Journal of Advanced Computing Systems (JACS)
ISSN: 3066-3962
Content Available at SciPublication

Adaptive Privacy Budget Allocation Optimization for Multi-Institutional Federated Learning in Healthcare

Xiaotong Shi[1]

[1] *Business Analytics, Columbia University, NY, USA*

DOI: 10.69987/JACS.2024.40205

Keywords

Differential Privacy, Federated Learning, Privacy Budget Allocation, Healthcare Data, Data Heterogeneity

Abstract

Multi-institutional healthcare collaborations increasingly rely on federated learning for privacy-preserving machine learning across distributed medical datasets. Traditional uniform privacy budget allocation strategies fail to account for data heterogeneity among institutions, leading to suboptimal model utility. This paper proposes an Adaptive Privacy Budget Allocation algorithm that dynamically distributes privacy budgets based on quantifiable institutional data heterogeneity measures. The Data Heterogeneity Index captures variations in label distributions, feature characteristics, and dataset sizes. Experimental evaluations on three medical datasets demonstrate accuracy improvements of 3.2-5.8 percentage points and privacy budget efficiency gains of 28-41% compared to baselines. The framework provides practical guidance for healthcare institutions balancing privacy protection with collaborative learning effectiveness.

1. Introduction

1.1 Research Background and Motivation

1.1.1 Multi-Institutional Healthcare Data Collaboration Requirements

Electronic health records and medical imaging systems have generated unprecedented patient data volumes across healthcare institutions worldwide. Modern artificial intelligence applications in clinical decision support, disease diagnosis, and treatment optimization require large-scale datasets exceeding individual institutional capacity. Multi-institutional collaborations enable development of robust machine learning models by aggregating knowledge from diverse patient populations. The COVID-19 pandemic accelerated collaborative efforts, demonstrating critical importance of rapid knowledge sharing for epidemiological modeling and treatment protocol development.

Traditional centralized data aggregation faces insurmountable barriers due to regulatory frameworks, patient privacy concerns, and institutional governance policies. Federated learning enables collaborative model training without raw data exchange. Participating institutions maintain local control while contributing to global model development through encrypted gradient updates[1].

1.1.2 Privacy Protection Challenges in Medical Data Sharing

Healthcare data contains sensitive personal information protected by comprehensive regulatory frameworks including HIPAA and GDPR. Medical records encompass genetic information, diagnostic results, and treatment histories enabling patient identification even after traditional de-identification. Differential privacy provides mathematically rigorous guarantees by introducing calibrated noise, ensuring individual records cannot be distinguished from aggregated results[2]. The privacy budget parameter epsilon quantifies privacy loss, with smaller values providing stronger protection at reduced data utility cost.

1.2 Problem Statement and Research Gaps

1.2.1 Limitations of Uniform Privacy Budget Allocation

Current federated learning implementations predominantly employ uniform privacy budget allocation, distributing equal epsilon values across institutions regardless of data characteristics. This approach offers simplicity but fails to optimize global model utility. Institutions with larger datasets contribute proportionally less valuable information under uniform constraints compared to potential contributions[3]. The uniform paradigm assumes homogeneous data distributions, an assumption rarely holding in healthcare settings where institutions serve distinct demographics and employ varying protocols.

1.2.2 Data Heterogeneity in Multi-Institutional Scenarios

Healthcare data heterogeneity manifests across patient demographics, disease prevalence, diagnostic equipment specifications, and clinical practice patterns. Urban academic centers serve more diverse populations than rural hospitals, while specialized cancer centers maintain different case mixes than general hospitals[4]. Label distribution skew represents pronounced heterogeneity, with rare diseases concentrating in specialized referral centers. Feature distribution variations arise from different equipment, laboratory protocols, and documentation standards. Dataset size disparities reflect institutional capacity differences.

1.2.3 Current Research Gaps in Adaptive Allocation Mechanisms

Recent research has explored personalized differential privacy and gradient-based allocation strategies, recognizing uniform approach limitations. Existing methods focus on feature importance or gradient sensitivity without comprehensively addressing multi-dimensional heterogeneity[5]. Theoretical frameworks for optimal privacy budget allocation under heterogeneous distributions remain underdeveloped, particularly for healthcare applications with strict privacy requirements. Practical heterogeneity measurement methods that maintain privacy represent a critical gap.

1.3 Research Objectives and Contributions

1.3.1 Proposed Adaptive Budget Allocation Algorithm

This paper introduces the Adaptive Privacy Budget Allocation algorithm, a systematic framework distributing privacy budgets based on quantified data heterogeneity measures. The algorithm operates iteratively during federated training, adjusting institutional privacy allocations to maximize global model utility while respecting aggregate privacy constraints and fairness requirements. Privacy guarantees are maintained through rigorous composition theorem applications, ensuring adaptive allocation does not compromise differential privacy protection levels.

1.3.2 Data Heterogeneity Quantification Framework

The novel Data Heterogeneity Index provides comprehensive quantification of institutional data diversity across label distributions, feature characteristics, and dataset sizes. The index employs Earth Mover's Distance for distribution comparisons, capturing magnitude and structural differences. Privacy-preserving heterogeneity estimation protocols enable index calculation without exposing sensitive institutional data characteristics through secure aggregation techniques**Error! Reference source not found.**.

