Data-Aware Layer Assignment for Secure and Efficient Communication in Federated Learning for Medical Image Analysis

Sai Sriram Gonthina[1], Sandip Roy[2], Pronaya Bhattacharya[3], Pratip Rana[4], Sachin Shetty[2]

[1] Department of Computer Science & Engineering, IIIT–Naya Raipur, Raipur, Chhattisgarh, India
[2] Center for Secure & Intelligent Critical Systems, Old Dominion University, Virginia, USA
[3] Department of Computer Science & Engineering, Amity University, Kolkata, India
[4] Department of Computer Science, Old Dominion University, Virginia, USA
gonthina22100@iiitnr.edu.in, sroy@odu.edu, pbhattacharya@kol.amity.edu, prana@odu.edu, sshetty@odu.edu

Abstract

Cross-silo medical imaging federations must contend with strict privacy, limited bandwidth, and non identically distributed (non-IID) data that destabilize training. Current federated learning (FL) architectures either carry the full model (e.g., FedAvg/FedProx) or use naive client/layer pruning and random sampling while ignoring both non-IID heterogeneity and per-layer utility. Based on these limitations, the paper presents a data-aware, layer-wise protocol that aligns communication with expected loss descent while bounding per-round client leverage. Each round, the server estimates per-layer influence from a tiny root set, and clients expose lightweight metadata to form data-quality scores. A capacity-constrained entropic transport matches high-influence layers to high-quality clients under redundancy and temporal coverage. Clients train all layers but upload exactly one with *train-all, send-one* principle. The server then performs per-layer robust aggregation on masked updates via secure aggregation. On the three cross-silo imaging benchmarks of Pneumonia CXR, Brain-Tumor MRI, and ISIC Skin Cancer, it demonstrates a strong threshold free detection quality (AUROC/AUPRC: 0.925/0.935, 0.996/0.988, 0.834/0.852, respectively) while also reducing the per round up-link by $\approx 1/n$ with respect to FedAvg (e.g., $\approx 10\times$ with 10 clients) by only receiving one layer per client. Indicating its viability for deployment-grade secure aggregation for hospital networks.

Introduction

Medical imaging is a leading application area for collaborative model training, yet strict governance and localization constraints often prevent centralized data pooling. Federated learning (FL) offers a principled alternative by enabling multi-institutional training without raw data exchange, and several comprehensive reviews and domain-specific studies have documented its feasibility across classification, detection, and segmentation tasks in clinical contexts (Guan et al. 2024; Sandhu et al. 2023; Rehman et al. 2023). However, cross-silo settings are inherently heterogeneous owing to challenges like differences in scanners, acquisition protocols, populations, and labeling practices, which induce strong non identically distributed (non-IID) effects that slows convergence and reduce out-of-distribution generalization when treated with uniform schedules (Alekseenko, Karargyris, and Padoy 2024; Rehman et al. 2023). These observations motivate designs that adapt to data heterogeneity while remaining mindful of privacy, robustness, and operational constraints typical of hospital networks.

A second, practical challenge is bandwidth. In medical imaging, high-resolution inputs drive large backbones, and communication, not computation, often dominates the training wall clock. This has spurred *partial-model synchronization* strategies that exchange only a subset of parameters each round. Layer-wise adaptive aggregation, like FedLAMA varies synchronization frequency at the layer level to reduce bandwidth with minimal accuracy loss (Lee, Zhang, and Avestimehr 2023). More recent results on partial network updates indicate that restricting each round's exchange to selected layers can accelerate convergence and cut message size, while preserving end accuracy (Wang et al. 2024). A broad survey of aggregation techniques likewise concludes that parameter importance is not uniform and that prioritizing *what* to exchange can outperform monolithic averaging (Qi et al. 2024; Guan et al. 2024). Formally, for a model with P parameters partitioned into L logical groups, exchanging a single group per round reduces uplink from $\Theta(P)$ to $\Theta(P/L)$. This happens when low-influence groups are deferred, and thus, the amortized cost can be even smaller. This turns communication into a first-class optimization variable.

Orthogonal to efficiency, robustness and fairness are central to clinically deployable FL. Byzantine-resilient aggregation methods (for example: coordinate-wise median, trimmed mean, and recent robust M-estimators) provide non-asymptotic guarantees and favorable breakdown behavior under adversarial clients. Although their performance depends on the data regime and on limiting any single participant's per-round influence (Li, Ngai, and Voigt 2024; Bao, Wu, and He 2024). Fairness-aware approaches have concurrently addressed gradient-direction conflicts at the *layer level* to avoid global updates that benefit only a subset of clients. This improves per-client parity while retaining accuracy (Pan et al. 2024). Empirical evidence from medical segmentation and personalization shows that explicitly accounting for cross-site distance and client drift yields better generalization and personalization than uniform schedules (Alekseenko, Karargyris, and Padoy 2024; Xie et al. 2024). Collectively, these results argue that *which layer is updated*

and *who updates it* should be decided jointly, under explicit fairness and robustness constraints.

Security requirements further constrain design choices (Banerjee et al. 2024), (Ahamed et al. 2025a), (Ahamed et al. 2025b). Secure aggregation, for example in SecAgg has become the de-facto protocol to conceal individual client updates while revealing only their sum to the server (Bonawitz et al. 2017). Newer schemes improve verifiability and efficiency at scale, which is especially relevant for imaging backbones and partial updates where large vectors are masked and combined frequently (Behnia et al. 2024). In parallel, end-to-end system studies and surveys in healthcare FL continue to emphasize that robustness, fairness, and verifiability must co-exist with realistic scheduling and networking assumptions to be viable in practice (Qi et al. 2024; Guan et al. 2024), (Ahamed et al. 2025c). Encouragingly, recent multi-institution evaluations comparing FL and centralized training (CL) under matched protocols report that FL can meet or approach CL performance at substantially lower data-sharing risk. In some tasks, FL even match CL with fewer training epochs when using appropriate optimizers and clipping strategies (Kim et al. 2025).

