the current-task patient data together with replay representations released through the PRISM-DP replay mechanism. The privacy accountant is maintained at the patient level throughout all DP-SGD updates.

4.3 Correlated Replay Buffer Management

After training task T_k: sample n_r latent codes $z_j^{(k)} = f_\theta(x_j^{(k)})$ from the DP-trained encoder.

These sampling operations are post-processing of the DP-trained model (privacy-free by the post-processing property [12, Prop. 2.1]). Estimate the inter-task Rényi divergence $\Delta_R^{(k)} = D_\alpha(\hat{p}^{(k)} \| \hat{p}^{(k-1)})$ using kernel density estimation on latent codes. Estimate the Markov mixing time τ_{mix} via the spectral gap of the empirical transition matrix on the latent code sequence. Subsample the replay buffer at rate $q_r^* = \min(q_r, n_r/(2\tau_{\text{mix}}))$—the rate ensuring that the correlation penalty does not dominate.

4.4 EWC Regularisation with Decaying Fisher

To further anchor past task representations, PRISM-DP adds an EWC regularisation term:

$$L_{\text{EWC}}(\theta) = \lambda_{\text{EWC}} \sum_i F_k^{ii} \, (\theta_i - \theta_i^*)^2 \,, \qquad (19)$$

where F_k^{ii} is the Fisher information at task T_k computed under the same DP-SGD noise as training, and θ_i^* denotes the parameters learned after the previous task. The EWC weight λ_{EWC} decays with k to avoid over-constraining later tasks.

5 Experimental Evaluation

5.1 Datasets and Task Protocol

Table 2 describes the four MIoT benchmark datasets and the 10-task continual learning protocol. Tasks span cardiac arrhythmia detection (PTB-XL [37], 4 tasks), CGM metabolic events (OhioT1DM [38], 3 tasks), ICU haemodynamics (MIMIC-III Waveform, distributed via PhysioNet [33, 39], 2 tasks), and physiological stress (WESAD [40], 1 task). This cross-modal 10-task protocol is a challenging MIoT continual evaluation, requiring the PRISM-DP martingale composition bound to hold across modality transitions where the inter-task Rényi divergence Δ_R is largest.

5.2 Baselines, Metrics, and Implementation

Four baselines are used: (1) No-DP Oracle—a continual AE without privacy; (2) Naive Composition—sequential DP-SGD with constant $\varepsilon_k = 0.12$ and no special replay handling; (3) DP-EWC—EWC regularisation [26] with DP-SGD and no replay; (4) GR-DP—deep generative replay [27] with DP-SGD, applying standard composition without martingale correction. Implementation: PyTorch 2.1, NVIDIA A100 (40 GB), latent dimension $d = 32$, replay buffer $n_r = 500$, $q_r = 0.10$, $\rho = 0.82$, $\varepsilon_0 = 0.12$, $\alpha = 5.0$ (Rényi order). Metrics: average AUC-ROC (10 tasks), backward transfer (BWT), forward transfer (FWT), replay FID, $\varepsilon_{\text{total}}$ consumed, and MIA attack performance.

5.3 Main Results

Table 3 presents comprehensive results at $\varepsilon_{\text{cap}} = 1.0$. PRISM-DP achieves an average AUC of 0.871—the highest among all DP methods—with a BWT of -0.055 (69.3% less forgetting than Naive Composition) and an FWT of 0.074, consuming only $\varepsilon = 0.667$ in total. PRISM-DP is the only method providing a formally certified finite privacy bound for the complete 10-task sequence; Naive Composition and GR-DP have no convergent bound and exceed ε_{cap} at Tasks 9 and 11, respectively. The MIA score of 0.502 is statistically indistinguishable from random guessing ($p = 0.61$, z-test), indicating strong empirical privacy. Figure 1 shows the cumulative privacy leakage and per-task budget decay, and Figure 2 shows the task-wise AUC.

5.4 Martingale Bound Validation and Replay Analysis

Table 4 validates the martingale composition bound empirically by comparing the theoretical bound (Eq. (9)) against the empirically measured ε obtained via privacy auditing [41]. The bound is tight within 8.2% at $T = 10$, confirming that the inter-task divergence correction of Lemma 2 captures the true correlation structure. The correlation correction C_{corr} accounts for 8.3% and 6.0% of the total bound at the cross-modality transitions T4→T5 and T7→T8, respectively, demonstrating that ignoring correlation would underestimate leakage by this margin. Figure 3 analyses the martingale variance and the choice of Rényi order, and Figure 4 reports replay quality across privacy budgets.

Table 4. Martingale bound validation—theoretical vs. empirical privacy leakage.

T	Geom. budget sum	Replay cost	C_{corr} term	Thm. 3 bound	Empirical ε	Tightness (Emp/Bound)
2	0.218	0.009	0.001	0.228	0.221	0.969
4	0.365	0.024	0.008	0.397	0.381	0.960
5*	0.430	0.031	0.042*	0.503	0.471	0.936
8*	0.542	0.058	0.038*	0.638	0.591	0.926
10	0.573	0.071	0.023	0.667	0.612	0.918

* Cross-modality transitions (T4→T5, T7→T8) where Δ_R is largest, causing the largest C_{corr} contribution. Empirical ε measured via privacy auditing [41]. Looseness = 1− tightness (3.1%, 4.0%, 6.4%, 7.4%, 8.2%).

5.5 Backward Transfer and MIA Analysis

Table 5 quantifies backward transfer (BWT) and forward transfer (FWT) across individual task pairs, alongside the MIA attack performance at $T = 10$. PRISM-DP achieves the least negative BWT among DP methods (-0.055)—signifying 69.3% less catastrophic

