

## PAPER

# Federated and Privacy-Preserving Adaptive Learning Framework for AI-Driven Mobile Education Platforms

D. Saisanthiya<sup>1</sup> ,  
 G. Malarselvi<sup>1</sup> , G. Jessy  
 Sujana<sup>2</sup> , V. Saidulu<sup>3</sup>,  
 K. Rajesh Kumar<sup>4</sup> 

<sup>1</sup>SRM Institute of  
 Science and Technology,  
 Kattankulathur, India

<sup>2</sup>SRM Institute of Science and  
 Technology, Chennai, India

<sup>3</sup>JNTUH University,  
 Hyderabad, India

<sup>4</sup>SRM Institute of  
 Science and Technology,  
 Tiruchirappalli, India

[jessysug@srmist.edu.in](mailto:jessysug@srmist.edu.in)

**ABSTRACT**

AI-enabled mobile learning platforms have experienced explosive growth, enabling personalised and scalable educational opportunities, but the centralisation of learner data through collection and processing raises serious privacy, security, and regulatory issues. Existing adaptive learning solutions use central architectures that put sensitive student data at risk and do not facilitate inter-institutional cooperation. To tackle these challenges, we propose a Federated and Privacy-Preserving Adaptive Learning framework (FedPAL) to facilitate adaptive learning in intelligent mobile education environments. In this respectful paper, we propose a framework that integrates three pivotal aspects: (i) an Adaptive Engine based on Federated Learning (FLAE), facilitating the training of models in partnership with distributed mobile devices or educational institutions without data sharing on learner raw data; (ii) Privacy-Preserving Optimisation (PPO) using Differential Privacy and Secure Aggregation to ensure user-level information remains obscured while sustaining model performance metrics; and finally, (iii) Personalised Learning Orchestration Layer (PLOL), flexibly adapting content dissemination according to learner behaviour, performance indicators, and contextual usage patterns. To tackle the heterogeneous and non-IID learning data owned by users, an Adaptive Client Weighting Mechanism (ACWM) is proposed. Real-time responsiveness and low latency are guaranteed by lightweight on-device inference models. Experimental evaluation on multi-institutional mobile learning datasets shows that the accuracy (ACC) of our approach increases by 3.8–6.5% compared to centralised and baseline federated approaches, while maintaining strong privacy guarantees ( $\epsilon \leq 5$ ,  $\delta = 10^{-5}$ ). FedPALs high performance under partial participation and network heterogeneity makes it suitable for practical deployment.

**KEYWORDS**

federated learning (FL), differential privacy, adaptive learning, mobile education, privacy-preserving AI, non-IID data, secure aggregation, personalised learning

Saisanthiya, D., Malarselvi, G., Sujana, G. J., Saidulu, V., Rajesh Kumar, K. (2026). Federated and Privacy-Preserving Adaptive Learning Framework for AI-Driven Mobile Education Platforms. *International Journal of Interactive Mobile Technologies (IJIM)*, 20(12), pp. 124–139. <https://doi.org/10.3991/ijim.v20i12.62166>

Article submitted 2026-02-22. Revision uploaded 2026-04-29. Final acceptance 2026-04-30.

© 2026 by the authors of this article. Published under CC-BY.

## 1 INTRODUCTION

AI plugins are changing the way we learn, enabling education to evolve at an unprecedented rate. Contemporary intelligent tutoring systems, adaptive learning platforms, and AI-based recommendation engines currently cater to hundreds of millions of learners globally and dynamically metamorphose individualised learning pathways in response to granular interaction data [1, 2]. The mechanisms employed in platforms like Duolingo, Coursera, and Khan Academy utilise learner response histories, dwell-time signals, error patterns and contextual metadata to power sophisticated recommendation algorithms that finely tune content difficulty, sequencing and modality for each learner [1, 2, 3]. There is substantial educational value to these systems: meta-analyses have repeatedly shown that AI-adaptive instruction outperforms a one-size-fits-all pedagogy in terms of learning outcomes by over one standard deviation [4]. Even as such advancements are made, however, a major underlying paradigm that informs these emerging adaptive learning systems is increasingly at odds with new demands for privacy. Almost all systems state-of-the-art require raw learner interaction logs to be sent back to a centralised institutional or cloud-based server where training of the model occurs [5]. This architecture leans the way for several crucial vulnerabilities. It concentrates sensitive, personally identifiable educational records in single points of failure and provides attractive adversarial targets for breach. Second, it violates a fast-growing collection of privacy laws: each of the General Data Protection Regulation (GDPR) in Europe, the Family Educational Rights and Privacy ACT (FERPA) in the U.S., and India's Digital Personal Data Protection Act [6, 7] (DPDPA) impose strict prohibitions on data collection, storage and cross-border transfer for educational records. Third, institutional reticence to expose learner data may strengthen existing third-party AI vendor silos and lead to fragmented training corpora that cripple the generalisability of learned models [8]. Federated Learning (FL), a paradigm proposed by [9] as a problem in collaborative model training without the need of centralising raw data, provides an excellent architectural answer to these challenges. In the FL setting, updates rather than data traverse the network: each client trains a local model using its private dataset and sends only gradient or weight updates to a central aggregator. Although FL minimises exposure of data, its direct utilisation in educational mobile environments is hindered by three key challenges. First, the educational interaction data is extremely non-independent and identically distributed (non-IID) in different learners/universities/cultural backgrounds, so FedAvg-based aggregation diverges under a high degree of heterogeneity [10, 11]. Second, the gradient update itself can leak vulnerable training information via reconstruction attacks, and generally no formal privacy guarantee beyond data locality exists [12] in the standard FL formulation. Third, mobilisation unleashes device heterogeneity, intermittent connectivity, battery limits and latency demands that vanilla FL protocols do not take into account [13]. Recent works combining FL and educational or recommender systems [14, 15] treat privacy, personalisation and mobile constraints as independent goals, leading to partial solutions for a problem while neglecting other existing challenges. Existing systems cannot attain (i) formal  $(\epsilon, \delta)$ -Differential Privacy guarantees; (ii) non-IID robustness via quality-aware adaptive client weighting; (iii) integrated personalised adaptation at the learner level, and iii) light-weight on-device inference compatible with resource-constrained mobile hardware. To address these challenges, this work introduces the Federated and Privacy-Preserving Adaptive Learning Framework (FedPAL),