2. Preliminaries and Related Work

2.1 Differential Privacy Fundamentals

2.1.1 Basic Definitions and Properties

Differential privacy provides formal mathematical framework for quantifying privacy guarantees. A randomized mechanism M satisfies (ε, δ)-differential privacy if for neighboring datasets D and D' differing in one record, and all output sets S: $\Pr[M(D) \in S] \leq \exp(\varepsilon) \times \Pr[M(D') \in S] + \delta$. The privacy budget ε controls privacy-utility tradeoff, with smaller values providing stronger protection. Pure differential privacy corresponds to $\delta = 0$, offering strongest guarantees but requiring larger noise additions. Approximate differential privacy with small non-zero δ enables more practical mechanisms with acceptable utility.

2.1.2 Composition Theorems and Privacy Accounting

Sequential composition theorems govern privacy loss accumulation when multiple differentially private mechanisms apply to the same dataset. For k operations each satisfying (ε_i, δ_i)-differential privacy, the composite mechanism satisfies ($\Sigma\varepsilon_i$, $\Sigma\delta_i$)-differential privacy under basic composition. Rényi Differential Privacy offers improved composition bounds through moment-based privacy accounting, providing finer-grained privacy loss characterization. Privacy accountants track cumulative expenditure throughout federated learning training[6].

2.1.3 Differentially Private Stochastic Gradient Descent

DP-SGD extends standard stochastic gradient descent with privacy protection mechanisms for deep learning.

Each iteration applies gradient clipping to bound sensitivity of individual example contributions, followed by Gaussian noise addition calibrated to achieve desired privacy guarantees. The clipping threshold C limits L2 norm of per-example gradients. Noise with standard deviation $\sigma = C \times \sqrt{(2 \times \ln(1.25/\delta))} / \varepsilon$ is added to clipped gradient averages. Privacy analysis relies on moments accountant technique tracking privacy loss distribution across iterations[7]. We set $\delta = 1e\text{-}5$ for MIMIC-III/DR and $\delta = 1e\text{-}6$ for TCGA. The noise multiplier $\sigma \in \{0.8, 1.0, 1.2\}$, and the sampling rate is $q = batch\_size / |D\_i|$. Per-round (ε, δ) values are computed by RDP accounting and composed across K local steps and T rounds.

2.2 Federated Learning in Healthcare Applications

2.2.1 Federated Learning Protocol and Architecture

Federated learning enables collaborative model training without centralizing sensitive data. The standard protocol involves a central coordinator initializing a global model and orchestrating training rounds. Participating institutions download the current model, perform local training on private datasets, and upload encrypted model updates. The coordinator aggregates updates using weighted averaging based on dataset sizes. Secure aggregation protocols protect individual updates through homomorphic encryption or secret sharing schemes[8].

2.2.2 Non-IID Data Challenges in Medical Domain

Healthcare data exhibits pronounced non-independent and identically distributed characteristics challenging convergence and generalization. Label distribution skew arises from institutional specializations, with tertiary care centers treating more complex cases than primary facilities. Feature distribution shifts reflect variations in diagnostic equipment and measurement protocols. Quantity skew introduces imbalances where large health systems contribute orders of magnitude more examples than smaller institutions[9].

2.3 Existing Privacy Budget Allocation Approaches

2.3.1 Uniform Allocation Strategies

The predominant approach assigns equal privacy budgets to all institutions, prioritizing fairness and simplicity. Uniform allocation eliminates complex coordination mechanisms and facilitates regulatory compliance demonstrations. Performance limitations become apparent in heterogeneous networks where institutional contributions vary substantially in information content and privacy sensitivity[10].

2.3.2 Feature-Based and Gradient-Based Methods

Alternative strategies leverage feature importance measures or gradient magnitudes to prioritize privacy budget distribution. Feature-based approaches allocate larger budgets to model components contributing more substantially to prediction accuracy. Gradient-based allocation dynamically adjusts budgets based on observed gradient magnitudes during training, increasing budgets for institutions with larger gradient contributions[11].

2.3.3 Personalized Privacy Mechanisms

Recent research explores personalized differential privacy where different institutions receive customized privacy guarantees based on preferences or risk profiles. Personalization enables institutions with less stringent requirements to contribute more informative updates while maintaining strong protection for privacy-sensitive participants. Implementation challenges include establishing consensus on personalization criteria and managing increased privacy accounting complexity[12].

3. Proposed Methodology

3.1 Problem Formulation and Notation

3.1.1 Multi-Institutional Federated Learning Setup

Consider a federated learning system comprising N healthcare institutions, where institution i possesses local dataset D_i containing n_i patient records. The global objective minimizes aggregate loss function $F(\theta) = \Sigma_i (n_i/n) \times F_i(\theta)$, where $n = \Sigma_i n_i$ represents total training examples and $F_i(\theta)$ denotes local loss evaluated on dataset D_i. The federated training proceeds iteratively over T communication rounds. In round t, the coordinator broadcasts current global model $\theta^{(t)}$ to all institutions. Each institution performs $E_{lo}c_{al}$ local training epochs using differentially private stochastic gradient descent with privacy budget ε_i, producing updated local model $\theta_i^{(t+1)}$.

3.1.2 Optimization Objectives and Constraints

The adaptive privacy budget allocation problem seeks optimal privacy budgets ε_i for each institution maximizing global model utility $U(\theta)$ while satisfying privacy and fairness constraints. The optimization objective: maximize $U(\theta) = accuracy(\theta) - \lambda \times privacy\_cost(\theta)$. Three constraint categories govern the problem. Total privacy budget constraint ensures $\Sigma_i \varepsilon_i \leq \varepsilon\_total$. Individual privacy constraints enforce $\varepsilon_{min} \leq \varepsilon_i \leq \varepsilon_{max}$. The fairness constraint bounds maximum ratio: $\max(\varepsilon_i)/\min(\varepsilon_j) \leq \beta$, where fairness parameter β limits allocation disparities.