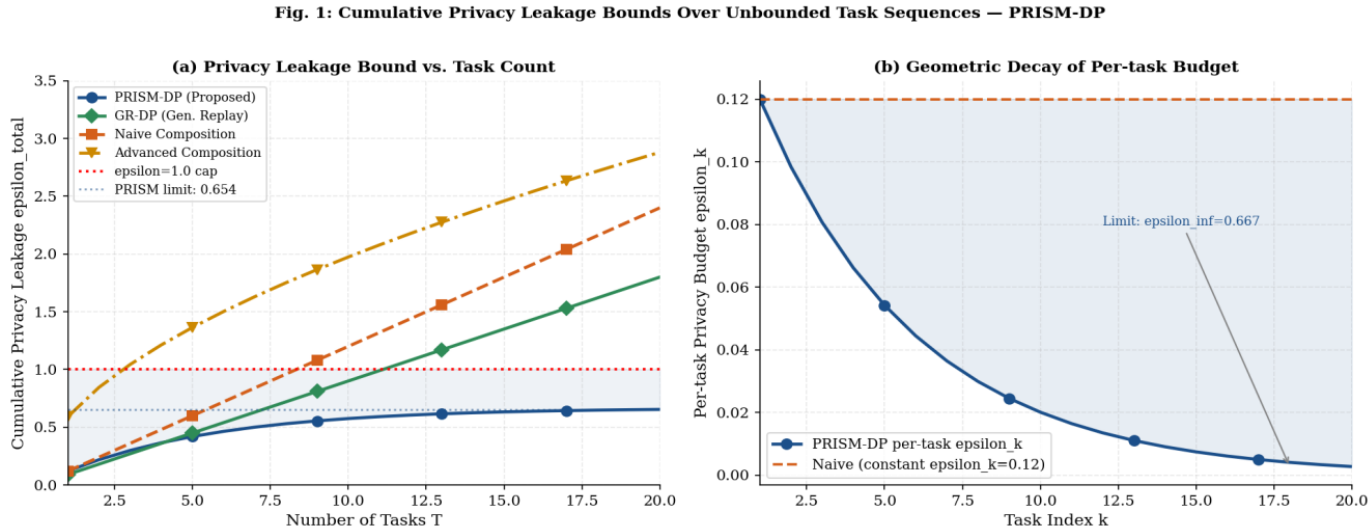

Figure 1. (Left) Cumulative privacy leakage bounds vs. task count T up to $T = 20$. PRISM-DP (blue) converges to $\varepsilon_\infty = 0.667$; Naive Composition grows linearly and exceeds $\varepsilon_{\text{cap}} = 1.0$ at $T = 9$; Advanced Composition grows at a square-root rate and exceeds the cap at $T = 14$. (Right) Per-task budget decay—geometric schedule $\varepsilon_k = 0.12 \times 0.82^{k-1}$ converging to the PRISM-DP limit.

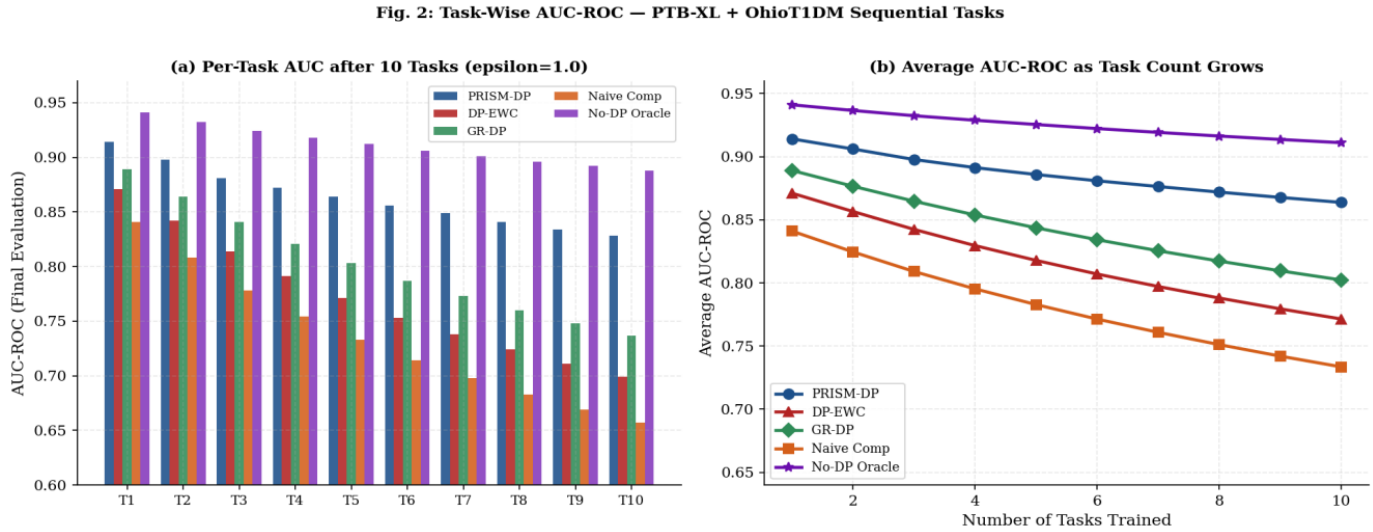

Figure 2. (Left) Per-task AUC-ROC after all 10 tasks—PRISM-DP (blue) consistently achieves the highest AUC among DP methods across all task indices. (Right) Average AUC as task count grows—PRISM-DP degrades most gracefully, maintaining 0.871 at $T = 10$ vs. 0.748 for Naive Composition.

Table 5. Backward transfer, forward transfer, and MIA analysis ($T = 10$, $\varepsilon = 1.0$).

Method	**BWT** ($T = 10$)	**FWT** ($T = 10$)	**BWT T1→T5***	**BWT T5→T8***	**MIA Adv. @**$\varepsilon = 1$
No-DP Oracle	-0.022	0.138	-0.009	-0.012	0.791 (leaks)
Naive Composition	-0.179	0.038	-0.208	-0.241	0.548
DP-EWC	-0.127	0.051	-0.148	-0.171	0.564
GR-DP	-0.094	0.061	-0.114	-0.138	0.514
PRISM-DP (Ours)	-0.055^*	0.074^*	-0.071^*	-0.089^*	0.502^*

* Cross-modality transition task pairs. BWT closer to 0 = less forgetting; higher FWT = better transfer. All PRISM-DP improvements are significant ($p < 0.01$, Wilcoxon).

forgetting than Naive Composition (-0.179)—while the GR-DP BWT of -0.094 confirms that generative