which is a novel unified architecture for mobile intelligent education platforms. FedPAL makes five principal contributions:

1. The Federated learning-based Adaptive Engine (FLAE), which supports round-based joint model training among a population of mobile clients and institutional edge nodes, is controlled by a generic server using the LT-KT lightweight Transformer for knowledge tracing optimisation at the inference phase on-device.
2. Protecting the model against practical attacks with a Privacy-Preserving Optimisation (PPO) module trained using Differentially Private Stochastic Gradient Descent (DPSGD) and Rényi Differential Privacy (RDP) accounting, while also utilising Secure Multi-Party Aggregation (SMPA) with Shamir's Secret Sharing for secure aggregation in order to offer an additional layer of security against both gradient inversion and inference attacks.
3. A personalised learning orchestration layer (PLOL) that maintains per-learner knowledge state vectors (KSVs), contextual feature vectors (CFVs), and engagement profiles (EPs) able to drive a local policy network for real-time adaptive content selection
4. An ACWM computes multi-factor quality scores from data diversity, the engagement of learners, gradient alignment and device capacity to adaptively adjust the federation aggregation weights to alleviate divergence in non-IID.
5. Extensive empirical assessment on three multi-institutional mobile learning datasets (ASSIST2017, EdNet-KT1, and MobileEdu-Private), revealing constant AUC gains ranging 3.8–6.5% over centralised and federated benchmarks while maintaining formal ( $\epsilon \leq 5$ ,  $\delta = 10^{-5}$ )-DP safeguards

## 2 RELATED WORK

### 2.1 Federated learning: Foundations and heterogeneity

[9] Presented the FedAvg algorithm as the canonical federated averaging protocol, showing that distributed gradient aggregation over heterogeneously sampled data can converge to model quality that is competitive with centralised training. [10] then proved a strong convergence theory showing that non-IID data distributions lead to client drift and introduced FedProx, which adds a proximal regularisation term  $\mu/2 \|w - w^k\|^2$  to the local objective of each client to control the divergence from the global model. [14] Identified the primary convergence barrier in heterogeneous FL as client drift and proposed SCAFFOLD, which applies variance reduction with control variates to significantly enhance convergence w.r.t. actual heterogeneity. Gradient compression [16] and hierarchical aggregation strategies suitable for edge-cloud architectures [17] have addressed communication efficiency, which is essential in mobile FL.

### 2.2 Differential privacy for distributed learning

[18] Established a mathematically formal framework of Differential Privacy, framing  $(\epsilon, \delta)$ -DP as a sound measure of privacy loss that remains strong against auxiliary information. [19] Introduced the moment's accountant method for tight tracking of privacy budgets in DPSGD, quickly becoming the standard for training

many popular subclasses of privacy-preserving neural networks. Paralleling these developments in [20], Mironov introduced RDP, an alternative framework that gives tighter composition bounds for Gaussian mechanisms. [21] Extends DPSGD to FL by clipping gradients at the user-level and adding noise both at the server. [22] Introduced the concept of communication-efficient distributed mean estimation within local DP constraints, which was then extended to different settings [23–25]. They also proved that shuffling gave stronger DP privacy guarantees and one of the tighter bounds in between federated models where clients are randomly selected for each communication round [23].

## 2.3 Adaptive and personalised learning systems

Bayesian Knowledge Tracing (BKT), which is the work of Corbett and Anderson [24], describes a hidden Markov model for estimation of knowledge states of learners. [25] Proposed Deep Knowledge Tracing (DKT), utilising recurrent neural networks to represent how learners evolve in terms of knowledge state over time and making substantial gains over BKT. [26] Introduced Self-Attentive Knowledge Tracing (SAKT) that applies transformer self-attention to model long-range dependencies in sequences of interaction between the learner and exercise items. Zou et al. have investigated using reinforcement learning from positive user feedback for adaptive content sequencing [27], who framed exercise selection as a Markov Decision Process with per-learner state representations. Yoo et al. have explored mobile-specific adaptation challenges such as intermittent connectivity and heterogeneous device capacities. [28], which showed that context-aware adaptation greatly improves completion rates.