The above discussions highlights the fact that communication reduction, robustness, and fairness, treated in isolation, do not fully address the realities of cross-silo medical imaging. Partial synchronization reduces bandwidth but usually ignores *who* should update *which* parts of the network in a given round. Similarly, robust aggregation curbs outliers but often operates at the full-model granularity, allowing any single client to exert wide influence. Finally, fairness methods steer the global update direction but are rarely coupled to communication budgets or to per-layer importance. To summarize, the challenges that remain in practical deployments are: (i) layers contribute unequally to loss descent, (ii) clients differ markedly in data utility and sample size, and (iii) hospitals require strict privacy with lightweight telemetry. A workable solution must therefore *couple* layer importance with client quality, explicitly bound per-round leverage, and integrate robustness and secure aggregation. It needs to keep the wire cost proportional to the *selected* layers rather than the full model.

Motivation and Novelty

Prior works have addresses issues of partial synchronization, robustness, and fairness in isolation. The schemes leaves a gap where scheduling decisions are decoupled from the current value of layers and the current utility of clients. In practice, layer relevance and client utility drift over rounds. Thus, static or heuristic schedules misallocate bandwidth, starve important layers, and induce oscillations in the global update. A complete end-to-end umbrella that couples marginal utility (per layer) with contributor quality (per client) enforces redundancy and coverage to smooth dynamics.

Motivated by aforementioned discussions, in this paper, we introduce a data-aware scheduler that turns these requirements into a single assignment problem. Per-round influence signals (per layer) and quality signals (per client) feed a capacity-constrained, entropic transport that maximizes expected loss descent per communicated parameter. Then it rounds to a stable one-update-per-client plan. This design explicitly bounds per-round leverage at the layer level and pairs it with redundancy and per-layer robust aggregation executed inside secure aggregation. This tightens breakdown behavior without exposing individual updates.

Contributions and Layout

The contributions of the article are enumerated as follows.

- A per-round scheduling is formulated as a single assignment that couples *layer influence* (server-side, EMA of gradient energy on a tiny root set) with *client data quality* (lightweight accuracy and size metadata). A capacity-constrained, entropic transport with redundancy and sliding-window coverage yields a soft plan that we discretize to one-layer-per-client, aligning bandwidth with expected loss descent.
- Based on the assignment, a train-all, send-one protocol is adopted so each client affects only one layer per round, then perform *per-layer* robust aggregation under secure aggregation.
- We provide a drop-in implementation, communication/compute profile, and ablations, The design is framework-agnostic and compatible with standard secure-aggregation backends used in cross-silo hospitals.

Section 2 formalizes the cross-silo setting and the scheduling objective. Section 3 details the pipeline: the layer-influence estimator, client data-quality scorer, and the entropic-OT matcher with discretization. Section 4 presents the evaluation details with discussions. Finally, section 5 concludes the article with future extension of the work.

System Model and Problem Formulation

The section discusses the system model and problem formulation in the cross-silo setting.

System Model

We are considering a cross-silo federated environment with a server and clients $\mathcal{C} = \{1, \ldots, C\}$ jointly training a deep model with parameters $\theta = \{\theta_\ell\}_{\ell \in \mathcal{L}}$, where $\mathcal{L}$ denotes logical layers/groups (e.g., residual blocks). Time proceeds in rounds $t = 0, 1, \ldots$. Each round, every selected client c initializes from the current global model θ^t, performs local training on its own private data $\mathcal{D}_c$ updating all layers, and uploads only one assigned layer's model delta. Let $L(c,t) \in \mathcal{L}$ denote the assigned layer for client c at round t, and $S_\ell^t = \{c \in \mathcal{C} : L(c,t) = \ell\}$ the set of clients assigned to layer ℓ. The layer-wise client delta and server update are

$$\Delta\theta_\ell^{(c,t)} = \theta_\ell^{(c,t)} - \theta_\ell^t, \tag{1}$$

$$\theta_\ell^{t+1} = \theta_\ell^t + \eta\, \mathsf{Agg}_\ell\Big(\big\{\Delta\theta_\ell^{(c,t)} : c \in S_\ell^t\big\}\Big), \tag{2}$$

where $\eta > 0$ is a server step size and $\mathsf{Agg}_\ell(\cdot)$ is a robust per-layer aggregator (e.g. coordinate-median or trimmed-mean). The untouched layers (with $S_\ell^t = \varnothing$) will be pulled forward, with potentially some momentum decay implemented.

To curb poisoning and decrease variance, we use redundancy $r \geq 2$. Thus, $|S_\ell^t| \geq r$ for critical layers and also

enforce coverage ensuring each ℓ is updated at least once in any block of H rounds. Communication uses secure aggregation (over the transmitted layer vectors) so individual $\Delta\theta_\ell^{(c,t)}$ are masked exposing the server to the aggregate

$$\mathbf{m}_\ell^{(t)} = \sum_{c\in S_\ell^t} \left(\Delta\theta_\ell^{(c,t)} + \text{mask}_{c,\ell}^{(t)}\right), \sum_{c\in S_\ell^t} \text{mask}_{c,\ell}^{(t)} = \mathbf{0}, \quad (3)$$

after which Agg_ℓ in (2) is applied to the recovered (unmasked) updates.