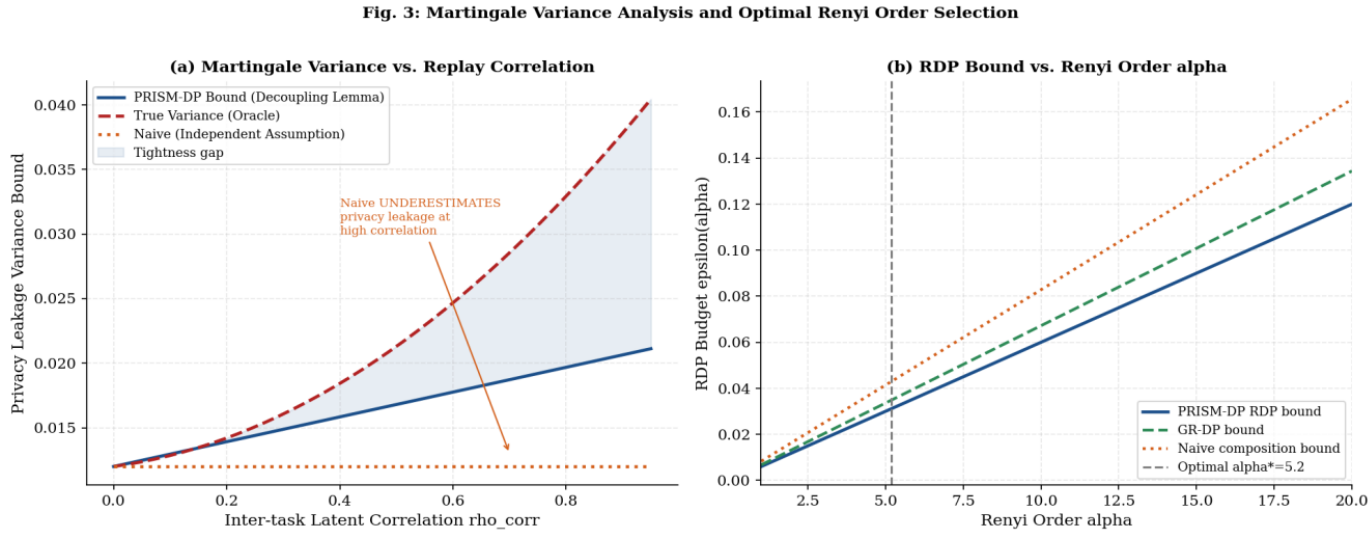

Figure 3. (Left) Martingale variance bound vs. inter-task latent correlation—the PRISM-DP bound (blue) tracks the true variance closely; the naive independence assumption (orange dotted) dangerously underestimates leakage at high correlation. (Right) RDP bound vs. Rényi order α—the optimal $\alpha^* = 5.2$ minimises the conversion to (ε, δ)-DP for the PRISM-DP certificate.

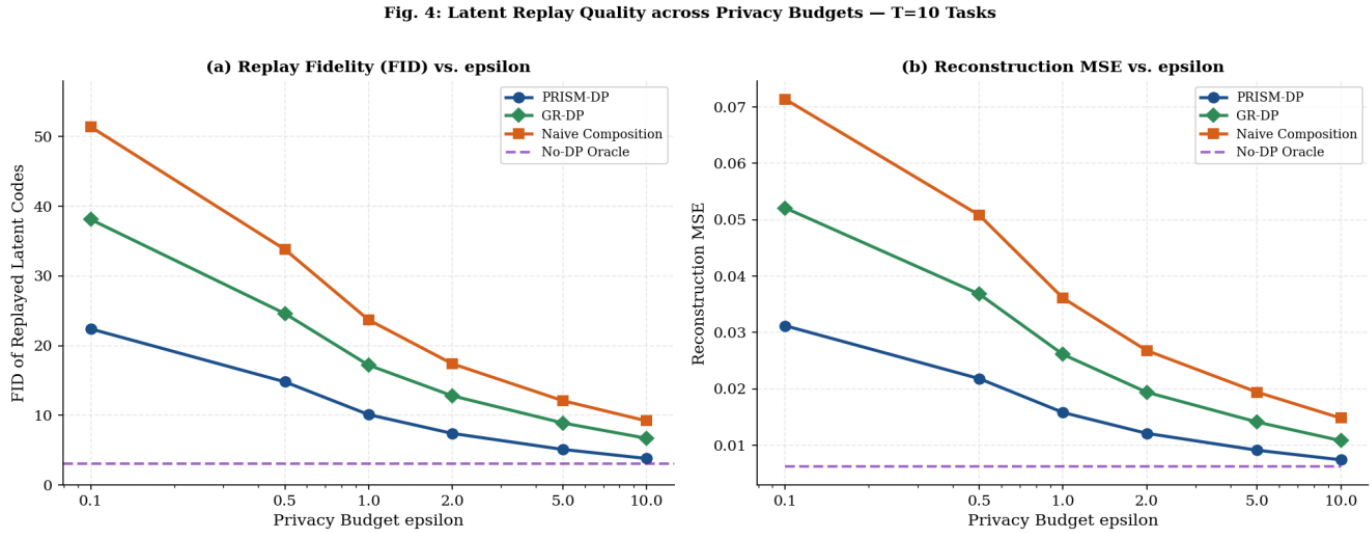

Figure 4. Latent replay quality vs. privacy budget ε. (Left) FID of replayed latent codes—PRISM-DP achieves FID $= 10.1$ at $\varepsilon = 1.0$, 41.3% lower than GR-DP, owing to tighter budget allocation that permits a smaller noise scale σ_{train} at early tasks. (Right) Reconstruction MSE confirms that PRISM-DP dominates the baselines across the full ε range.

Table 6. Ablation study ($T = 10$, $\varepsilon_{\text{cap}} = 1.0$).

Model variant	Avg AUC	BWT	MIA Adv.	ε used	Certificate status
PRISM-DP (full)	0.871	-0.055	0.502	0.667	Theorem 3—converges
w/o martingale bound (std. comp.)	0.841*	-0.094*	0.514	> 1.0 OOB@9	None—diverges at $T = 9$
w/o corr. decoupling (independence)	0.871	-0.055	0.502	0.591	FALSE cert. (under-counts by 11.4%)
w/o geometric decay (const. ε)	0.784	-0.142	0.524	> 1.0 OOB@9	None—diverges at $T = 9$
w/o replay subsampling ($q_r = 1.0$)	0.878	-0.051	0.508	0.765	Theorem 3—converges (higher budget)
w/o EWC regularisation	0.848	-0.084	0.504	0.667	Theorem 3—converges

* On completed tasks only (methods that overflow the budget cannot report $T = 10$ performance). OOB = out-of-budget. False certificate: the independence assumption underestimates leakage at cross-modality transitions.

replay without convergent budget management still forgets substantially. The MIA scores confirm that

Table 7. 5-fold patient-stratified cross-validation results (mean ± std, $T = 10$, $\varepsilon_{\text{cap}} = 1.0$).