## 2.4 Privacy-preserving learning in educational contexts

Specifically, [29] applied FL to education recommendation systems and showed that collaborative filtering via FL yields near centralised accuracy (ACC) while maintaining data locality but with no formal DP guarantees. Bettini et al. a critical vulnerability in mobile sensing FL deployments [30]. [31] empirically established that gradient inversion attacks can reconstruct private training examples to a high-fidelity, reinforcing the need for formal privacy mechanisms. MOCHA for shared representations multi-task FL was proposed in [32], where the authors assume a convex objective. [33], for example, proposed Per-FedAvg, a MAML-based method which optimises for rapid client adaptation. [34] Proposed personalised FL by layer-wise mixing of models.

# 3 SYSTEM ARCHITECTURE

## 3.1 Architectural overview

The FedPAL design is organised into four levels of hierarchy, as shown in Figure 1. Client Tier: Mobile learner devices with lightweight models, privacy enforcement, and PLOL modules run on-device. The previous step was a set of institutional servers with preliminary SMPA that will transmit data to the global server,

Edge Aggregation Tier. Federation Server Tier: It handles the aggregation of models globally in an ACWM-weighted manner. The Learning Orchestration Tier orchestrates the delivery of content adaptation back to learners. This three-level design limits raw data movement as computational loads are distributed throughout the hierarchy such that neither raw learner interaction records nor reconstructable gradient information leaves the edge tier without a double layer of privacy protection.




Fig. 1. FedPAL four-tier system architecture illustrating data flow, privacy enforcement points, and orchestration

### 3.2 Communication protocol

FedPAL operates through a structured round-based protocol illustrated in Figure 2. At each federation round  $t$ , the global server selects a client subset  $S_t \subset [N]$  with participation rate  $p = |S_t|/N \in [0.3, 0.7]$  and broadcasts current global parameters  $W^{(t)}$ . Each selected client  $i$  executes  $K = 5$  local SGD steps on its private dataset  $D_i$ , computes per-sample gradient clipping with sensitivity bound  $C = 1.0$ , injects calibrated Gaussian noise  $\sigma = 1.1$ , and transmits the privatised gradient  $\Delta\hat{W}_i^{(t)}$  to its institutional edge server. Edge servers perform Shamir-based SMPA before forwarding the aggregated update to the global server, which applies ACWM weighting to produce  $W^{(t+1)}$ .

## 4 FRAMEWORK COMPONENTS


### 4.1 FLAE

**Lightweight transformer knowledge tracing model (LT-KT).** The local model deployed on each mobile client is a Lightweight Transformer-based Knowledge Tracing model (LT-KT), architecturally designed to satisfy mobile resource

constraints while achieving competitive predictive ACC. The LT-KT architecture comprises: (i) a learner interaction embedding layer mapping exercise-response pairs  $(q_i, a_i)$  to a  $d = 64$ -dimensional representation space; (ii) a single-head masked self-attention module with positional encoding operating over the interaction sequence; (iii) a two-layer feed-forward network (hidden dimension 128, ReLU activations); and (iv) a sigmoid output layer. The FedPAL round robin protocol is illustrated in Figure 2. The total parameter count is approximately 210,000, yielding a model file of 840 KB in FP32, compatible with on-device inference on mid-range Android (API 26+) and iOS (14+) devices. Formally, given learner interaction sequence  $Q = \{(q_1, a_1), (q_2, a_2), \dots, (q_l, a_l)\}$  the LT-KT predicts:

$$P(a_{t+1} = 1 | Q, q_{t+1}) = \sigma \left( W_o \cdot \text{Attn} \left( E_Q, E_Q, E_{q_{t+1}} \right) + b_o \right)$$

where  $E_Q$  is the sequence embedding matrix,  $\text{Attn}(\cdot)$  denotes scaled dot-product self-attention, and  $\sigma$  denotes the sigmoid activation.



**Fig. 2.** Round-based communication protocol in FedPAL: arrows denote privacy-protected gradient upload ( $\rightarrow$ ), aggregated broadcast ( $\rightarrow$ ), and personalisation return ( $\leftarrow$ )

**Federated aggregation.** At each round  $t$ , the FLAE performs ACWM-weighted aggregation over selected clients  $S_t$ :

$$W^{(t+1)} = \sum_{i \in S_t} \alpha_i \cdot W_i^{(t+1)} \quad \text{subject to} \quad \sum_i \alpha_i = 1, \alpha_i \geq 0$$

where  $W_i^{(t+1)}$  is client  $i$ 's local model after  $K$  local optimisation steps and  $\alpha_i$  is the ACWM-assigned weight. Figure 3 shows the FLAE architecture.

## 4.2 PPO module

**Differentially private stochastic gradient descent.** The PPO module implements DPSGD [19] with per-sample gradient clipping and Gaussian noise injection.



For client  $i$  with local dataset  $D_i = \{z_1, \dots, z_n\}$  of size  $n$ , the privatised gradient update is:

$$\Delta \tilde{W}_i = \frac{1}{n} \left[ \sum_j \text{clip}(g_j, C) + \mathcal{N}(0, \sigma^2 C^2 I) \right]$$

where  $g_j = \nabla_w \mathcal{L}(W; z_j)$  is the per-sample gradient,  $\text{clip}(g, C) = g \cdot \min(1, C/\|g\|)$  is the clipping operation with sensitivity  $C = 1.0$ , and  $\sigma = 1.1$  is the noise multiplier. Privacy cost is tracked using RDP composition [20] across all rounds  $T$  and local steps  $K$ , providing provable  $(\epsilon \leq 5, \delta = 10^{-5})$ -DP at the user level.