Problem Formulation

The goal is to determine who publishes which layer each round so we can maximize global loss descent while protecting against bad clients. We introduce (i) *layer influence* I_ℓ^t calculated server-side (e.g., EMA of per-layer gradient energy on tiny root data set) and (ii) *client data-quality* Q_c^t (e.g., validation utility with shrinkage/EMA plus a sample-size term). Let $x_{\ell c}^t \in \{0,1\}$ be assignment ($x_{\ell c}^t = 1 \Leftrightarrow L(c,t) = \ell$). We view scheduling as an influence×quality match with redundancy and coverage.

$$\begin{aligned}\max_{\{x_{\ell c}^t\}} \quad & \sum_{\ell\in\mathcal{L}}\sum_{c\in\mathcal{C}} I_\ell^t\, Q_c^t\, x_{\ell c}^t \\ & - \lambda_1 \sum_{\ell\in\mathcal{L}} \Big[r - \textstyle\sum_{c\in\mathcal{C}} x_{\ell c}^t\Big]_+^2 \\ & - \lambda_2 \sum_{\ell\in\mathcal{L}} \left[1 - \sum_{\tau=t-H+1}^{t} \sum_{c\in\mathcal{C}} x_{\ell c}^\tau\right]_+ \qquad (4) \\ \text{s.t.} \quad & \sum_{\ell\in\mathcal{L}} x_{\ell c}^t \le 1, \quad \forall c\in\mathcal{C}, \qquad x_{\ell c}^t \in \{0,1\}. \qquad (5)\end{aligned}$$

where $[\cdot]_+ = \max(0,\cdot)$; the first penalty promotes per-layer redundancy $\ge r$ and the second enforces coverage over a sliding window H. By matching high-influence layers with high-quality clients, the effective gradient signal is concentrated and variance is reduced, thus leading to faster convergence.

Security and robustness are strengthened by limiting leverage in each round to a single layer (leveraging both reweighting and overlapping multiple clients through redundancy via robust M-estimation). Writing $\mathcal{U}_\ell^t = \{\Delta\theta_\ell^{(c,t)} : c \in S_\ell^t\}$ and a robust loss ρ_ℓ, the per-layer aggregate solves which bounds the influence of a Byzantine fraction $0 < \beta_\ell < \frac{1}{2}$ per layer (per layer henceforth referring to as the breakdown point of ρ_ℓ), but preserves the learning progress, guaranteed by the assignment program (4)–(5).

Proposed Approach

In this section, we first present our data-aware layer-wise federated training pipeline. We then present the train-all, send-one with per-layer robust aggregation and redundancy, and finish with communication, compute, and complexity analysis.

Overview

Our data-aware layered federated training consists of three embeddings: (i) *layer-influence estimator*, (ii) *data-quality scorer*, and (iii) *layer–client matcher with robust aggregation*, as shown in Fig. 1. On the server, the layer influence estimator computes the influence of a layer based on the level of influence and a small root set, and maintains exponentially smoothed influence scores. At the client, the data-quality scorer condenses a local validation performance and sample-size into a single scalar without revealing any raw data. Based on these signals and some protocol parameters (e.g., redundancy r and coverage window length H), the layer-client matcher gives a next-round list of (client, layer) pairs, one layer per client, via an influence×quality assignment; clients train all layers, but only upload the delta from their assigned layer using secure aggregation, and the server applies a robust per-layer or client-level reducer to modify the global model.

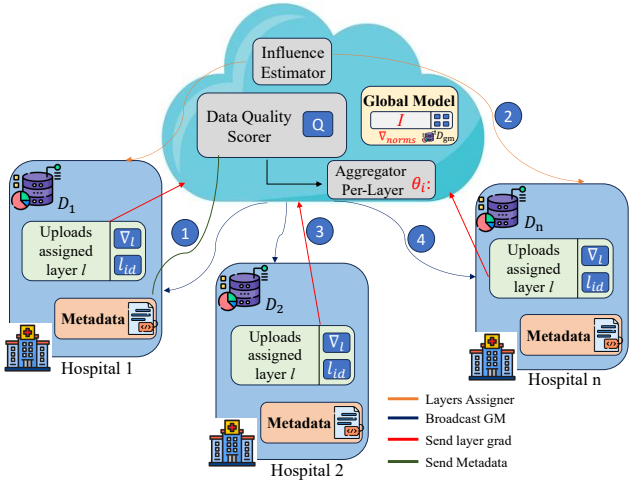

Figure 1: **Schematic illustration:** the server estimates per-layer influence (on a small root set) and client data quality, assigns one layer per client, then aggregates per-layer updates to refresh the global model.

Data Aware Layer Assignment

This module decides *which client uploads the gradient of which layer* in round t so that the server receives exactly one layer update per participating client while prioritizing layers that most influence the loss and clients whose data is most informative.

Server side layer influence: Let $g_\ell(x,y;\theta^t) = \nabla_{\theta_\ell}\mathcal{L}(f_{\theta^t}(x),y)$ be the gradient of the loss w.r.t. layer ℓ at the current global model. Using a tiny server root minibatch $B_t \subset \mathcal{D}_0$, we estimate and smooth the per-layer influence and then normalize across layers as shown in Equation 6.

$$\begin{aligned}\widehat{I}_\ell^t &= \frac{1}{|B_t|}\sum_{(x,y)\in B_t} \big\|g_\ell(x,y;\theta^t)\big\|_2, \\ I_\ell^t &= (1-\rho)\, I_\ell^{t-1} + \rho\, \widehat{I}_\ell^t, \\ I_\ell^t &\leftarrow \frac{I_\ell^t}{\varepsilon + \sum_{k\in\mathcal{L}} I_k^t}. \qquad (6)\end{aligned}$$

Server side client data quality: Each client c returns only lightweight metadata $m_c^t = (a_c^t, n_c, u_c^t)$ consisting of validation accuracy a_c^t on a small challenge set (or local hold-out), local sample count n_c, and an optional utility signal u_c^t (e.g., recent loss drop per unit time). The server forms a stabilized, size-aware score and applies EMA,

$$\bar{a}^{\,t} = \frac{1}{|\mathcal{C}|}\sum_{j\in\mathcal{C}} a_j^t, \qquad \widetilde{a}_c^t = (1-\lambda)\, a_c^t + \lambda\, \bar{a}^{\,t}, \tag{7}$$

$$s_c = \frac{n_c}{\max_j n_j}, \qquad q_c^t = \widetilde{a}_c^t + \alpha\, s_c + \mu\, u_c^t, \tag{8}$$

$$Q_c^t = (1-\sigma)\, Q_c^{t-1} + \sigma\, q_c^t, \qquad Q_c^t \leftarrow \mathrm{clip}\big(Q_c^t, 0, 1\big). \tag{9}$$

We henceforth use I_ℓ^t and Q_c^t as the normalized influence and quality scores inside the assignment.