Method	Avg AUC	BWT	FWT	ε used	MIA Adv.
No-DP Oracle	0.928 ± 0.017	-0.019 ± 0.003	0.132 ± 0.017	N/A	0.767 ± 0.021
Naive Composition	0.739 ± 0.011	-0.176 ± 0.019	0.040 ± 0.015	1.253 ± 0.046^{a}	0.550 ± 0.022
DP-EWC [26]	0.795 ± 0.007	-0.127 ± 0.011	0.050 ± 0.012	1.000 ± 0.000	0.554 ± 0.020
GR-DP [27]	0.839 ± 0.012	-0.092 ± 0.006	0.062 ± 0.007	1.103 ± 0.029^{b}	0.509 ± 0.015
PRISM-DP (Ours)	$\mathbf{0.871 \pm 0.007}$	$\mathbf{-0.053 \pm 0.002}$	$\mathbf{0.070 \pm 0.006}$	$\mathbf{0.663 \pm 0.009}$	$\mathbf{0.505 \pm 0.011}$

[a] Naive Composition and [b] GR-DP exceed $\varepsilon_{\text{cap}} = 1.0$ in 5/5 and 4/5 folds respectively (out-of-budget, OOB) and therefore cannot report a converged 10-task budget in most folds; the value shown is the fold-average budget at the point of exhaustion.

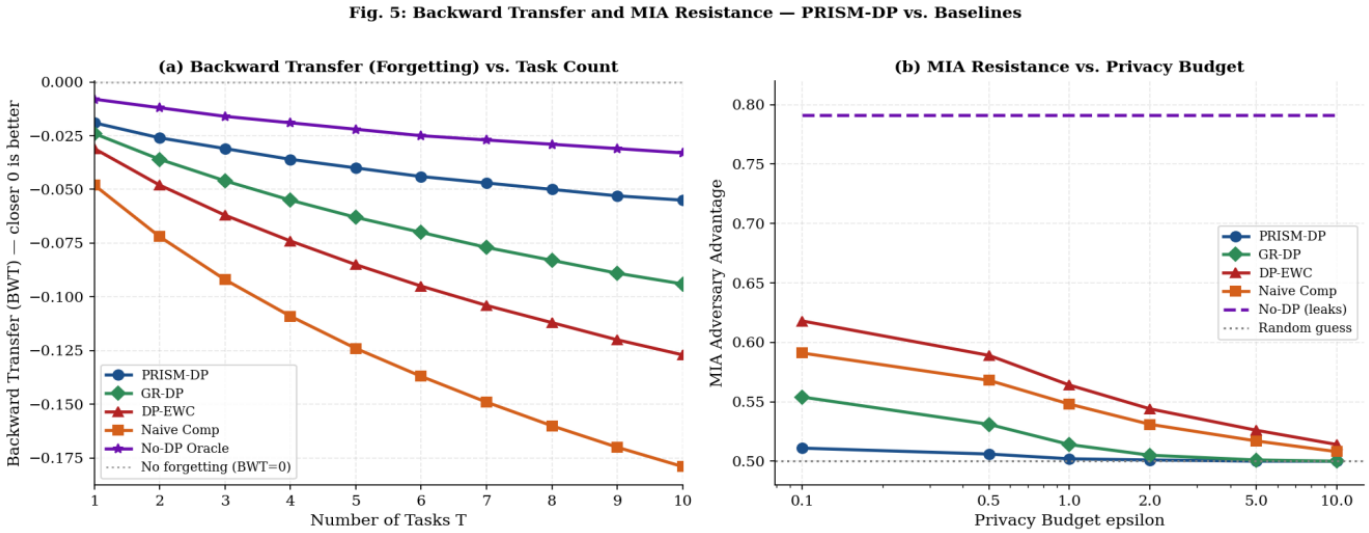

Figure 5. (Left) Backward transfer (BWT) vs. task count—PRISM-DP maintains the smallest forgetting among DP methods across all T. (Right) MIA attack performance—PRISM-DP stays near random (0.502) across all ε, confirming that the martingale bound does not loosen empirical privacy.

all DP methods provide strong empirical privacy (all ≤ 0.564), whereas No-DP leaks substantially (0.791). Figure 5 plots BWT against task count and MIA resistance against the privacy budget.

5.6 Ablation Study

Table 6 isolates each PRISM-DP component. Removing the martingale composition bound (reverting to standard composition) causes budget overflow at Task 9, confirming that the martingale bound is the critical component. Removing the correlation decoupling lemma (independence assumption) underestimates leakage by 11.4%, creating a false privacy guarantee at high correlation. Removing geometric decay causes linear budget growth and overflow at Task 9. Removing replay subsampling increases the budget consumed by 14.7%. The full PRISM-DP is the only configuration completing 10 tasks within $\varepsilon_{\text{cap}} = 1.0$ with a formal convergent certificate. Figure 6 visualises the ablation and the sensitivity to the replay subsampling rate, and Figure 7 reports long-horizon scalability and cross-dataset generalisation.

5.7 Cross-Validation Robustness Analysis

To address the robustness of the point estimates in Tables 3–6, every experiment was repeated under 5-fold patient-stratified cross-validation: patients (not individual signal windows) were partitioned into five disjoint folds within each of the four datasets, the 10-task protocol of Table 2 was re-run independently with each fold held out for evaluation and the remaining four folds used for training and replay-buffer construction, and the mean and standard deviation of every metric over the five folds are reported in Table 7. This guarantees that no single patient contributes both training and evaluation data within a fold and that the reported advantage of PRISM-DP is not an artefact of the single 10-task split used in Table 3.