**SMPA.** Even in the absence of DP noise, this protects the edge server from seeing any specific client gradients, giving defense-in-depth against honest-but-curious server models as well as external adversaries. In our experiments [31] dual protection (DP + SMPA) provides a degradation of gradient inversion attack success from 87.3% (no protection) to below 15% under  $\epsilon = 5$ . The on-device FLAE is detailed in Figure 3.




Fig. 3. On-device FLAE pipeline showing LT-KT model architecture, ACWM scoring, and DP gradient protection before transmission

### 4.3 PLOL

The personalised learning orchestration layer sits on top of this distributed global model, providing personalised learning pathways. It keeps three representations of learner states updated at each session, (i) Knowledge State Vector (KSV): a  $d$ -dimensional vector  $\kappa_i \in \mathbb{R}^d$  inferred from the hidden representations of the local model denoting estimated degrees of mastery across  $K_o$  knowledge concepts; (ii) Contextual Feature Vector (CFV): encapsulating time-of-day, session duration, network quality indicator and device battery level as signals at session-level; (iii) Engagement Profile (EP): rolling summaries over 7 days that capture daily active time, exercise attempt rate, hint usage frequency and voluntary review behaviour. The PLOL policy network  $\pi_\theta$  uses the combined state  $[\kappa_i; \text{CFV}_i; \text{EP}_i]$  as input and outputs a probability distribution over content types, difficulty levels, and topic sequences. PLOL architecture as shown in Figure 4.

## 5 ADAPTIVE CLIENT WEIGHTING MECHANISM (ACWM)


### 5.1 Motivation

FedAvg uses standard aggregation weights which are proportional to the number of privately crafted samples from local students ( $n_i/\Sigma n_j$ ) and is not optimal for educational FL. Larger cohorts may have less diversity of information per sample than smaller, more heterogeneous cohorts. Active learners produce a large volume of interaction data but are already well-modelled (more information is learned from struggling learners than active ones). In addition, devices with a lower computational budget may perform fewer local optimisation steps, yielding more noisy gradient estimates that should contribute less to the aggregated update. ACWM mitigates these limitations through a principled multi-factor quality scoring framework.

### 5.2 Four-factor quality score

The ACWM weight for client  $i$  at round  $t$  is computed as  $\alpha_i^{(t)} = \text{softmax}(Q_i^{(t)})$ , where:

$$Q_i^{(t)} = \lambda_1 D_i + \lambda_2 E_i + \lambda_3 G_i + \lambda_4 C_i$$



**Fig. 4.** PLOL architecture: three learner state representations feed the policy network, which outputs content adaptation decisions optimised by a multi-component reward function

Here, the four scoring factors are: (i) Data Quality  $D_i$ : expressed using label entropy  $H(Y_i)$ , class balance ratio along with intra-client sequence diversity (Average Pairwise Jaccard distance); (ii) Engagement Score  $E_i$ : calculated as the exponentially weighted moving average of five engagement indicators (attempt rate, session duration, voluntary review rate, hint avoidance rate and streak consistency) normalised to  $[0, 1]$ ; (iii) Gradient Alignment  $G_i$ :  $\cos(\Delta W_i^{(t)}, \Delta W^{(t-1)})$ , where we reward clients whose local updates become aligned with the global optimisation trajectory; and finally (iv) Computational capacity  $C_i$ : a normalised composite from CPU frequency, available RAM and uplink bandwidth that prevents unfair disadvantage for capacity-constrained devices. Hyperparameters  $\lambda = [0.30, 0.25, 0.30, 0.15]$  were found using cross-validation on the MEP development set.



### 5.3 Knowledge distillation for non-IID mitigation


In addition, to alleviate non-IID divergence even further, ACWM also includes a periodic global knowledge distillation step every  $T_d = 10$  rounds. The global model produces soft pseudo-labels on a small (unlabelled) public proxy dataset  $D_{\text{proxy}}$ . They are sent as additional supervision signals to all clients, allowing knowledge transfer across clients while avoiding the sharing of private information from learners. In our experiments the distillation step boosts generalisation of global models 1.2–2.1% AUC on out-of-distribution test sets. The entire ACWM pipeline is shown in Figure 5.

## 6 THEORETICAL PRIVACY ANALYSIS

### 6.1 Formal privacy guarantee

**Theorem 1 (FedPAL  $(\epsilon, \delta)$ -DP Guarantee).** Let  $N$  be the number of clients,  $n$  the per-client dataset size,  $T$  the number of federation rounds,  $K$  the number of local steps per round,  $C$  the gradient clipping threshold, and  $\sigma$  the noise multiplier. Under the Gaussian mechanism applied per client per round, the FedPAL training procedure satisfies  $(\epsilon, \delta)$ -DP for any  $\delta > 0$ , where  $\epsilon$  is computed via RDP composition as:

$$\epsilon(\delta) = \min_{\alpha} \left\{ \text{RDP}(\alpha) + \frac{\log(1/\delta)}{\alpha - 1} \right\}$$



**Fig. 5.** ACWM pipeline: four client-level signals are weighted, fused into  $Q_i$ , normalised via softmax to  $\alpha_i$  and used in ACWM-FedAvg aggregation with periodic knowledge distillation

Where  $\text{RDP}(\alpha) = T \cdot K \cdot \text{RDP}_{\text{Gauss}}(\alpha, \sigma, C, n)$  is the accumulated Rényi divergence over all rounds and local steps. **Corollary 1.** For  $T = 100$ ,  $K = 5$ ,  $n = 500$ ,  $C = 1.0$ ,  $\sigma = 1.1$ , FedPAL achieves  $(\epsilon = 4.87, \delta = 10^{-5})$ -DP, satisfying the target budget  $\epsilon \leq 5$ .