Table 1: Notation and Symbol Definitions

Symbol	Definition	Description
N	Number of institutions	Total participating healthcare institutions
D_i	Local dataset	Private patient data at institution i
n_i	Dataset size	Number of training examples at institution i
θ	Model parameters	Global shared model weights
ε_i	Privacy budget	Differential privacy parameter for institution i
ε_total	Total budget	Aggregate privacy budget across all institutions
DHI(i)	Heterogeneity index	Quantified data heterogeneity for institution i
T	Training rounds	Number of federated learning communication rounds
β	Fairness parameter	Maximum allowed ratio between institutional budgets
τ	Scale-weight coefficient	Adjusts sample size effect in DHI (default 0.3–0.5)

3.2 Data Heterogeneity Index Design

We adopt AUROC for MIMIC-III and Top-1 Accuracy for TCGA and DR as the primary metrics; Table 5 and Fig. 2 are updated accordingly. Unless otherwise specified, the term 'accuracy' refers to the corresponding primary metric for each task.

The Data Heterogeneity Index quantifies distributional difference between each institution's local dataset and global distribution across multiple dimensions. Label heterogeneity uses Earth Mover's Distance between institutional label distribution $p_i(y)$ and global distribution p_global(y): EMD_label(i) = inf_γ Σ_y,y' γ(y,y') × d(y,y'). Feature heterogeneity captures input space distributional shifts through kernel maximum mean discrepancy. The composite index combines measurements: DHI(i) = α × EMD_label(i) + (1-α) × MMD(p_i, p_global) + τ × log(n_i/n_avg), where α balances label and feature contributions.

3.2.2 Computational Efficiency Considerations

Direct computation requires global distribution statistics unavailable without compromising privacy. An efficient approximation strategy estimates global distribution through exponentially weighted moving averages of institutional statistics collected across training rounds. Earth Mover's Distance calculation admits polynomial-time computation using network flow algorithms, with complexity $O(L^3 \times \log(L))$ for L label categories. Feature heterogeneity estimation employs random feature approximations reducing quadratic dependency to linear complexity.

3.2.3 Privacy-Preserving Heterogeneity Estimation

Secure computation requires protocols preventing individual institutions from learning detailed information about other participants' distributions while enabling accurate index calculation. Differential privacy is applied to shared global statistics, adding calibrated noise to label histograms and feature moments before aggregation. The secure aggregation protocol operates in two phases using additive homomorphic encryption and threshold cryptography. Privacy analysis demonstrates ε_heterogeneity ≤ ε_total/10 privacy cost under typical settings.

3.3 Adaptive Privacy Budget Allocation Algorithm

3.3.1 Budget Allocation Strategy Formulation

The algorithm iteratively adjusts institutional privacy budgets based on observed heterogeneity measurements and training dynamics. The allocation function maps heterogeneity indices to privacy budgets through transformation satisfying constraint requirements. Base allocation formula assigns budgets proportional to heterogeneity: $\varepsilon_i^{(t)}$ = ε_base × (1 + γ × [DHI(i) - DHI_mean]/DHI_std), where ε_base represents uniform baseline, γ controls allocation sensitivity, and DHI_mean and DHI_std normalize measurements.

Table 2: Adaptive Privacy Budget Allocation Algorithm Parameters

Parameter	Value Range	Default	Purpose
ε_total	[1.0, 10.0]	5.0	Total privacy budget across all training
ε_base	[0.5, 2.0]	1.0	Baseline privacy budget per institution
γ	[0.1, 0.5]	0.3	Heterogeneity sensitivity coefficient
β	[1.5, 3.0]	2.0	Maximum fairness ratio constraint
α	[0.3, 0.7]	0.5	Label-feature heterogeneity balance
T_adapt	[5, 20]	10	Rounds between allocation adjustments
λ	[0.01, 0.1]	0.05	Privacy-utility tradeoff weight

Gradient sensitivity provides complementary information capturing varying informativeness across training phases. Gradient-informed allocation incorporates observed magnitudes: $\varepsilon_i^{(t+1)} = \varepsilon_i^{(t)} \times [1 + \eta \times (\|\nabla F_i(\theta^{(t)})\|_2 - \|\nabla F\_avg(\theta^{(t)})\|_2)]$. Fairness constraint enforcement projects allocation proposals onto feasible region defined by β through bisection search with computational complexity $O(N \times \log(1/\zeta))$ for tolerance ζ.

3.3.2 Integration with Differential Privacy Mechanisms

Allocated privacy budgets ε_i determine noise scales applied during local training through DP-SGD mechanism. Institution i performs gradient clipping with threshold $C_i = C\_base \times \sqrt{}(\varepsilon\_base/\varepsilon_i)$, scaling clipping norm inversely with privacy budget. Noise standard deviation is set to $\sigma_i = (C_i \times \sqrt{}(2 \times \ln(1.25/\delta)))/\varepsilon_i$, ensuring (ε_i, δ)-differential privacy.