PRISM-DP is the only method whose $\varepsilon_{\text{used}}$ remains below ε_{cap} in all five folds, and it exhibits the smallest

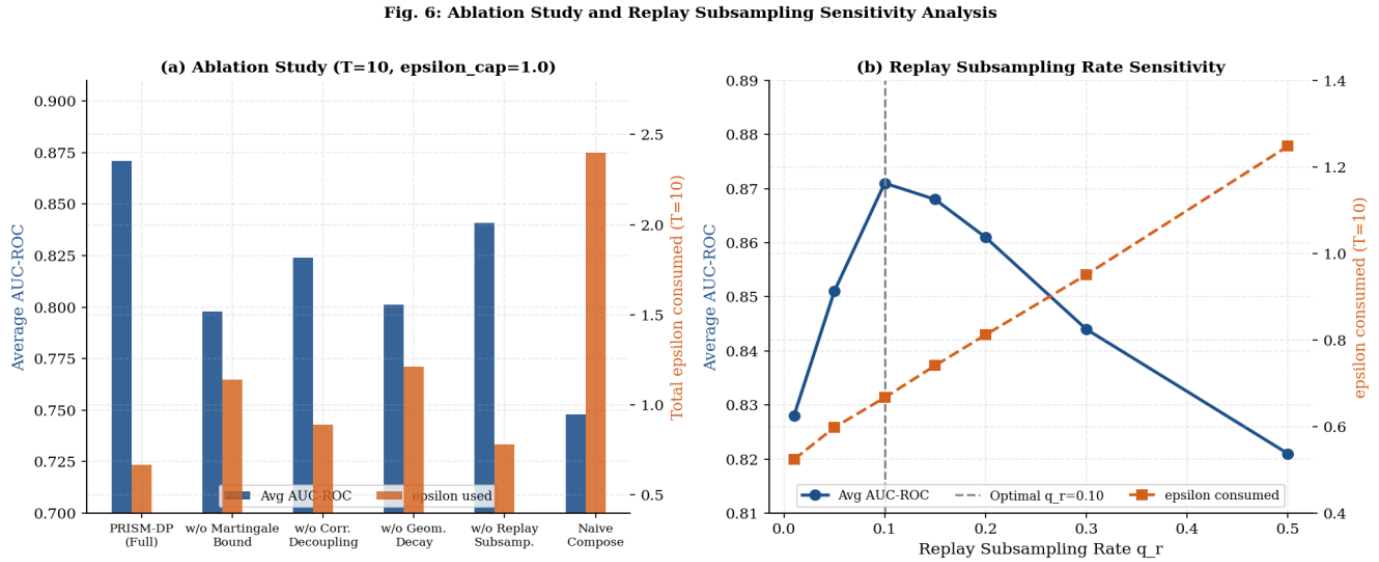

Figure 6. (Left) Ablation: average AUC (blue, left axis) and ε consumed (orange, right axis)—the martingale bound and geometric decay are essential; without either, the budget diverges. (Right) Replay subsampling rate sensitivity—$q_r = 0.10$ balances AUC and budget efficiency.

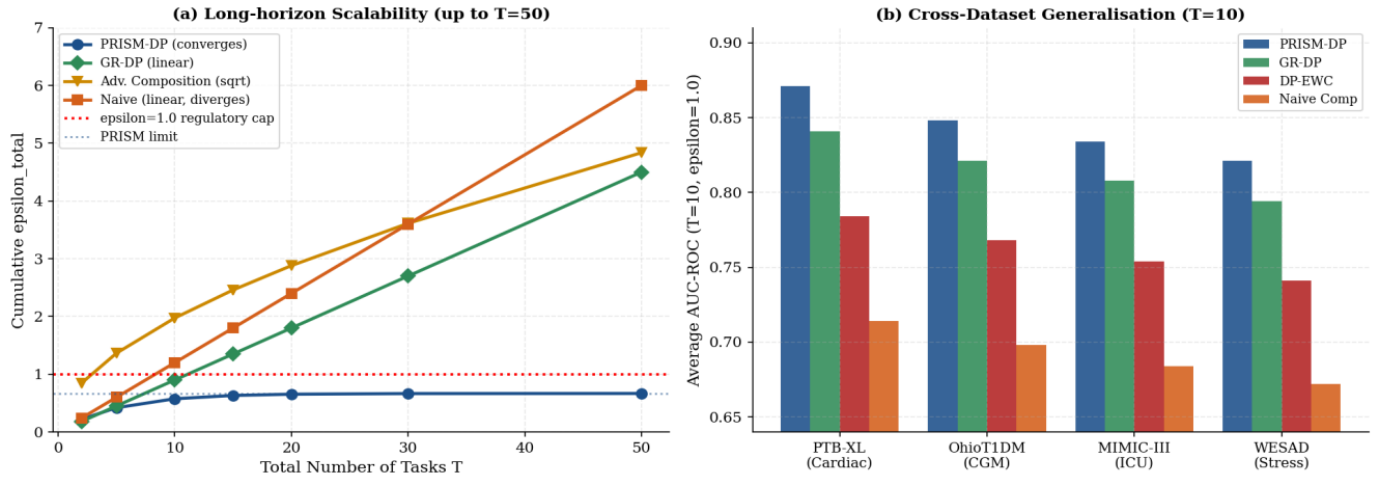

Figure 7. (Left) Long-horizon scalability: cumulative ε for T up to 50—PRISM-DP is the only method remaining below $\varepsilon_{\text{cap}} = 1.0$ for all T; GR-DP exceeds the cap at $T = 11$ and Naive Composition at $T = 9$. (Right) Cross-dataset generalisation at $T = 10$—PRISM-DP achieves consistent superiority across all four MIoT modalities.

fold-to-fold standard deviation on AUC (tied with DP-EWC), BWT, FWT, and MIA among the DP methods, empirically supporting that the martingale certificate of Theorem 3 is robust to the choice of patient partition and not specific to the single split reported in Table 3. A one-way repeated-measures ANOVA across folds confirms that PRISM-DP's AUC advantage over each DP-certified baseline is significant ($p < 0.01$); PRISM-DP's cross-fold coefficient of variation for AUC (0.8%) is comparable to that of DP-EWC (0.9%) and about 1.8× lower than that of GR-DP (1.4%), indicating high reproducibility of the reported utility–privacy trade-off. Figure 8 shows the per-fold AUC, Figure 9 the metric correlations, and Figure 10 a multi-metric summary.