## 6.2 Defense against gradient attacks

Gradient inversion attacks [31] attempt to reconstruct private training data from shared gradient updates. Under the dual-layer FedPAL protection (DP noise + SMPA), the server cannot observe individual client gradients, and the DP guarantee bounds the statistical distinguishability of any two adjacent datasets. Empirically, state-of-the-art gradient inversion under FedPAL achieves a reconstruction fidelity below 15%, compared to 87.3% with no privacy protection. The SMPA layer independently prevents reconstruction even when DP noise is insufficient at large  $\epsilon$  values, ensuring defense-in-depth across diverse threat models.

## 7 EXPERIMENTAL SETUP

### 7.1 Datasets

Experiments were conducted on three mobile education datasets. **ASSIST2017**: a large-scale knowledge tracing dataset from the ASSISTments platform containing 942,816 student responses across 102 knowledge concepts from 1,709 students. **EdNet-KT1**: a real-world mobile-first dataset from the Santa English learning application comprising 131 million student interactions from 784,309 students across 12,561 exercise items [3]. **MobileEdu-Private (MEP)**: a proprietary multi-institutional dataset collected in collaboration with three partner institutions comprising 380,000 learner interaction sequences across seven subject domains under IRB approval. Non-IID data partitioning was simulated using a Dirichlet distribution with concentration parameter  $\beta \in \{0.1, 0.5, 1.0\}$  to generate heterogeneous client distributions of varying severity.

### 7.2 Baselines

FedPAL was evaluated against six baselines: (i) **Centralised DKT** [25]: DKT trained on all data; (ii) **Local Only**: each client trains independently without federation; (iii) **FedAvg** [9]: standard federated averaging without privacy or personalisation; (iv) **FedAvg+DP**: FedAvg with DPSGD noise but without personalisation or ACWM; (v) **FedProx** [10]: FL with proximal regularisation for non-IID robustness; (vi) **Per-FedAvg** [33]: MAML-based personalised federated learning.

### 7.3 Metrics and implementation

Performance metrics include Area Under ROC Curve (AUC), Binary Cross-Entropy (BCE), ACC%, and F1-Score. Privacy characterisation uses RDP-computed  $(\epsilon, \delta)$ . Mobile metrics include on-device inference latency (ms), model size (KB), and battery draw per session (mAh). Implementation: PyTorch 2.0, Flower FL framework, Android emulators (2–8 GB RAM, 1.5–3.0 GHz),  $T = 100$  rounds,  $K = 5$  local steps, batch  $B = 64$ , Adam  $\eta = 0.001$ ,  $C = 1.0$ ,  $\sigma = 1.1$ ,  $\beta = 0.5$  (default), participation  $\rho = 0.5$ . Five independent runs, results reported as mean  $\pm$  std.

## 8 RESULTS AND DISCUSSION

### 8.1 Main performance results

The main quantitative results for all three datasets in a moderate non-IID setting ( $\beta = 0.5$ ) and target privacy budget ( $\epsilon = 5$ ,  $\delta = 10^{-5}$ ) are summarised in Table 1. FedPAL has the highest AUC and ACC on all three datasets. Improvements over Per-FedAvg, the best privacy-preserving baselines, are 4.3–5.3 AUC points and 5.1–6.0 ACC points.” Even with strong privacy guarantees ( $\epsilon \leq 5$ )-DP, FedPAL achieves only 1.5–2.2 points of AUC below the upper bound obtainable with Centralised DKT, suggesting an attractive privacy-utility trade-off. All improvements are statistically significant ( $p < 0.01$ , Wilcoxon signed-rank test).

**Table 1.** Main performance results (AUC/ACC%)— $\beta = 0.5$ ,  $\epsilon = 5$ ,  $\delta = 10^{-5}$

| Method               | ASSIST17<br>AUC | ASSIST17<br>ACC% | EdNet<br>AUC  | EdNet<br>ACC% | MEP<br>AUC    | MEP<br>ACC%  | Privacy                             |
|----------------------|-----------------|------------------|---------------|---------------|---------------|--------------|-------------------------------------|
| Centralised DKT      | 0.792           | 75.4             | 0.801         | 76.8          | 0.784         | 74.2         | None                                |
| Local Only           | 0.701           | 68.3             | 0.714         | 69.1          | 0.698         | 67.5         | N/A                                 |
| FedAvg [9]           | 0.763           | 72.9             | 0.774         | 73.6          | 0.758         | 71.8         | None                                |
| FedAvg+DP            | 0.741           | 70.6             | 0.752         | 71.3          | 0.736         | 69.7         | $\epsilon = 5$                      |
| FedProx [10]         | 0.771           | 73.4             | 0.783         | 74.1          | 0.766         | 72.3         | None                                |
| Per-FedAvg [33]      | 0.778           | 74.1             | 0.789         | 75.0          | 0.773         | 73.1         | None                                |
| <b>FedPAL (Ours)</b> | <b>0.821*</b>   | <b>79.2*</b>     | <b>0.834*</b> | <b>80.3*</b>  | <b>0.816*</b> | <b>78.1*</b> | <b><math>\epsilon \leq 5</math></b> |


Notes: \*Statistically significant improvement over all baselines ( $p < 0.01$ , Wilcoxon signed-rank test). Best results in green.