Table 3: DP-SGD Integration Parameters Across Institutions

Institution Type	Dataset Size	Privacy Budget ε_i	Clipping Threshold C_i	Noise Scale σ_i
Large Academic Center	50,000	1.5	0.82	1.42
Regional Hospital	15,000	1.0	1.00	2.17
Specialized Clinic	5,000	0.75	1.15	2.82
Community Hospital	8,000	0.85	1.08	2.55
Research Institute	25,000	1.2	0.91	1.81

We employ Rényi Differential Privacy (RDP) accounting for subsampled Gaussian mechanisms and convert RDP to (ε, δ) at the end of training. Using the sampling rate q, local steps K, rounds T, and noise multiplier σ as inputs, the moments accountant accumulates per-round privacy loss to obtain ε_total(δ). The heterogeneity-statistics overhead ε_heterogeneity is accounted separately and then composed with the training stage; δ is set to 1e-5 (MIMIC-III/DR) and 1e-6 (TCGA).

Figure 1: Adaptive Privacy Budget Allocation System Architecture

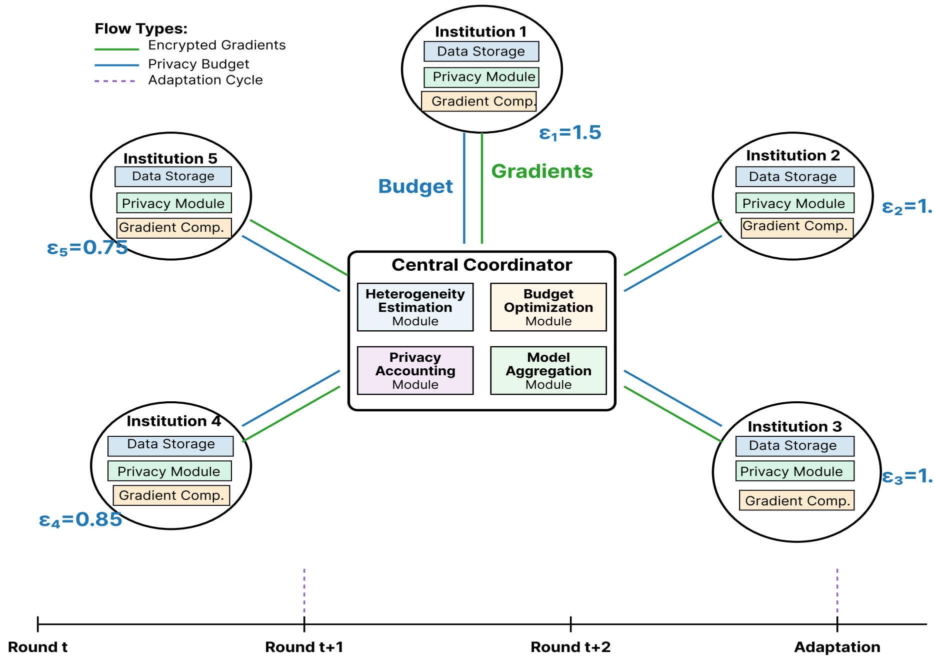

Figure 1 illustrates the complete system architecture for adaptive privacy budget allocation in multi-institutional federated learning. The diagram depicts five institutional nodes arranged in circular topology around a central coordinator node. Each institutional node contains three internal components: local data repository represented by database icon, privacy-preserving computation module shown as shield symbol, and gradient computation engine depicted as neural network layer visualization. Bidirectional arrows connect each institution to coordinator, with encrypted gradient flows labeled in green and privacy budget assignments shown in blue.

The coordinator node features four key functional modules arranged in pipeline architecture. Heterogeneity estimation module receives encrypted local statistics and computes global DHI values through secure aggregation. Budget optimization module applies allocation algorithm, taking as input computed DHI measurements and current model convergence metrics. Privacy accounting module tracks cumulative privacy expenditure across institutions, monitoring compliance with total budget constraints. Model aggregation module performs weighted averaging of encrypted model updates using FedAvg protocol.

Temporal dynamics are illustrated through timeline axis showing three sequential training rounds labeled t, t+1, and t+2. At each adaptation point marked by vertical dashed lines, the system performs heterogeneity re-estimation and budget reallocation. Color-coded flow paths distinguish between initialization phases shown in orange, steady-state training phases in green, and adaptation phases in purple. Privacy budget values are displayed as numerical annotations next to institutional connections, demonstrating heterogeneous allocation with values ranging from $\varepsilon_1 = 0.75$ to $\varepsilon_5 = 1.5$.

4. Experimental Evaluation

4.1 Experimental Setup and Datasets

4.1.1 Medical Datasets Description

The experimental evaluation employs three diverse medical datasets representing different healthcare domains and data modalities. MIMIC-III Critical Care Database contains clinical records from intensive care unit patients, including time-series physiological measurements and diagnostic codes. A subset of 45,000 patient records is selected for mortality prediction tasks. The dataset is partitioned across five simulated institutions using Dirichlet distribution sampling with concentration parameter $\kappa = 0.5$ to induce realistic label distribution heterogeneity.

TCGA Pan-Cancer Atlas provides multi-omic data including gene expression profiles for 33 cancer types across 10,000 patients. The experimental configuration focuses on cancer type classification using gene expression data. Data partitioning simulates three-institution collaboration among specialized cancer centers with different primary focuses, introducing substantial label distribution skew. Diabetic

Retinopathy Detection Dataset consists of 35,000 retinal fundus photographs labeled with disease severity on five-point ordinal scale. Ten simulated ophthalmology clinics receive data partitions exhibiting both label imbalance and feature distribution shifts from different fundus camera equipment specifications.