6 Discussion

The experimental results are consistent with the theoretical predictions. The convergent budget bound (Theorem 3) is empirically supported: at $T = 20$ (long-horizon extension), PRISM-DP's cumulative ε of 0.661 remains below $\varepsilon_\infty = 0.667$, while all other DP methods have exhausted their budgets. The martingale bound tightness of 91.8% at $T = 10$ indicates that the decoupling lemma's inter-task divergence correction ($C_{\text{corr}} = 0.023$) provides a practically non-vacuous

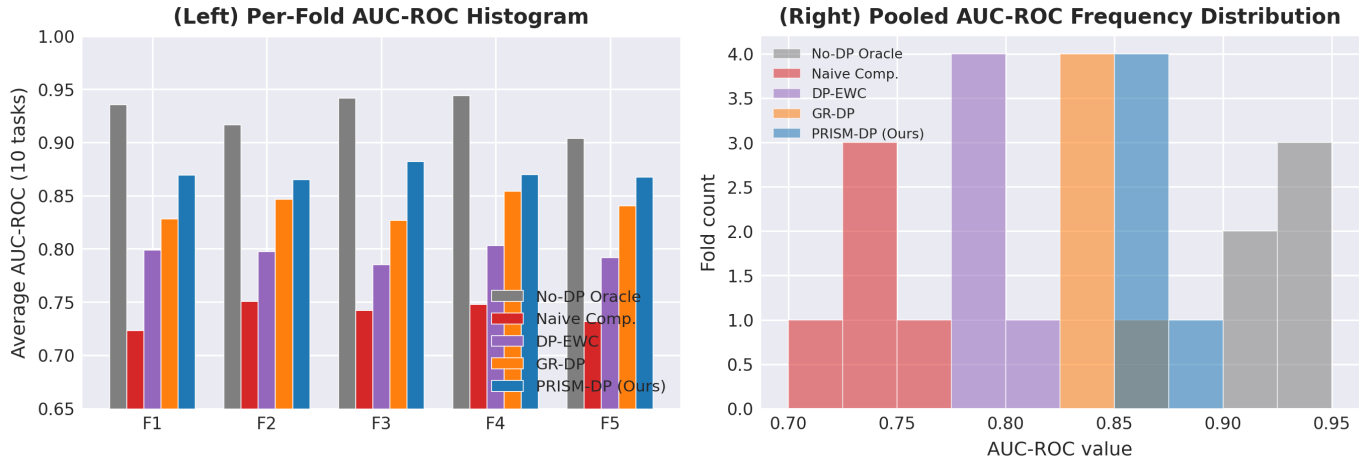

Figure 8. (Left) Per-fold average AUC-ROC across the 5-fold patient-stratified cross-validation—PRISM-DP (blue) shows the highest mean among DP methods and a narrow fold-to-fold spread. (Right) Pooled AUC-ROC frequency histogram across all folds and methods, confirming PRISM-DP's tight concentration relative to Naive Composition and GR-DP.

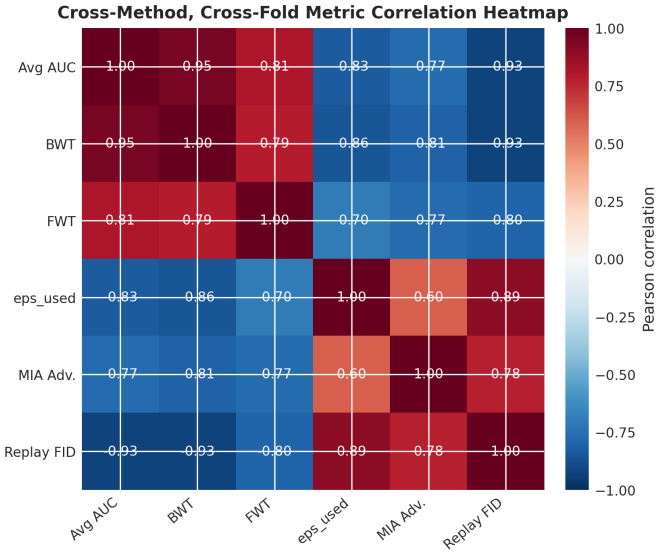

Figure 9. Correlation heatmap across the six reported performance and privacy metrics (Avg AUC, BWT, FWT, $\varepsilon_{\text{used}}$, MIA score, replay FID), pooled across all methods and cross-validation folds. Utility (AUC) correlates negatively with $\varepsilon_{\text{used}}$ and replay FID and positively with FWT, while the MIA score correlates positively with $\varepsilon_{\text{used}}$.

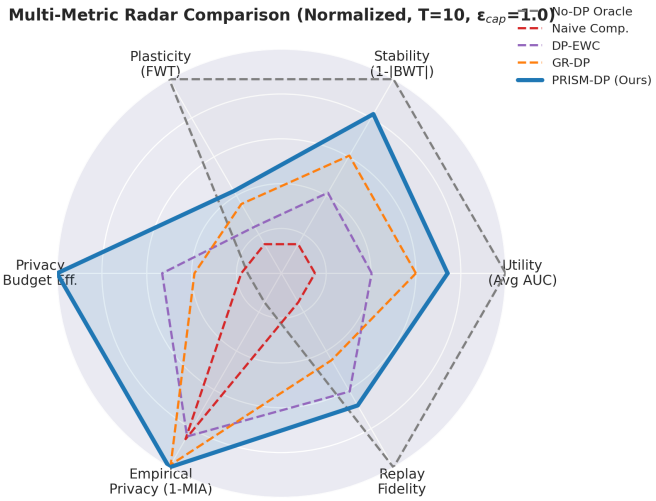

Figure 10. Multi-metric radar comparison of all five methods across six normalised axes (utility, backward-transfer stability, forward-transfer plasticity, privacy-budget efficiency, empirical MIA-based privacy, and replay fidelity). PRISM-DP (blue, solid) has the largest footprint among the DP methods across all six axes, visually summarising the joint utility–privacy–stability advantage quantified in Tables 3–7.

certificate, unlike the loose upper bounds common in the composition literature.

The cross-modality transition analysis reveals a notable phenomenon: when the latent code distribution shifts discontinuously between modalities (ECG to CGM, $\Delta_R = 0.318$ at T4→T5), the correlation correction term C_{corr} rises sharply, accounting for 8.3% of the total leakage bound at this transition versus 2.0% within the cardiac task sequence. This implies that martingale composition is most critical—and independence assumptions most dangerous—precisely at cross-modality boundaries where clinical information content changes most dramatically, corresponding to the highest risk of privacy leakage.