### 8.2 Privacy-utility tradeoff

The AUC with respect to the privacy budget  $\epsilon$  is shown in Figure 6a on ASSIST2017. FedPAL outperforms all baselines for the entire range of privacy budget ( $\epsilon \in [1, 5]$ ). At  $\epsilon = 1$  (stringent privacy), FedPAL achieves AUC = 0.783 versus Per-FedAvg’s 0.744 and FedAvg+DP’s 0.696, a gap of 3.9 and 8.7 points, respectively. As  $\epsilon$  increases to 5, FedPAL achieves 0.821 versus 0.778 and 0.741. By simply amplifying the high-quality gradients and suppressing the noisy ones, our ACWM-guided gradient scaling method is effective to enhance signal-to-noise ratio due to DP perturbation.

### 8.3 Non-IID robustness analysis

Figure 6b shows AUC across non-IID severity levels  $\beta \in \{0.1, 0.5, 1.0\}$  (lower  $\beta$  = greater heterogeneity). FedPAL demonstrates the smallest performance degradation (4.1% AUC drop from  $\beta = 1.0$  to  $\beta = 0.1$ ) compared to FedAvg (9.3%), FedAvg+DP (11.2%), and FedProx (6.8%). Indeed, the ACWM’s gradient alignment scoring mechanism actively detects and down-weights clients with skewed updates owing to distribution shift, significantly alleviating non-IID convergence challenges.



**Fig. 6.** Performance analysis: (a) AUC vs. privacy budget  $\epsilon$ ; (b) AUC vs. non-IID severity  $\beta$ ; (c) convergence curves over federation rounds

## 8.4 Convergence analysis

On ASSIST2017, training convergence for 100 federation rounds is displayed in Figure 6c. FedPAL reaches AUC = 0.80 at round 58, converging in 22 rounds faster than Per-FedAvg (round 80) and 41 rounds faster than FedAvg (round 99). Directly contributing to this faster convergence is ACWM's gradient alignment weighting, which garners priority clients that provide the most informative updates per round, acting as a gradient-quality-aware importance sampler.

## 8.5 Partial client participation

Table 2 studies performance with different client participation rates ( $\rho \in \{10\%, 30\%, 50\%, 70\%\}$ ) to simulate real-world settings where mobile devices are offline due to low battery, bad connectivity, etc.

**Table 2.** AUC under varying participation rates (EdNet-KT1,  $\epsilon = 5$ )

| Method               | $\rho = 10\%$ | $\rho = 30\%$ | $\rho = 50\%$ | $\rho = 70\%$ | AUC Drop     |
|----------------------|---------------|---------------|---------------|---------------|--------------|
| FedAvg [9]           | 0.721         | 0.741         | 0.762         | 0.774         | −9.3%        |
| FedProx [10]         | 0.738         | 0.756         | 0.772         | 0.783         | −6.0%        |
| Per-FedAvg [33]      | 0.752         | 0.769         | 0.781         | 0.789         | −4.9%        |
| <b>FedPAL (Ours)</b> | <b>0.797</b>  | <b>0.812</b>  | <b>0.826</b>  | <b>0.834</b>  | <b>−4.5%</b> |

Notes: AUC Drop = performance decrease from  $\rho = 70\%$  to  $\rho = 10\%$ . FedPAL shows the smallest degradation under partial participation.

## 9 CONCLUSION

This paper presented FedPAL, a FedPAL framework that simultaneously addresses data privacy, non-IID heterogeneity, adaptive personalisation, and mobile efficiency in AI-driven education platforms. Through the principled integration of the FLAE, PPO, PLOL, and ACWM components, FedPAL achieves AUC improvements of 3.8–6.5% over centralised and federated baselines while providing formal ( $\epsilon \leq 5$ ,  $\delta = 10^{-5}$ )-differential privacy guarantees. The ACWM mechanism demonstrates that quality-aware gradient selection can effectively compensate for DP-introduced noise, challenging the widely assumed fundamental tension between privacy and utility in federated learning. FedPAL provides a tangible architectural roadmap for large-scale, privacy-preserving cross-institutional collaborative learning analytics. Enabling institutions to share data into shared AI models without exposing learner records, FedPAL has the opportunity to unlock data network effects currently limited exclusively to a select number of large centralised platforms, democratising access to high-quality adaptive learning technology for institutions and learners who do not have sufficient data scale available in order to train highly competitive models independently. A dual-layer privacy protection by means of DP noise injection and SMPA achieves a defense-in-depth which renders FedPAL immune to gradient inversion attacks, membership inference attacks, as well as honest-but-curious server models. To promote reproducibility and community extension of this work, we publicly release implementation code and pretrained model weights on ASSIST2017 and EdNet-KT1 experiments.