The server decides *who uploads which layer* at round t using layer influence and client data quality. Given per-layer scores $I_\ell^t > 0$ and per-client scores $Q_c^t > 0$, we form a weight matrix $W^t \in \mathbb{R}_+^{|\mathcal{L}|\times|\mathcal{C}|}$ with entries $W_{\ell c}^t = I_\ell^t\, Q_c^t$ and encode coverage and fairness via time-aware modifiers $u_\ell^t = 1 + \beta \min\{1, \frac{t - t_\ell^{\mathrm{upd}}}{H}\}$ and $v_c^t = 1 - \gamma\, \chi_c^t$, where t_ℓ^{upd} is the last round layer ℓ was updated and χ_c^t counts recent critical-layer assignments of client c. We then compute a *soft assignment* X^t by solving an entropic, capacity-constrained transport that respects per-layer redundancy $a_\ell = r$ and per-client capacity $b_c = 1$:

$$\begin{aligned} X^t = \arg\max_{X\ge 0} \Big\langle \underbrace{\log\big(\epsilon + (u^t \odot I^t)\,(v^t \odot Q^t)^\top\big)}_{\widetilde{W}^t},\, X \Big\rangle \\ - \ \tau \sum_{\ell,c} X_{\ell c}(\log X_{\ell c} - 1) \\ - \ \lambda \big\| X - \overline{X}^{\,t-1} \big\|_F^2 \\ \text{s.t.} \quad X\mathbf{1} = a, \qquad X^\top \mathbf{1} = b. \end{aligned} \tag{10}$$

Here $\epsilon > 0$ stabilizes logs, $\tau > 0$ is the entropic temperature, $\lambda \ge 0$ smooths via $\overline{X}^{\,t-1}$, $a = (r, \dots, r)^\top$, and $b = (1, \dots, 1)^\top$. With kernel $K = \exp(\widetilde{W}^t/\tau)$, the Sinkhorn form is

$$X^t = \mathrm{diag}(p)\, K\, \mathrm{diag}(q), \tag{11}$$
$$p \leftarrow a \oslash (Kq), \quad q \leftarrow b \oslash (K^\top p).$$

We discretize via maximum-weight b matching (or randomized Birkhoff rounding). Let $\mathcal{M}$ be $\{0,1\}$ matrices with $\sum_c x_{\ell c} = r$ and $\sum_\ell x_{\ell c} \le 1$; then

$$X_{\mathrm{disc}}^t \in \arg\max_{X\in\mathcal{M}} \langle \widetilde{W}^t, X\rangle, \qquad \mathbb{E}[X_{\mathrm{disc}}^t] \approx X^t, \tag{12}$$

and $L(c,t) = \{\ell :\ X_{\mathrm{disc},\,\ell c}^t = 1\}$. This couples *what* to update with *who* should update it, while enforcing redundancy and one-layer per client capacity.

Client Update and Per-Layer Aggregation

Given the discrete assignment X_{disc}^t from the previous subsection, each selected client c receives the current global model θ^t and performs local training on its private shard for E epochs (Adam, step size 10^{-4}, cross-entropy), updating *all* layers to produce $\theta^{(c,t)}$. After training, client c evaluates a held-out split to obtain validation accuracy a_c^t and reports metadata (a_c^t, n_c). The server then computes a *size-aware* quality score as a convex blend of accuracy and normalized sample mass,

$$s_c = \frac{n_c}{\max_{j\in\mathcal{C}} n_j}, \tag{13}$$

$$Q_c^t = (1-\omega)\, a_c^t + \omega\, s_c, \quad \omega \in [0,1]. \tag{14}$$

(optionally clipped to $[0,1]$) to stabilize noisy estimates from small shards while rewarding data coverage.

On the server, layers are grouped by parameter-name prefixes (e.g., `conv1`, `bn1`, `layer1.0.`, ...).Per-round layer influence is computed by averaging the ℓ_2 norm of gradients over a few server-side batches:

$$I_\ell^t = \frac{1}{|B_t|}\sum_{(x,y)\in B_t}\sum_{p\in\ell} \big\|\nabla_{\theta_p}\mathcal{L}(f_{\theta^t}(x), y)\big\|_2, \tag{15}$$
$$\text{where } B_t \subset \mathcal{D}_{\mathrm{srv}} \text{ and } |B_t| \le M_{\mathrm{inf}}.$$

Here M_{inf} is a small cap on the number of server batches used for influence estimation (e.g., $M_{\mathrm{inf}} \in [5, 10]$) to control compute and stabilize variance.