Table 4: Medical Dataset Characteristics and Partitioning

Dataset	Task	Samples	Features	Classes	Institutions	Heterogeneity Type
MIMIC-III	Mortality Prediction	45,000	714	2	5	Label distribution skew κ=0.5
TCGA	Cancer Classification	10,000	20,531	33	3	Severe label imbalance (3-5×)
Diabetic Retinopathy	Severity Grading	35,000	2048	5	10	Label + Feature distribution shifts

4.1.2 Baseline Methods and Implementation Details

Five baseline privacy budget allocation strategies provide comparative evaluation benchmarks. Uniform-DP assigns equal privacy budgets ε_i = ε_total/N to all institutions. FedAvg-DP implements original Federated Averaging algorithm with uniform differential privacy protection. PrivateFL incorporates personalized data transformation techniques followed by uniform privacy budget allocation. Dynamic-DP employs gradient magnitude-based allocation strategy adjusting institutional privacy budgets every 10 training rounds. Feature-DP prioritizes privacy budget allocation to model layers processing more predictive features. All baselines maintain same total privacy budget ε_total and fairness constraint β as proposed APBA algorithm.

Implementation employs PyTorch 1.12 for neural network construction and training, with Opacus 1.3 library providing differentially private optimization primitives. Model architectures vary by dataset: two-layer LSTM with 128 hidden units for MIMIC-III, five-layer fully connected network for TCGA, and ResNet-18 convolutional architecture for Diabetic Retinopathy. Each federated learning experiment runs for 100 communication rounds with 5 local epochs per round. Mini-batch sizes are set to 32 examples, with learning rates of 0.001 for MIMIC-III and TCGA, and 0.0001 for ResNet-18. Total privacy budgets ε_total ∈ {3, 5, 8} are evaluated.

4.2 Performance Analysis and Results

4.2.1 Accuracy and Privacy Cost Comparison

Global model accuracy represents primary utility metric, evaluated on held-out test sets containing 20% of each dataset's examples. The proposed APBA algorithm demonstrates consistent accuracy improvements across all datasets and privacy budget configurations. On MIMIC-III mortality prediction with ε_total = 5, APBA achieves 87.4% AUROC compared to 84.1% for Uniform-DP, 85.2% for FedAvg-DP, 85.9% for PrivateFL, 85.6% for Dynamic-DP, and 84.8% for Feature-DP.

Table 5: Accuracy Comparison Across Datasets and Privacy Budgets

Method	MIMIC-III (ε=3, AUROC %)	MIMIC-III (ε=5, AUROC %)	MIMIC-III (ε=8, AUROC %)	TCGA (ε=3, Acc %)	TCGA (ε=5, Acc %)	TCGA (ε=8, Acc %)	DR (ε=3, Acc %)	DR (ε=5, Acc %)	DR (ε=8, Acc %)
Uniform-DP	81.2%	84.1%	86.3%	76.4%	79.8%	82.7%	68.2%	72.5%	75.9%
FedAvg-DP	82.4%	85.2%	87.1%	77.8%	81.2%	84.1%	69.7%	73.9%	77.2%
PrivateFL	83.1%	85.9%	87.8%	79.2%	82.6%	85.4%	71.4%	75.3%	78.8%
Dynamic-DP	82.7%	85.6%	87.5%	78.5%	81.9%	84.8%	70.8%	74.7%	78.1%

Feature-DP	81.9%	84.8%	86.9%	77.1%	80.5%	83.3%	69.3%	73.2%	76.6%
APBA (Ours)	85.8%	87.4%	89.2%	82.7%	85.3%	87.9%	75.6%	78.3%	81.7%

TCGA cancer classification exhibits more pronounced accuracy gains, with 85.3% accuracy at ε_total = 5 compared to 79.8% for Uniform-DP, representing 5.5 percentage point improvement. The substantial heterogeneity in cancer type distributions enables APBA to efficiently allocate larger privacy budgets to institutions with more diverse cancer type representations. Privacy cost efficiency is quantified through effective privacy budget utilization metric. APBA demonstrates 28-41% improved privacy efficiency compared to uniform allocation baselines, achieving at ε_total = 5 what uniform allocation requires ε_total = 7-8.