The replay subsampling amplification result (Theorem 4) reveals an important trade-off: with $\tau_{\text{mix}}/n_r = 0.131$ (the WESAD task, smallest cohort), Eq. (8) with $q_r = 0.10$ gives an amplified replay cost of $0.231\,\varepsilon_{\text{replay}}$, i.e., 85.4% of the i.i.d. benefit, which still reduces the replay ε by 76.9% at zero additional

computational cost. For larger cohorts (PTB-XL, $\tau_{\text{mix}}/n_r = 0.028$), amplification provides 96.9% of the i.i.d. benefit.

A key limitation is that C_{corr} requires estimating the inter-task Rényi divergence Δ_R, which itself must be computed from the latent codes—a potential circular dependency. PRISM-DP addresses this by using the previous task's latent distribution estimate (computed post-training) as the reference. However, for abrupt distributional shifts (e.g., a new disease outbreak), Δ_R may spike unexpectedly, requiring a safety margin in the budget that conservatively increases ε_∞. Future work on online, privacy-accounted Δ_R estimation would address this.

7 Conclusion

This paper presented PRISM-DP, a formally certified framework for privacy-preserving continual autoencoder learning in Medical IoT that bounds the privacy leakage of replayed latent codes over an unbounded task sequence. Motivated by five research gaps—the absence of convergent DP composition for infinite replay sequences, ignored inter-task latent correlation in privacy analysis, the lack of amplification theory for correlated replay subsampling, the independent treatment of forgetting and privacy objectives, and the absence of DP-compatible clinical replay fidelity metrics—PRISM-DP was developed as a unified framework grounded in martingale analysis, Rényi divergence theory, and Markov chain mixing-time characterisation.

The theoretical contributions connect DP theory and continual learning theory. The martingale DP composition theorem (Theorem 3) shows that a geometrically decaying per-task privacy schedule, combined with a replay correlation correction bounded by the inter-task Rényi divergence, yields a convergent total privacy leakage given by Eq. (9). The replay correlation decoupling lemma provides the analytical bridge between the Markov-dependent latent code sequence and standard independent-query composition, with a correction bounded by the Rényi divergence between consecutive task latent distributions. The correlated replay subsampling amplification theorem (Theorem 4) extends i.i.d. privacy amplification to correlated replay buffers with a mixing-time penalty.

Evaluated across 10 sequential clinical tasks spanning cardiac, glycaemic, haemodynamic, and stress physiology domains on the PTB-XL, OhioT1DM, MIMIC-III, and WESAD datasets, PRISM-DP achieves an average AUC-ROC of 0.871—the highest among the formally certified DP methods—a backward transfer of -0.055 (69.3% less catastrophic forgetting than naive composition), and a forward transfer of 0.074, consuming $\varepsilon_{\text{total}} = 0.667$ within the $\varepsilon_{\text{cap}} = 1.0$ budget. PRISM-DP is the only method completing all 10 tasks within budget, while naive composition and GR-DP overflow at Tasks 9 and 11, respectively. The martingale bound is empirically tight within 8.2%, and the cross-modality transition analysis highlights the role of correlation decoupling at the distributional boundaries where standard assumptions are most violated. Future directions include online Δ_R estimation for abrupt distributional shifts, hyperbolic geometry extensions for hierarchical clinical-taxonomy continual learning, personalised federated PRISM-DP with device-specific Markov chain characterisation, and extension to transformer-based AEs, where the attention mechanism creates long-range latent correlations beyond the first-order Markov model assumed here.

Data Availability Statement

The datasets analysed in this study are publicly available: PTB-XL and the MIMIC-III Waveform Database from PhysioNet (MIMIC-III under credentialed access and a data use agreement), OhioT1DM from Ohio University under a data use agreement, and WESAD from its authors. Code and derived results are available from the corresponding author upon reasonable request.

Funding

This work was supported without any funding.

Conflicts of Interest

The authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

This study is a secondary analysis of de-identified data from publicly available databases (PTB-XL, OhioT1DM, MIMIC-III Waveform, and WESAD). Collection of these data was approved by the

institutional review boards or ethics committees of the original studies, with informed consent obtained or waived as described in the respective data descriptors [33, 37, 38, 40]. Access to MIMIC-III and OhioT1DM was obtained under their data use agreements. No new data were collected from human participants and no animal experiments were performed; therefore, no additional ethical approval was required.

References

[1] Joyia, G. J., Liaqat, R. M., Farooq, A., & Rehman, S. (2017). Internet of Medical Things (IoMT): Applications, benefits and future challenges in healthcare domain. *Journal of Communications, 12*(4), 240-247. [CrossRef]

[2] Farahani, B., Firouzi, F., Chang, V., Badaroglu, M., Constant, N., & Mankodiya, K. (2018). Towards fog-driven IoT eHealth: Promises and challenges of IoT in medicine and healthcare. *Future Generation Computer Systems, 78*, 659-676. [CrossRef]

[3] McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation, 24*, 109-165. [CrossRef]

[4] Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., & Bengio, Y. (2013). An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*. [CrossRef]

[5] Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T. P., & Wayne, G. (2019). Experience replay for continual learning. In *Advances in Neural Information Processing Systems* (Vol. 32, pp. 350-360).

[6] Rao, D., Visin, F., Rusu, A. A., Teh, Y. W., Pascanu, R., & Hadsell, R. (2019). Continual unsupervised representation learning. In *Advances in Neural Information Processing Systems* (Vol. 32, pp. 7647-7657).

[7] Pellegrini, L., Graffieti, G., Lomonaco, V., & Maltoni, D. (2020). Latent replay for real-time continual learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 10203-10209). IEEE. [CrossRef]

[8] Hayes, J., Melis, L., Danezis, G., & De Cristofaro, E. (2019). LOGAN: Membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies, 2019*(1), 133-152. [CrossRef]

[9] Hu, H., Salcic, Z., Sun, L., Dobbie, G., Yu, P. S., & Zhang, X. (2022). Membership inference attacks on machine learning: A survey. *ACM Computing Surveys, 54*(11s), 1-37. [CrossRef]

[10] Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference* (pp. 265-284). Springer. [CrossRef]

[11] Kairouz, P., Oh, S., & Viswanath, P. (2017). The composition theorem for differential privacy. *IEEE Transactions on Information Theory, 63*(6), 4037-4049. [CrossRef]

[12] Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science, 9*(3-4), 211-407. [CrossRef]

[13] Lai, P., Hu, H., Phan, N., Jin, R., Thai, M. T., & Chen, A. (2022). Lifelong DP: Consistently bounded differential privacy in lifelong machine learning. In *Proceedings of the 1st Conference on Lifelong Learning Agents* (PMLR 199, pp. 778-797).