Groups are sorted by I_ℓ^t (descending). Iterating in this order, the server selects for each group ℓ the best *available* client by data quality,

$$c^\star(\ell) \ \in \ \arg\max_{c\in\mathrm{Avail}} Q_c^t, \tag{16}$$

assigns that client to ℓ (one layer per client until the pool is exhausted, then the pool is reset), and updates the group's

Algorithm 1 Data Aware Layer Assignment

θ^t, layers $\mathcal{L}$, clients $\mathcal{C}$, root set $\mathcal{D}_0$, metadata $\{m_c^t = (a_c^t, n_c, u_c^t)\}$, last-updated $\{t_\ell^{\mathrm{upd}}\}$, fairness counts $\{\chi_c^t\}$, params $r, H, \rho, \lambda, \tau, \epsilon, \alpha, \mu, \sigma, \beta, \gamma$ X_{disc}^t, $\{S_\ell^t\}$, $L(c,t)$

▷ *Layer influence (EMA, norm on root minibatch)* Sample $B_t \subset \mathcal{D}_0$; $\widehat{I}_\ell^t \leftarrow \frac{1}{|B_t|}\sum_{(x,y)\in B_t} \|\nabla_{\theta_\ell}\mathcal{L}(f_{\theta^t}(x), y)\|_2, \forall \ell$; $I_\ell^t \leftarrow (1-\rho) I_\ell^{t-1} + \rho\, \widehat{I}_\ell^t$; $I_\ell^t \leftarrow I_\ell^t / (\epsilon + \sum_k I_k^t)$;

▷ *Client quality (shrinkage + EMA)* $\bar{a}^{\,t} \leftarrow \frac{1}{|\mathcal{C}|}\sum_j a_j^t$, $s_{\max} \leftarrow \max_j n_j$; $c \in \mathcal{C}$ $\widetilde{a}_c^t \leftarrow (1-\lambda) a_c^t + \lambda \bar{a}^{\,t}$; $s_c \leftarrow n_c / s_{\max}$; $q_c^t \leftarrow \widetilde{a}_c^t + \alpha s_c + \mu u_c^t$; $Q_c^t \leftarrow (1-\sigma) Q_c^{t-1} + \sigma q_c^t$; $Q_c^t \leftarrow \mathrm{clip}(Q_c^t, 0, 1)$;

▷ *Transport weights and capacities* $u_\ell^t \leftarrow 1 + \beta \min\{1, (t - t_\ell^{\mathrm{upd}})/H\}$, $v_c^t \leftarrow 1 - \gamma \chi_c^t$; $\widetilde{W}^t \leftarrow \log\big(\epsilon + (u^t \odot I^t)(v^t \odot Q^t)^\top\big)$; $K \leftarrow \exp(\widetilde{W}^t/\tau)$; $a \leftarrow (r, \dots, r)$, $b \leftarrow (1, \dots, 1)$;

▷ *Sinkhorn scaling (entropic OT)* $p \leftarrow \mathbf{1}$, $q \leftarrow \mathbf{1}$; marginals not met $p \leftarrow a \oslash (Kq)$; $q \leftarrow b \oslash (K^\top p)$; $X^t \leftarrow \mathrm{diag}(p) K\, \mathrm{diag}(q)$;

▷ *Discretization and schedule* $X_{\mathrm{disc}}^t \leftarrow$ MaxWeightBMatching($\widetilde{W}^t$; $a = r, b = 1$) **or** RandomizedBirkhoff(X^t); $S_\ell^t \leftarrow \{c :\ X_{\mathrm{disc},\ell c}^t = 1\}, \forall \ell$; $L(c,t) \leftarrow \{\ell :\ X_{\mathrm{disc},\ell c}^t = 1\}, \forall c$;