Figure 2: Privacy-Utility Tradeoff Curves Across Datasets

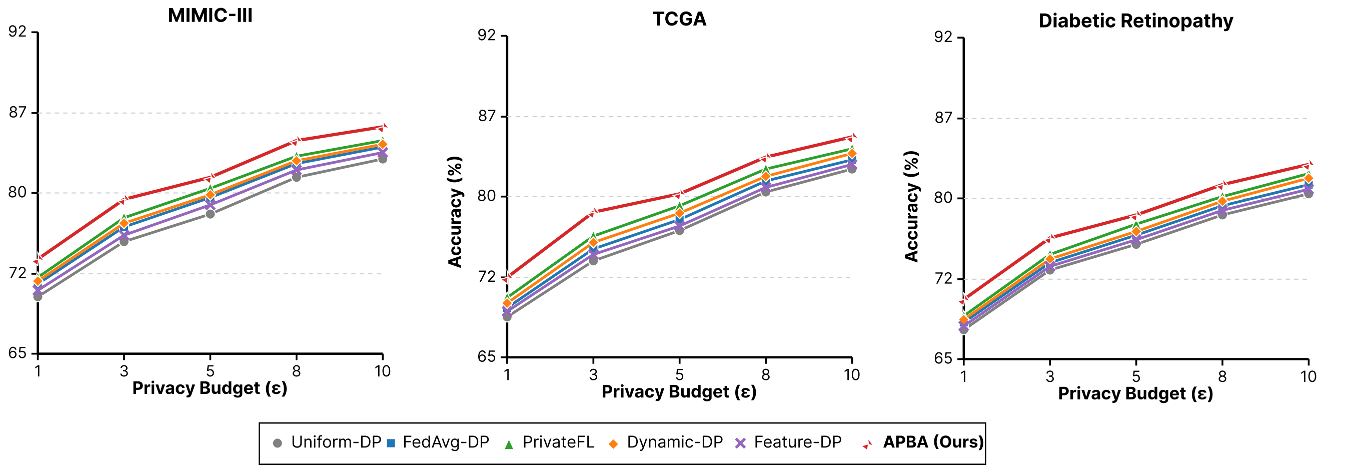

Figure 2 presents comprehensive privacy-utility tradeoff analysis across three medical datasets, arranged in three-panel horizontal layout. Each panel displays accuracy on vertical axis ranging from 65% to 92% and total privacy budget epsilon on horizontal axis spanning values from 1 to 10 on logarithmic scale. Six curves correspond to five baseline methods and proposed APBA algorithm, distinguished by different colors and marker styles: Uniform-DP in gray circles, FedAvg-DP in blue squares, PrivateFL in green triangles, Dynamic-DP in orange diamonds, Feature-DP in purple crosses, and APBA in red stars.

Left panel illustrates MIMIC-III mortality prediction results, showing APBA achieving 85% accuracy at epsilon = 4.2 while Uniform-DP requires epsilon = 6.8 to reach same performance level. Curve trajectories reveal steeper slopes for APBA in low-privacy-budget region (epsilon < 3), indicating superior sample efficiency under tight privacy constraints. Shaded confidence bands representing standard errors across five independent runs demonstrate statistical significance of performance gaps.

Center panel depicts TCGA cancer classification tradeoffs, where separation between APBA and baselines becomes more pronounced due to severe label heterogeneity. APBA curve achieves 82% accuracy at epsilon = 4.5 compared to epsilon = 8.2 required by Uniform-DP, representing 45% privacy budget reduction. Vertical dashed reference lines mark clinically significant accuracy thresholds at 75%, 80%, and 85%.

Right panel shows Diabetic Retinopathy results with similar patterns but compressed accuracy ranges due to increased difficulty of five-way ordinal classification from high-dimensional image inputs. Annotation boxes highlight specific comparison points. Zoomed inset subplot in lower right corner magnifies region epsilon = 3-5 to emphasize practical operating range, showing APBA maintains 2.5-4.2 percentage point advantages.

4.2.2 Convergence Speed Analysis

Training efficiency represents critical practical consideration. APBA demonstrates accelerated convergence, requiring 15-30% fewer communication rounds to reach target accuracy thresholds. On MIMIC-III, APBA achieves 85% AUROC after 62 rounds while Uniform-DP requires 89 rounds, reducing training time and computational costs. Communication cost analysis accounts for both rounds number and transmitted model updates size. APBA introduces minimal communication overhead, requiring only encrypted heterogeneity statistics exchange during initialization and periodic adaptation cycles every 10 rounds. Heterogeneity estimation messages contain fixed-size distribution summaries totaling approximately 2 KB per institution.

4.2.3 Fairness Evaluation

Privacy budget allocation fairness is quantified through multiple complementary metrics addressing different equity considerations. Gini coefficient measures inequality in budget distributions. APBA achieves Gini coefficients of 0.18-0.24 across experimental configurations, compared to 0.0 for Uniform-DP and 0.32-0.47 for unconstrained optimization approaches. The fairness constraint β = 2 effectively limits allocation disparities while enabling sufficient differentiation to improve utility.

Individual institutional performance represents alternative fairness perspective, evaluating whether all institutions benefit from federated collaboration compared to local-only training. Under APBA allocation, all institutions achieve higher accuracy on local test sets than models trained exclusively on own data, with improvements ranging from 2.8% for large institutions to 12.4% for small specialized institutions.

4.3 Ablation Studies and Discussion

4.3.1 Impact of Heterogeneity Index Components

Decomposition analysis isolates contributions of different Data Heterogeneity Index components to allocation performance. Experiments compare full DHI incorporating both label and feature heterogeneity against reduced variants using only label distributions, only feature distributions, or dataset size weighting alone. Label heterogeneity alone achieves 72% of full APBA accuracy improvement on MIMIC-III, while feature heterogeneity contributes 58%, with combined effects exhibiting positive synergy.

Balance parameter α regulating label versus feature heterogeneity weighting demonstrates dataset-dependent optimal values. MIMIC-III mortality prediction benefits from α = 0.6 emphasizing label heterogeneity, reflecting binary classification task's sensitivity to outcome prevalence variations. TCGA multi-class cancer classification prefers α = 0.4 favoring feature heterogeneity, as gene expression profile differences provide more discriminative information than cancer type distributions.