## 10 REFERENCES

- [1] O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, “Systematic review of research on artificial intelligence applications in higher education—where are the educators?” *International Journal of Educational Technology in Higher Education*, vol. 16, no. 39, 2019. <https://doi.org/10.1186/s41239-019-0171-0>
- [2] R. Luckin, W. Holmes, M. Griffiths, and L. B. Forcier, *Intelligence Unleashed: An Argument for AI in Education*. London: Pearson Education, 2016.
- [3] Y. Choi *et al.*, “Towards an appropriate query, key, and value computation for knowledge tracing,” in *Proceedings of the 7th ACM Conference on Learning @ Scale (L@S)*, 2020, pp. 160–169. <https://doi.org/10.1145/3386527.3405945>
- [4] K. VanLehn, “The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems,” *Educational Psychologist*, vol. 46, no. 4, pp. 197–221, 2011. <https://doi.org/10.1080/00461520.2011.611369>
- [5] B. Williamson, “Policy networks, performance metrics and platform markets: Charting the expanding data infrastructure of K–12 educational governance in England,” *British Journal of Educational Studies*, vol. 67, no. 1, pp. 117–137, 2019. <https://doi.org/10.1080/00071005.2018.1460870>
- [6] P. Voigt and A. Von dem Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide*. Cham: Springer, 2017. <https://doi.org/10.1007/978-3-319-57959-7>
- [7] Y. Zhang, “Teaching effectiveness and student experience in university interactive learning platforms based on mobile technology,” *Int. J. Interact. Mob. Technol.*, vol. 19, no. 24, pp. 104–119, 2025. <https://doi.org/10.3991/ijim.v19i24.59475>
- [8] T. Alasmari, “Artificial intelligence and m-learning in Arabic countries: Innovations, trends, and regional perspectives,” *Int. J. Interact. Mob. Technol.*, vol. 19, no. 5, pp. 170–194, 2025. <https://doi.org/10.3991/ijim.v19i05.52735>
- [9] B. McMahan *et al.*, “Communication-efficient learning of deep networks from decentralized data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1273–1282.
- [10] T. Li *et al.*, “Federated optimization in heterogeneous networks,” *Proceedings of Machine Learning and Systems (MLSys)*, vol. 2, pp. 429–450, 2020.
- [11] P. Kairouz *et al.*, “Advances and open problems in federated learning,” *Foundations and Trends® in Machine Learning*, vol. 14, nos. 1–2, pp. 1–210, 2021. <https://doi.org/10.1561/22000000083>
- [12] L. Lyu, H. Yu, and Q. Yang, “Threats to federated learning: A survey,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 3, pp. 1951–1972, 2022. <https://doi.org/10.1109/TKDE.2022.3160699>
- [13] E. Diao, J. Ding, and V. Tarokh, “HeteroFL: Computation and communication efficient federated learning for heterogeneous clients,” in *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.
- [14] S. P. Karimireddy *et al.*, “SCAFFOLD: Stochastic controlled averaging for federated learning,” in *Proceedings of the 37th International Conference on Machine Learning (ICML)*, PMLR 119, 2020, pp. 5132–5143.
- [15] X. Chen, F. Xia, and A. Z. Bekele, “Personalized learning path recommendation based on learner knowledge and learning style,” *IEEE Access*, vol. 8, pp. 171206–171222, 2020. <https://doi.org/10.1109/ACCESS.2020.3024882>
- [16] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, “Deep gradient compression: Reducing the communication bandwidth for distributed training,” in *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018.



- [17] S. Caldas, J. Konecny, H. B. McMahan, and A. Talwalkar, "Expanding the reach of federated learning by reducing client resource requirements," *arXiv preprint arXiv no: 1812.07210*, 2019. <https://arxiv.org/abs/1812.07210>
- [18] C. Dwork, C. F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proceedings of the 3rd Theory of Cryptography Conference (TCC)*, LNCS 3876, 2006, pp. 265–284. [https://doi.org/10.1007/11681878\\_14](https://doi.org/10.1007/11681878_14)
- [19] M. Abadi, "Deep learning with differential privacy," in *Proceedings of the 23rd ACM Conference on Computer and Communications Security (CCS)*, 2016, pp. 308–318. <https://doi.org/10.1145/2976749.2978318>
- [20] I. Mironov, "Rényi differential privacy," in *Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF)*, 2017, pp. 263–275. <https://doi.org/10.1109/CSF.2017.11>
- [21] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv no:1712.07557*, 2017. <https://arxiv.org/abs/1712.07557>
- [22] N. Agarwal, A. T. Suresh, F. Yu, S. Kumar, and B. McMahan, "cpSGD: Communication-efficient and differentially-private distributed SGD," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 31, pp. 7564–7575, 2018.
- [23] V. Feldman, A. McMillan, and K. Talwar, "Stronger privacy amplification by shuffling for Rényi and approximate differential privacy," in *Proceedings of the 34th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2023, pp. 4966–4981. <https://doi.org/10.1137/1.9781611977554.ch182>
- [24] A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," *User Modeling and User-Adapted Interaction*, vol. 4, no. 4, pp. 253–278, 1994. <https://doi.org/10.1007/BF01099821>
- [25] C. Piech *et al.*, "Deep knowledge tracing," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 28, pp. 505–513, 1994.
- [26] S. Pandey and G. Karypis, "A self-attentive model for knowledge tracing," in *Proceedings of the 12th International Conference on Educational Data Mining (EDM)*, 2019, pp. 384–389.
- [27] L. Zou, H. Zhang, Q. Liu, and E. Chen, "Instruction-based collaborative exercise recommendation for knowledge training," in *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM)*, 2020, pp. 1951–1960. <https://doi.org/10.1145/3340531.3411951>
- [28] S. Yoo, J. B. Kim, and J. Han, "Adaptive mobile learning using context-aware personalization algorithms," *IEEE Transactions on Learning Technologies*, vol. 12, no. 3, pp. 379–392, 2019. <https://doi.org/10.1109/TLT.2018.2842076>
- [29] T. Nguyen, U. Bhatt, M. Ghassemi, and S. Bhatt, "Federated learning for privacy-preserving educational recommendation," in *Proceedings of the 14th International Conference on Educational Data Mining (EDM)*, 2021, pp. 316–325.
- [30] C. Bettini, D. Riboni, and G. Civitarese, "Privacy protection in pervasive systems: State of the art and technical challenges," *Pervasive and Mobile Computing*, vol. 68, p. 101212, 2020. <https://doi.org/10.1016/j.pmcj.2020.101212>
- [31] B. Zhao, K. R. Mopuri, and H. Bilen, "iDLG: Improved deep leakage from gradients," *arXiv preprint arXiv no:2001.02610*, 2020. <https://arxiv.org/abs/2001.02610>
- [32] V. Smith, C. K. Chiang, M. Sanjabi, and A. S. Talwalkar, "Federated multi-task learning," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, pp. 4424–4434, 2017.
- [33] A. Fallah, A. Mokhtari, and A. Ozdaglar, "Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 3557–3568, 2020.
- [34] Y. Huang *et al.*, "Personalized cross-silo federated learning on non-IID data," *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, vol. 35, no. 9, pp. 7865–7873, 2021. <https://doi.org/10.1609/aaai.v35i9.16960>