X_{disc}^t, $\{S_\ell^t\}$, $L(c,t)$

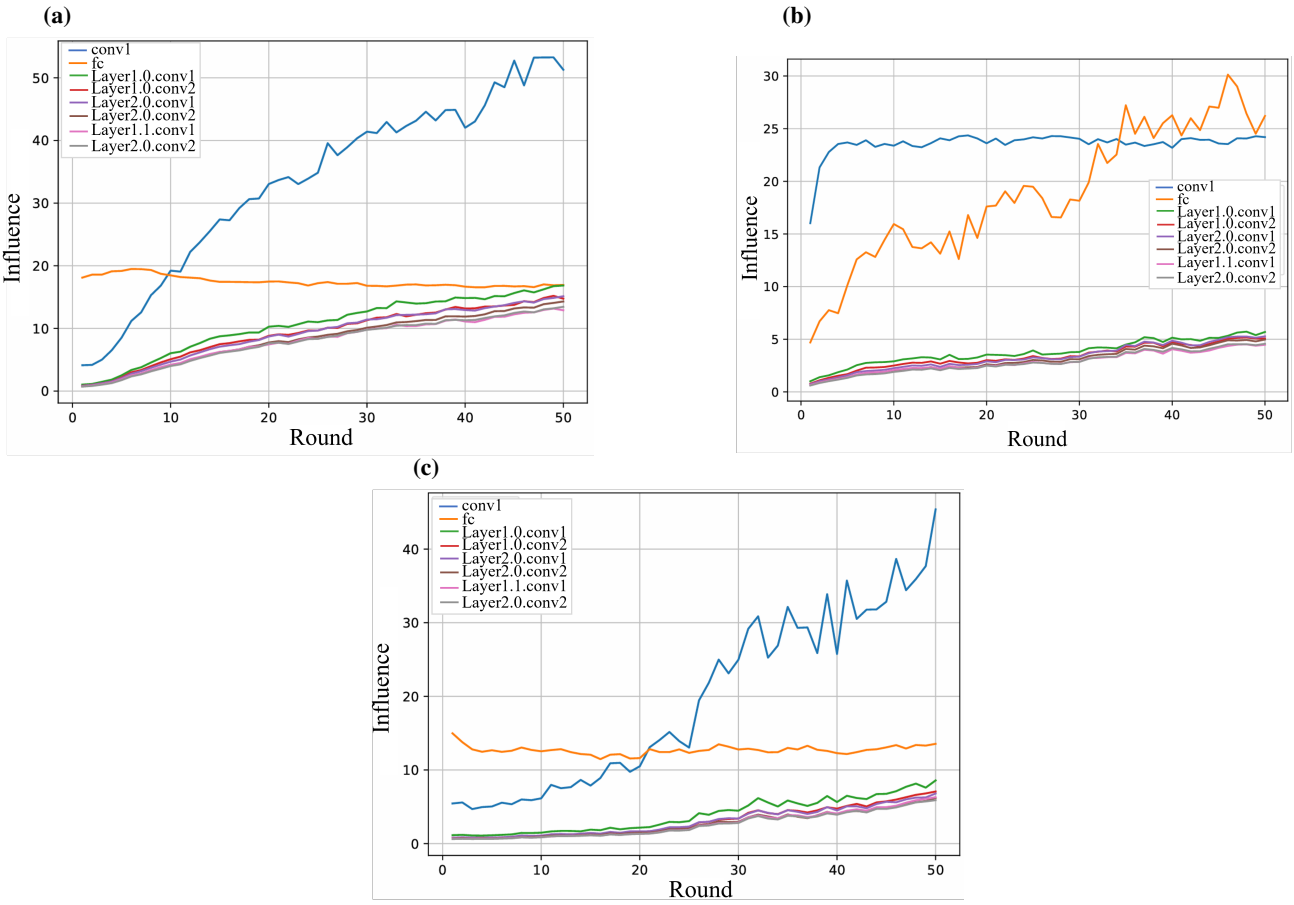

Figure 2: Top-k per-layer influence trajectories (EMA of ℓ_2 gradient norms per logical layer) across federated rounds under our data-aware layer scheduler: **(a)** Brain Tumor MRI, **(b)** Pneumonia CXR, **(c)** Skin Cancer (ISIC). The curves illustrate an early superiority of shallow layers that is then overtaken by mid/deep layers as the representations stabilize; the sporadic spikes illustrate when layers are reassigned critical to quality clients, and the overall decay and reduction in variance indicate convergence.

parameters by a convex blend of global and client weights:

$$\theta_\ell^{t+1} \;=\; (1-\alpha)\,\theta_\ell^{t} \;+\; \alpha\,\theta_\ell^{(c^\star(\ell),t)}, \quad \alpha \in (0,1) \;\; (\texttt{ALPHA}). \tag{17}$$

All other groups use their previous values (carry-forward). This realizes the implemented "train-all, send-one (server-chosen)" protocol with *one client per group* and *blended per-group aggregation* rather than explicit delta averaging or redundancy.

Experimental Evaluation

Experimental Setup

Dataset Description and model We evaluate on three medical image classification benchmarks: (1) *Pediatric Chest X-ray (AP)* from Guangzhou Women and Children's Medical Center (Kermany, Zhang, and Goldbaum 2018); (2) *ISIC Skin Lesion* (benign vs. malignant) with 1,440 benign and 1,197 malignant dermoscopic images (Rotemberg et al. 2021); and (3) *Brain Tumor MRI* with four classes namely glioma (1,321), meningioma (1,339), no-tumor (1,595), and pituitary (1,457) (Nickparvar 2021). All images are resized to 224×224; grayscale modalities are replicated to three channels and normalized with ImageNet statistics. Unless noted otherwise, we federate 10 clients via an uneven, class-skewed partitioner (each client eats between 2–15% of the corpus) and utilize an 80/20 train/validation split locally. The backbone is *ResNet-18* tenured from scratch with cross-entropy (He et al. 2016).

Analysis w.r.t per-layer influence Across all three datasets as shown in Fig. 2, the top-k layer-influence trajectories start off with relatively high influence in the stem and early blocks and incrementally transition mid/deep blocks and the classifier head later in the experiment, as features stabilize. Brain Tumor MRI traverses the influence space sooner, suffering larger amplitude oscillation while Pneumonia CXR travels a more gradual path with smaller dynamic range, both consistent with the potential visual signals. Impulses intermittently indicate rounds when the influence of the influential layers is reassigned to stronger clients by the scheduler. After these rounds, both magnitude

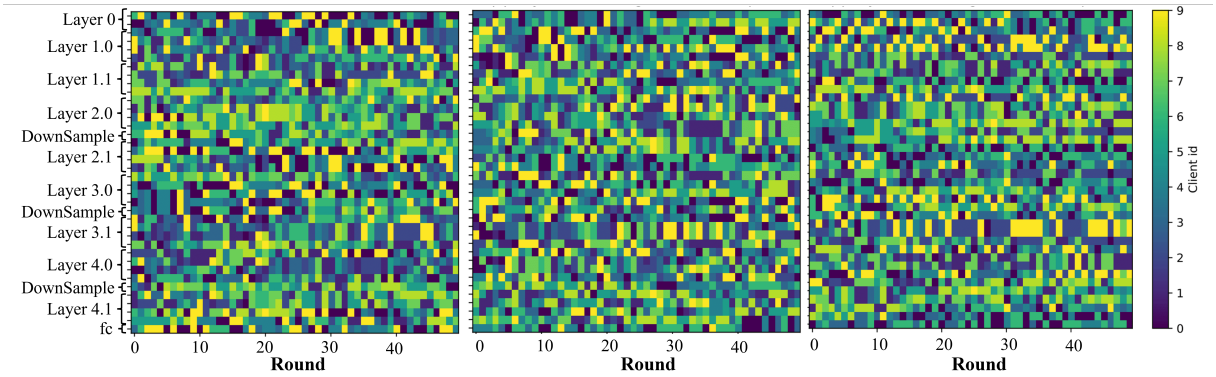

Figure 3: **Layer–client assignment heat maps** for (a) Skin Cancer (ISIC) , (b) Brain Tumor MRI, and (c) Pneumonia CXR, where each pixel displays which client uploaded the gradient of a given layer (conv. and batch normalization) at round t (x-axis), color-coded by client ID. The patterns rotate, rather than lock in, indicating selection driven by influence with fairness, when critical layers repeatedly align to high-quality clients, while updates are not being monopolized by a single client.

and cross-layer variance dissipate, implying convergence in later rounds. These trajectories afford evidence to our design goal: to learn which layers matter when, and assign to whom can best improve them.