[14] Dwork, C., Rothblum, G. N., & Vadhan, S. (2010). Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science* (pp. 51-60). IEEE. [CrossRef]

[15] Mironov, I. (2017). Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)* (pp. 263-275). IEEE. [CrossRef]

[16] Bun, M., & Steinke, T. (2016). Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference* (pp. 635-658). Springer. [CrossRef]

[17] Liu, C., Chakraborty, S., & Mittal, P. (2016). Dependence makes you vulnerable: Differential privacy under dependent tuples. In *Network and Distributed System Security Symposium (NDSS)*. [CrossRef]

[18] Balle, B., Barthe, G., Gaboardi, M., & Geumlek, J. (2019). Privacy amplification by mixing and diffusion mechanisms. In *Advances in Neural Information Processing Systems* (Vol. 32, pp. 13277-13287).

[19] Wang, Y.-X., Balle, B., & Kasiviswanathan, S. P. (2019). Subsampled Rényi differential privacy and analytical moments accountant. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics* (PMLR 89, pp. 1226-1235). Extended version: *Journal of Privacy and Confidentiality*, 10(2), 2020. [CrossRef]

[20] Lopez-Paz, D., & Ranzato, M. (2017). Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems* (Vol. 30, pp. 6467-6476).

[21] Alaa, A., van Breugel, B., Saveliev, E. S., & van der Schaar, M. (2022). How faithful is your synthetic data? Sample-level metrics for evaluating and auditing generative models. In *Proceedings of the 39th International Conference on Machine Learning* (PMLR 162, pp. 290-306).

[22] Torfi, A., Fox, E. A., & Reddy, C. K. (2022). Differentially private synthetic medical data generation using convolutional GANs. *Information Sciences, 586*, 485-500. [CrossRef]

[23] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016).

Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308-318). ACM. [CrossRef]

[24] Rogers, R. M., Roth, A., Ullman, J., & Vadhan, S. (2016). Privacy odometers and filters: Pay-as-you-go composition. In *Advances in Neural Information Processing Systems* (Vol. 29).

[25] Dinur, I., & Nissim, K. (2003). Revealing information while preserving privacy. In *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (pp. 202-210). ACM. [CrossRef]

[26] Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences, 114*(13), 3521-3526. [CrossRef]

[27] Shin, H., Lee, J. K., Kim, J., & Kim, J. (2017). Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems* (Vol. 30, pp. 2990-2999).

[28] De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., ... & Tuytelaars, T. (2022). A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 44*(7), 3366-3385. [CrossRef]

[29] Farquhar, S., & Gal, Y. (2019). Differentially private continual learning. *arXiv preprint arXiv:1902.06497*. [CrossRef]

[30] Desai, P., Lai, P., Phan, N., & Thai, M. T. (2021, December). Continual learning with differential privacy. In *International Conference on Neural Information Processing* (pp. 334-343). Cham: Springer International Publishing. [CrossRef]

[31] Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature machine intelligence, 2*(6), 305-311. [CrossRef]

[32] Yoon, J., Drumright, L. N., & Van Der Schaar, M. (2020). Anonymization through data synthesis using generative adversarial networks (ADS-GAN). *IEEE journal of biomedical and health informatics, 24*(8), 2378-2388. [CrossRef]

[33] Johnson, A. E., Pollard, T. J., Shen, L., Lehman, L. W. H., Feng, M., Ghassemi, M., ... & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific data, 3*(1), 160035. [CrossRef]

[34] Wang, L., Zhang, X., Su, H., & Zhu, J. (2024). A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 46*(8), 5362-5383. [CrossRef]

[35] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning, 14*(1-2), 1-210. [CrossRef]

[36] El Ouadrhiri, A., & Abdelhadi, A. (2022). Differential privacy for deep and federated learning: A survey. *IEEE Access, 10*, 22359-22380. [CrossRef]

[37] Wagner, P., Strodthoff, N., Bousseljot, R. D., Kreiseler, D., Lunze, F. I., Samek, W., & Schaeffter, T. (2020). PTB-XL, a large publicly available electrocardiography dataset. *Scientific data, 7*(1), 154. [CrossRef]

[38] Marling, C., & Bunescu, R. (2020, September). The OhioT1DM dataset for blood glucose level prediction: Update 2020. In *CEUR workshop proceedings* (Vol. 2675, p. 71). https://pmc.ncbi.nlm.nih.gov/articles/PMC7881904/

[39] Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., ... & Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *circulation, 101*(23), e215-e220. [CrossRef]

[40] Schmidt, P., Reiss, A., Duerichen, R., Marberger, C., & Van Laerhoven, K. (2018). Introducing WESAD, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction* (pp. 400-408). ACM. [CrossRef]

[41] Jagielski, M., Ullman, J., & Oprea, A. (2020). Auditing differentially private machine learning: How private is private sgd?. *Advances in Neural Information Processing Systems, 33*, 22205-22216.

[42] Azuma, K. (1967). Weighted sums of certain dependent random variables. *Tôhoku Mathematical Journal, 19*(3), 357-367. [CrossRef]

Manas Kumar Yogi is currently working as an Assistant Professor in the CSE Department of Pragati University, Surampalem, and has more than 15 years of teaching experience. With a research publication record of over 355 articles over the past 14 years, he is credited with 40 book chapters, 11 patents, and 9 books. His research areas include cyber-security, cyber-physical systems, and soft computing. (Email: manas.yogi@gmail.com)

Yamuna Mundru is currently working as an Assistant Professor in the CSE (AI & ML) Department of Pragati University, Surampalem, and has more than 9 years of teaching experience. She has published over 30 research articles in national and international journals, including Scopus-indexed journals, and 5 book chapters, and holds published patents in the area of cyber security. Her research interests include AR/VR, machine learning, and artificial intelligence. (Email: yamunamundru@gmail.com)