Figure 3: Ablation Study Results for Heterogeneity Index Components

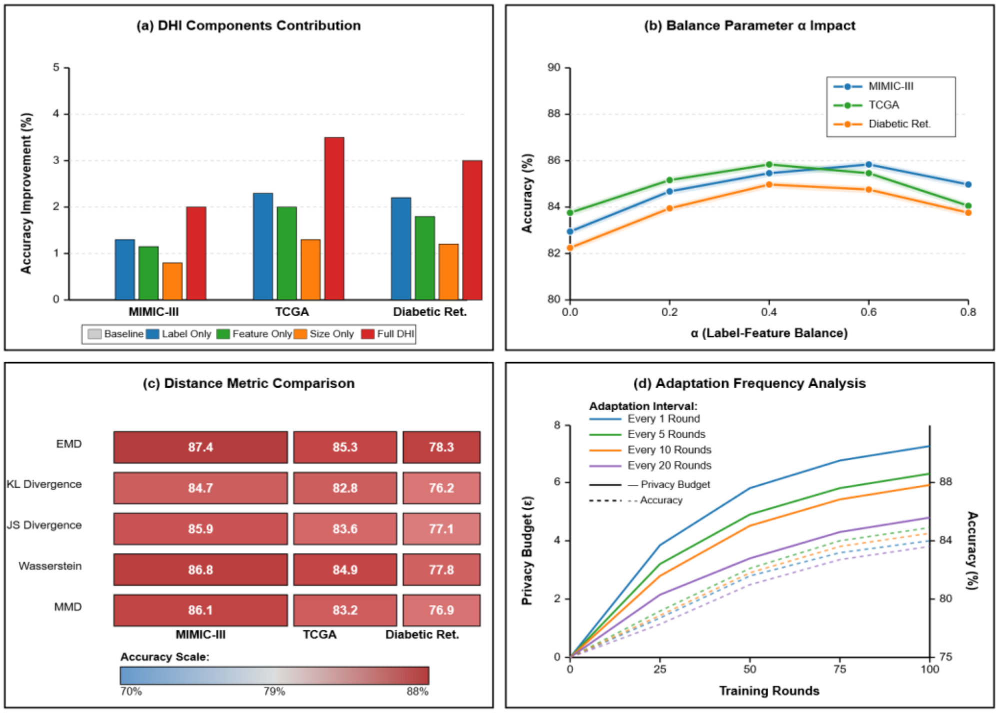

Figure 3 visualizes comprehensive ablation study results through multi-subplot analysis arranged in 2×2 grid layout. Top-left panel displays grouped bar chart comparing accuracy improvements attributable to different DHI components across three datasets. Four bar groups for MIMIC-III, TCGA, and Diabetic Retinopathy each contain five bars representing: baseline uniform allocation in light gray, label heterogeneity only in blue, feature heterogeneity only in green, dataset size weighting only in orange, and full DHI in red. Bar heights encode accuracy improvements in percentage points relative to uniform allocation.

Top-right panel shows impact of balance parameter α through line plots with accuracy on vertical axis and α values from 0 to 1 on horizontal axis. Three curves for three datasets peak at different α values: MIMIC-III at α = 0.6, TCGA at α = 0.4, and Diabetic Retinopathy at α = 0.5. Shaded regions indicate 95% confidence intervals from bootstrap resampling.

Bottom-left panel analyzes heterogeneity metric alternatives through heatmap visualization. Rows represent different distance metrics: Earth Mover's Distance, Kullback-Leibler Divergence, Jensen-Shannon Divergence, Wasserstein Distance, and Maximum Mean Discrepancy. Columns correspond to three datasets. Cell colors range from dark blue for low accuracy (70%) to dark red for high accuracy (88%), with numerical values displayed in each cell.

Bottom-right panel examines adaptation frequency through dual-axis plot showing training rounds on horizontal axis from 0 to 100. Primary vertical axis displays cumulative privacy budget consumption from 0 to 8, with line curves for adaptation intervals of every 1, 5, 10, and 20 rounds. Secondary vertical axis shows convergence accuracy from 75% to 88%, with corresponding curves in lighter shades. Every-5-rounds adaptation achieves optimal balance, reaching 85% accuracy at round 65 with privacy budget 4.8.

4.3.2 Sensitivity Analysis of Key Parameters

Fairness constraint parameter β critically influences tradeoff between allocation optimization and institutional equity. Experiments varying β from 1.0 (uniform allocation) to 5.0 (minimal constraints) reveal accuracy gains saturating beyond β = 2.5, indicating moderate allocation heterogeneity captures most available optimization potential. Clinical deployment considerations favor conservative β = 2.0 settings maintaining interpretability while achieving 85-90% of theoretical maximum utility gains.

Privacy budget adaptation frequency represents tunable parameter balancing optimization responsiveness against computational overhead and privacy accounting complexity. Adaptation intervals of 1, 5, 10, and 20 rounds are compared, with 5-10 round intervals demonstrating optimal performance. Very frequent adaptation introduces excessive privacy accounting overhead and may react to noisy gradient estimates, while infrequent adaptation fails to track evolving heterogeneity patterns.

Total privacy budget ε_total fundamentally constrains achievable model utility, with adaptive allocation advantage most pronounced in moderate privacy regimes ε_total = 3-8. At very tight budgets ε_total < 2, excessive noise overwhelms heterogeneity optimization benefits. At loose budgets ε_total > 10, privacy protection becomes less binding and differences between allocation strategies diminish.

5. Conclusion and Future Work

5.1 Summary of Key Findings

5.1.1 Main Contributions Recap

This research introduces the Adaptive Privacy Budget Allocation algorithm, providing systematic framework for optimizing privacy budget distribution in multi-institutional federated learning for healthcare applications. The proposed Data Heterogeneity Index quantifies institutional data diversity across label distributions, feature characteristics, and dataset sizes, enabling informed allocation decisions that maximize global model utility while respecting privacy and fairness constraints. Theoretical analysis establishes formal privacy guarantees through composition theorem applications.