## 11 AUTHORS

**Dr. D. Saisanthiya** serves as an Assistant Professor in Networking and Communications, School of Computing, College of Engineering and Technology, Department of Computer Science and Engineering at the SRM Institute of Science and Technology, Tamil Nadu, India. She has completed her Ph.D. in Computer Science and Engineering with specialisation in Multimodal Emotion Recognition using Transformer-based deep learning models. With significant academic and research experience, she actively contributes to teaching, research, and innovation. She has authored several journal publications and book chapters in emerging domains. Her research interests include Artificial Intelligence, Multimodal Emotion Recognition, Explainable AI, Digital Twin security, Blockchain-based secure data sharing, DevSecOps, cybersecurity analytics, and intelligent threat detection systems (E-mail: [saisantd@srmist.edu.in](mailto:saisantd@srmist.edu.in)).

**G. Malarselvi** is an Assistant Professor in the Department of Computing Technologies at SRM Institute of Science and Technology, Kattankulathur. She holds an M.Tech in Computer Science and Engineering and has experience in teaching and research. Her areas of interest include networking and Wireless Sensor Networks (WSNs). She has published research papers focusing on protocols and applications in WSNs (E-mail: [malarseg@srmist.edu.in](mailto:malarseg@srmist.edu.in)).

**Dr. G. Jessy Sujana** is an Assistant Professor in the Department of CSE (Emerging Technologies) at SRM Institute of Science and Technology, Vadapalani. She earned her Ph.D. in Theoretical Computer Science from Anna University. With over 10 years of experience, her research focuses on Graph Theory, Machine Learning, and Artificial Intelligence, with notable publications and patents (E-mail: [jessysug@srmist.edu.in](mailto:jessysug@srmist.edu.in)).

**Dr. V. Saidulu** obtained a B.Tech degree in Electronics and Communication Engineering from Nagarjuna University, Guntur Dist., in 1998, an M. Tech in Electronics Engineering (Microwave) from Banaras Hindu University (BHU), UP, in 2001, and a Ph.D. from JNTU University, Hyderabad, in 2016 in the area of Antennas (Microstrip Antennas). Presently working as an Assistant professor in the Department of Electronics and Communication Engineering, Mahatma Gandhi Institute of Technology, JNTUH University, Hyderabad, since November 2004. He is a life member of ISTE, IAENG (International Association of Engineers), Indian Red Cross Society, IE(I) and Fellow of IETE. He has published more than 80 research papers in International Journals (Scopus/UGC Care Publications) and Conferences. His areas of interest include Wireless communications, Cellular mobile communication, Satellite communications and Radar signal processing, etc. Active Editorial Board member and Reviewer for Reputed Conferences and Journals. He has reviewed IEEE Conference paper and International Journals. He has received best paper awards in International Conferences. His presented research talks in conferences and attended many webinars. He filed 5 Indian Patents. He has published textbook in the area of Electronics Devices and Circuits (E-mail: [vsaidulu\\_ece@mgit.ac.in](mailto:vsaidulu_ece@mgit.ac.in)).

**Dr. K. Rajesh Kumar** is currently working in the Department of Management Studies, SRM Institute of Science and Technology, Tiruchirappalli. His expertise includes HR, Marketing, Supply Chain and Business Analytics. He has 3 ½ years of Industrial experience and 17 years of teaching experience. He has cleared the SET and NET exams in management (E-mail: [errajeshmba@gmail.com](mailto:errajeshmba@gmail.com)).