Threshold-free detection quality We report AUROC (area under the ROC, i.e. TPR–FPR trade-off across thresholds) and AUPRC (area under the Precision–Recall curve, more informative under class imbalance). Pneumonia CXR achieves AUROC = 0.9249, AUPRC = 0.9347; Brain Tumor MRI is near-perfect with AUROC = 0.9955, AUPRC = 0.9881; Skin Cancer (ISIC) is moderate at AUROC = 0.8344, AUPRC = 0.8515. The tighter AUROC/AUPRC pairing on the CXR/BRAIN MRI suggests strong ranking and precision under skew, while dermatology remains now significantly harder (fine-grained classes, variable imaging) and opportunities this continues to leave open for attribution to additional richer augmentations or domain priors. Overall, the figures imply strong discrimination without threshold tuning across tasks.

Communication comparison with FedAvg Let $\|\theta\|$ denote the model size (in bytes) and n be the number of clients. In FedAvg, the per-round *uplink* is

$$C^{\uparrow}_{\text{FedAvg}} = n\,\|\theta\|. \tag{18}$$

In our scheme, each client uploads exactly one disjoint layer so the server receives about one model in total:

$$C^{\uparrow}_{\text{ours}} \approx \|\theta\|. \tag{19}$$

Hence, $\frac{C^{\uparrow}_{\text{ours}}}{C^{\uparrow}_{\text{FedAvg}}} \approx \frac{1}{n}$, i.e., a relative reduction of $1 - \frac{C^{\uparrow}_{\text{ours}}}{C^{\uparrow}_{\text{FedAvg}}} \approx \frac{n-1}{n}$ in per-round uplink compared to FedAvg.

Assignment Entropy & Fairness We see a significant amount of assignment entropy in Fig. 3, that is, in most layers the participants assigned for layers shift across rounds frequently, which indicates that it was predominantly dynamic scheduled, not a static partition. For the early rounds we see short dwell times for shallow layers, while medium and deeper layers show short 2-3 round streaks after that, indicating temporal coherence in that the influence/staleness is at its peak. The color pattern across rows/columns are typically evenly allocated, indicating relatively uniform levels of participation and abiding by considered fairness constraints. The short vertical bands indicate that participants are never committing to many rounds in sequence, and we have a reasonably reliable defense against a potential "lock-in" of clients with minimal threats of single poisoning.

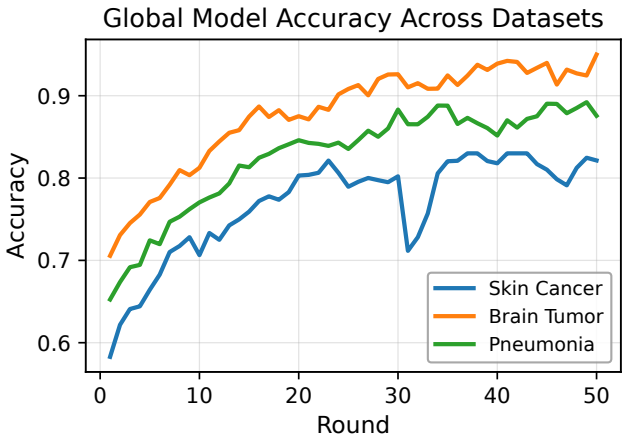

Figure 4: Global model accuracy vs. federated rounds on three datasets: Skin Cancer (ISIC), Brain Tumor MRI, and Pneumonia (Chest X-ray).

Conclusion

We presented a data-aware, layer-wise federated learning framework in this paper specifically designed for cross-silo, medical imaging. Our scheduler combines per-layer influence (expected mean absolute value of energy during training on a tiny server root set) with per-client data quality and solves a capacity-constrained continuous entropic assignment problem to decide every round who will upload which layer; clients can train-all, then send-one, and the

server facilitates a robust per-layer aggregation under secure aggregation with bounded leverage per round. Across Pneumonia CXR, Brain Tumor MRI, and ISIC Skin Cancer with AUROC = 0.925, 0.9955, 0.8344 respectively, the method demonstrated effectiveness, all while forming a decrease, and taking per round uplink a capability-over-time, and increasing parity and stability of assignments. These results encompass alignment of extrapolated expected loss descent and contribute to deployment-friendly efficiency, robustness, and fairness for hospital networks. Future research will assess budgeted considerations for multi-layer assignments and personalization heads, along with the application of formal convergence and robustness guarantees under secure aggregation.

Appendix

The complete source code for our proposed data-aware layer assignment framework, along with all experimental configurations, dataset preprocessing scripts, and hyperparameter settings, has been made publicly available. The repository includes the necessary instructions to replicate the results presented in this paper for all three medical imaging benchmarks. The implementation can be accessed at our GitHub repository: https://github.com/gsai21/Data-Aware-Layer-Assignment-for-Federated-Learning-for-Medical-Image-Analysis.

References

Ahamed, S. F.; Banerjee, S.; Roy, S.; Kapoor, A.; Vucovich, M.; Choi, K.; Rahman, A.; Bowen, E.; and Shetty, S. 2025a. Privacy Drift: Evolving Privacy Concerns in Incremental Learning. In *2025 International Conference on Computing, Networking and Communications (ICNC)*, 510–515. IEEE.