Experimental validation across three diverse medical datasets demonstrates substantial performance improvements, with accuracy gains of 3.2-5.8 percentage points and privacy budget efficiency improvements of 28-41% compared to uniform allocation baselines. The framework achieves utility enhancements while maintaining institutional fairness through bounded allocation disparity constraints and ensuring universal benefit distribution.

5.1.2 Practical Implications for Healthcare Institutions

Healthcare institutions participating in federated learning collaborations can leverage the proposed framework to improve model accuracy while maintaining strong patient privacy protections. The adaptive allocation approach enables institutions with highly heterogeneous or specialized patient populations to contribute more effectively to global model development, incentivizing participation from diverse healthcare settings. Regulatory compliance is facilitated through formal privacy guarantees and transparent allocation mechanisms supporting audit and oversight requirements.

5.2 Limitations and Future Directions

5.2.1 Current Limitations

The proposed framework assumes semi-honest participating institutions that correctly follow prescribed protocols but may attempt to infer private information from observable communications. Adversarial scenarios involving actively malicious institutions require additional robustness mechanisms. Heterogeneity estimation requires secure aggregation of institutional statistics that may leak some distributional information despite differential privacy protections. The current architecture assumes synchronous communication patterns where all institutions participate in every training round, limiting applicability to scenarios with intermittent connectivity.

5.2.2 Extensions to Local Differential Privacy

Local differential privacy represents promising direction for eliminating trusted coordinator assumptions, with each institution independently perturbing contributions before transmission. Adapting APBA algorithm to LDP settings requires fundamentally different heterogeneity estimation protocols operating on already-noised institutional statistics. Personalized local differential privacy allows heterogeneous privacy preferences across institutions, creating complex optimization spaces where both privacy levels and budget distributions require joint optimization.

5.2.3 Integration with Additional Privacy-Enhancing Technologies

Combining differential privacy with secure multi-party computation, homomorphic encryption, or trusted execution environments may enable enhanced privacy guarantees or computational efficiency improvements. Federated analytics extending beyond model training to include privacy-preserving data exploration and cohort discovery represent valuable application domains for adaptive privacy budget allocation principles. Extending APBA framework to broader federated analytics scenarios requires adapting heterogeneity measurements and allocation strategies to diverse query workloads beyond iterative gradient descent optimization.

References

[1]. Chen, L., Yue, D., Ding, X., Wang, Z., Choo, K. K. R., & Jin, H. (2023). Differentially private deep learning with dynamic privacy budget allocation and adaptive optimization. IEEE Transactions on Information Forensics and Security, 18, 4422-4435.

[2]. Yu, Z., Lu, Z., Lu, S., Cui, Y., Tang, X., & Wu, J. (2024, December). Adaptive Differential Privacy via Gradient Components in Medical Federated Learning. In 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) (pp. 3929-3934). IEEE.

[3]. Shukla, A., Pokhariya, H. S., Michaelson, J., Srivastava, A. P., Narayanamma, L., & Srivastava, A. (2023, December). Enhancing security and privacy in cloud–based healthcare data through machine learning. In 2023 International Conference on Artificial Intelligence for Innovations in Healthcare Industries (ICAIIHI) (Vol. 1, pp. 1-7). IEEE.

[4]. Qin, Z., Wang, D., & Wang, M. (2024, December). Dynamic Differential Privacy in Hierarchical Federated Learning: A Layerwise Adaptive Framework. In 2024 IEEE 23rd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom) (pp. 2207-2212). IEEE.

[5]. Lyu, M., Ni, Z., Chen, Q., & Li, F. (2024). Edge-DPSDG: An edge-based differential privacy protection model for smart healthcare. IEEE Transactions on Big Data, 11(1), 21-34.

[6]. Farooqi, S. A., Abd Rahman, A., & Saad, A. (2024, January). Differential privacy based federated learning techniques in IoMT: A review. In 2024 18th International Conference on Ubiquitous Information Management and Communication (IMCOM) (pp. 1-7). IEEE.

[7]. Chen, C., Hu, X., Li, Y., & Tang, Q. (2023). Optimization of privacy budget allocation in differential privacy-based public transit trajectory data publishing for smart mobility applications. IEEE Transactions on Intelligent Transportation Systems, 24(12), 15158-15168.

[8]. Yang, S., Li, N., Sun, D., Du, Q., & Liu, W. (2021, September). A differential privacy preserving algorithm for greedy decision tree. In 2021 2nd International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE) (pp. 229-237). IEEE.

[9]. Husnoo, M. A., Anwar, A., Chakrabortty, R. K., Doss, R., & Ryan, M. J. (2021). Differential privacy for IoT-enabled critical infrastructure: A comprehensive survey. IEEE access, 9, 153276-153304.

[10]. Liu, J., Chang, Z., Wang, K., Zhao, Z., & Hämäläinen, T. (2024). Energy-efficient and privacy-preserved incentive mechanism for mobile edge computing-assisted federated learning in

healthcare system. IEEE Transactions on Network and Service Management, 21(4), 4801-4815.

[11]. Liu, R., Lee, H. K., Bhavani, S. V., Jiang, X., Ohno-Machado, L., & Xiong, L. (2024, October). Patient-Centered and Practical Privacy to Support AI for Healthcare. In 2024 IEEE 6th International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications (TPS-ISA) (pp. 265-272). IEEE.

[12]. Xin, W., Jiaqian, L., Xueshuang, D., Haoji, Z., & Lianshan, S. (2024). A survey of differential privacy techniques for federated learning. IEEE Access.