Ahamed, S. F.; Banerjee, S.; Roy, S.; Vucovich, M.; Quinn, D.; Choi, K.; Rahman, A.; Hu, A.; Bowen, E.; and Shetty, S. 2025b. Accuracy-privacy trade-off in the mitigation of membership inference attack in federated learning. In *2025 International Conference on Computing, Networking and Communications (ICNC)*, 243–248. IEEE.

Ahamed, S. F.; Roy, S.; Banerjee, S.; Vucovich, M.; Choi, K.; Rahman, A.; Hu, A.; Bowen, E.; and Shetty, S. 2025c. Evaluating Query Efficiency and Accuracy of Transfer Learning-based Model Extraction Attack in Federated Learning. In *2025 International Wireless Communications and Mobile Computing (IWCMC)*, 643–648. IEEE.

Alekseenko, J.; Karargyris, A.; and Padoy, N. 2024. Distance-Aware Non-IID Federated Learning for Generalization and Personalization in Medical Imaging Segmentation. In Burgos, N.; Petitjean, C.; Vakalopoulou, M.; Christodoulidis, S.; Coupe, P.; Delingette, H.; Lartizien, C.; and Mateus, D., eds., *Proceedings of The 7nd International Conference on Medical Imaging with Deep Learning*, volume 250 of *Proceedings of Machine Learning Research*, 33–47. PMLR.

Banerjee, S.; Roy, S.; Ahamed, S. F.; Quinn, D.; Vucovich, M.; Nandakumar, D.; Choi, K.; Rahman, A.; Bowen, E.; and Shetty, S. 2024. MIA-BAD: An Approach for Enhancing Membership Inference Attack and its Mitigation with Federated Learning. In *2024 International Conference on Computing, Networking and Communications (ICNC)*, 635–640. IEEE Computer Society.

Bao, W.; Wu, J.; and He, J. 2024. BOBA: Byzantine-Robust Federated Learning with Label Skewness. In Dasgupta, S.; Mandt, S.; and Li, Y., eds., *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, 892–900. PMLR.

Behnia, R.; Riasi, A.; Ebrahimi, R.; Chow, S. S. M.; Padmanabhan, B.; and Hoang, T. 2024. Efficient Secure Aggregation for Privacy-Preserving Federated Machine Learning. In *2024 Annual Computer Security Applications Conference (ACSAC)*, 778–793.

Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H. B.; Patel, S.; Ramage, D.; Segal, A.; and Seth, K. 2017. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, 1175–1191. New York, NY, USA: Association for Computing Machinery. ISBN 9781450349468.

Guan, H.; Yap, P.-T.; Bozoki, A.; and Liu, M. 2024. Federated learning for medical image analysis: A survey. *Pattern Recognition*, 151: 110424.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.

Kermany, D.; Zhang, K.; and Goldbaum, M. 2018. Labeled Optical Coherence Tomography (OCT) and Chest X-Ray Images for Classification. Mendeley Data. Pediatric chest X-ray subset from Guangzhou Women and Children's Medical Center.

Kim, J.; Lee, H.; Park, J.; Park, S. H.; Lee, M.; Sunwoo, L.; Kim, C. K.; Kim, B. J.; Kim, D.-E.; and Ryu, W.-S. 2025. In-silo federated learning vs. centralized learning for segmenting acute and chronic ischemic brain lesions. *Intelligence-Based Medicine*, 12: 100283.

Lee, S.; Zhang, T.; and Avestimehr, A. S. 2023. Layer-Wise Adaptive Model Aggregation for Scalable Federated Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(7): 8491–8499.

Li, S.; Ngai, E. C.-H.; and Voigt, T. 2024. An Experimental Study of Byzantine-Robust Aggregation Schemes in Federated Learning. *IEEE Transactions on Big Data*, 10(6): 975–988.

Nickparvar, M. 2021. Brain Tumor MRI Dataset. Kaggle.

Pan, Z.; Li, C.; Yu, F.; Wang, S.; Wang, H.; Tang, X.; and Zhao, J. 2024. FedLF: Layer-Wise Fair Federated Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(13): 14527–14535.

Qi, P.; Chiaro, D.; Guzzo, A.; Ianni, M.; Fortino, G.; and Piccialli, F. 2024. Model aggregation techniques in federated learning: A comprehensive survey. *Future Generation Computer Systems*, 150: 272–293.

Rehman, M. H. u.; Hugo Lopez Pinaya, W.; Nachev, P.; Teo, J. T.; Ourselin, S.; and Cardoso, M. J. 2023. Federated learning for medical imaging radiology. *British Journal of Radiology*, 96(1150): 20220890.

Rotemberg, V.; Kurtansky, N.; Betz-Stablein, B.; and et al. 2021. A patient-centric dataset of images and metadata for skin lesions. *Scientific Data*, 8(1): 82.

Sandhu, S. S.; Gorji, H. T.; Tavakolian, P.; Tavakolian, K.; and Akhbardeh, A. 2023. Medical Imaging Applications of Federated Learning. *Diagnostics*, 13(19).

Wang, H.; Liu, X.; Niu, J.; Guo, W.; and Tang, S. 2024. Why Go Full? Elevating Federated Learning Through Partial Network Updates. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 99773–99799. Curran Associates, Inc.

Xie, L.; Lin, M.; Liu, S.; Xu, C.; Luan, T.; Li, C.; Fang, Y.; Shen, Q.; and Wu, Z. 2024. pFLFE: Cross-silo Personalized Federated Learning via Feature Enhancement on Medical Image Segmentation. In Linguraru, M. G.; Dou, Q.; Feragen, A.; Giannarou, S.; Glocker, B.; Lekadir, K.; and Schnabel, J. A., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, 599–610. Cham: Springer Nature Switzerland. ISBN 978-3-031-72117-5.