

PERSONALIZED FEDERATED LEARNING WITH RESIDUAL FISHER INFORMATION FOR MEDICAL IMAGE SEGMENTATION

Meilu Zhu, Yuxing Li, Zhiwei Wang, Edmund Y. Lam

Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong, China

Federated learning enables multiple clients (institutions) to collaboratively train machine learning models without sharing their private data. To address the challenge of data heterogeneity across clients, personalized federated learning (pFL) aims to learn customized models for each client. In this work, we propose pFL-ResFIM, a novel pFL framework that achieves client-adaptive personalization at the parameter level. Specifically, we introduce a new metric, Residual Fisher Information Matrix (ResFIM), to quantify the sensitivity of model parameters to domain discrepancies. To estimate ResFIM for each client model under privacy constraints, we employ a spectral transfer strategy that generates simulated data reflecting the domain styles of different clients. Based on the estimated ResFIM, we partition model parameters into domain-sensitive and domain-invariant parts. A personalized model for each client is then constructed by aggregating only the domain-invariant parameters on the server. Extensive experiments on public datasets demonstrate that pFL-ResFIM consistently outperforms state-of-the-art methods.

1. INTRODUCTION

Federated Learning (FL) allows multiple clients (hospitals) to collaboratively learn a shared model while keeping data localized, making it valuable for privacy-sensitive medical scenarios [1]. However, the performance of the standard FL framework is often compromised by statistical heterogeneity across clients. When each client learns from distinct data distributions, their local updates may conflict, causing the aggregated global model to perform poorly [2, 3, 4, 5].

Personalized federated learning (pFL) has emerged as an alternative framework that enables selective knowledge transfer across the federation [6]. pFL produces separate models tailored to each client to adapt to local data, rather than a shared global model. One representative type of pFL algorithms is decoupling parameters of a model into the domain-invariant and domain-sensitive subnetworks and keeping the latter locally to avoid inter-client interference. For example, FedBN [7] and SioBN [8] consider batch normalization layers to be domain-sensitive and keep them updated locally without

aggregating them at the server. In addition, FedPer[9] only aggregates feature extractors because they are less affected by domain discrepancy. However, these methods achieve limited performance due to two issues. (1) Layer-level selection strategy is empirically-driven and cannot balance inter-client knowledge sharing and local personalization. (2) As the model training procedure is dynamic, the client-specific subnetwork should vary across optimization iterations.

To address these problems, we introduce pFL-ResFIM, a new pFL framework based on Residual Fisher Information Matrix (ResFIM), which aims to achieve client-adaptive and atom-level personalization. Specifically, we first introduce a ResFIM metric to measure the sensitivity of model parameters to domain discrepancy at the parameter level. To estimate the ResFIM of each client model in FL setting without compromising data privacy, we propose to leverage a spectral transfer based strategy to generate simulated data with domain styles of different clients. Based on the estimated ResFIM, the parameters of each client model can be separated into domain-sensitive and domain-invariant parts. A personalized model for each client is then constructed by aggregating only the domain-invariant parameters on the server. The contributions of this work are summarized as follows: (1) We propose pFL-ResFIM, a new personalized FL framework that uses the Fisher Information Matrix to achieve client-adaptive and atom-level personalization. (2) We introduce a ResFIM metric that effectively measures the sensitivity of client model parameters to domain discrepancies among clients.

2. METHODOLOGY

In this section, we present our pFL-ResFIM, which aims to achieve personalized federated learning guided by Residual Fisher Information Matrix (ResFIM). The core idea of pFL-ResFIM is to use ResFIM to identify domain-sensitive parameters of client models. We propose to preserve these parameters at the client side for model personalization and only aggregate the remaining for knowledge sharing across clients.

2.1. Residual Fisher Information Matrix

Fisher Information Matrix (FIM) is a cornerstone of statistical estimation theory, quantifying the amount of information

Corresponding author: Edmund Y. Lam (elam@eee.hku.hk)

This work is supported by the Innovation and Technology Fund (ITS/341/23) of Hong Kong, China.




Fig. 1. The architecture of the pFL-ResFIM framework.

an observable random variable carries about an unknown parameter. In deep learning, FIM is adopted to capture the sensitivity of the model's output distribution to its parameters, which asymptotically approximates the second-order curvature of the loss landscape around a given point. Formally, for a model with parameters W , FIM is defined as:

$\mathbf{F}_W = \mathbb{E}_x [\mathbb{E}_{y \sim f_W(y|x)} \nabla_W \mathcal{L}(f_W(x), y) \cdot \nabla_W \mathcal{L}(f_W(x), y)^\top]$, where x is a sample with the label y . $\mathbf{F}_W \in \mathbb{R}^{|W| \times |W|}$ is the symmetrical matrix and can be efficiently approximated via the diagonal approximation to reduce the computational overhead [10]. $\nabla_W \mathcal{L}(\cdot)$ represents the gradient of the loss function with respect to the parameters W . Intuitively, $\mathbf{F}_W$ captures how much changing the parameters affects the model's output, thereby measuring the importance of each parameter.

However, FIM can not sufficiently capture the behavior of parameters in cross-domain scenes, particularly in FL, where parameters can exhibit significantly different sensitivities to varying data distributions. Consequently, a more nuanced metric is required to quantify the variation of parameters across different domains. To this end, we propose to calculate the Residual Fisher Information Matrix (ResFIM), denoted as $\Delta \mathbf{F}_W$, to directly measure parameter sensitivity to domain discrepancies. It is defined as:

$$\Delta \mathbf{F}_W = \frac{|\mathbf{F}_W(x, y) - \mathbf{F}_W(\bar{x}, y)|}{\max(|\mathbf{F}_W(x, y) - \mathbf{F}_W(\bar{x}, y)|, \epsilon)}, \quad (1)$$

where ϵ is a small constant to prevent division by zero. x and $\bar{x}$ are samples from different domains. $|\mathbf{F}_W(x, y) - \mathbf{F}_W(\bar{x}, y)|$ computes the difference between the FIM across domains, reflecting the model's varying sensitivity to parameter changes in different data distributions. A higher $\Delta \mathbf{F}_{W(p)}$ indicates that the parameter $W(p)$ is more sensitive to domain changes.

2.2. Personalized Federated Learning with ResFIM

Next, we describe how to integrate ResFIM into a FL system to achieve client-adaptive and atom-level personalization.

Overview of pFL-ResFIM We consider a pFL system comprising a central server and C clients, collaboratively training a set of personalized models $\{W_c\}_{c=1}^C$. Each client c possesses a private dataset $\mathcal{D}_c = \{(x_i^c, y_i^c)\}_{i=1}^{N_c}$ with N_c samples. The objective the system is formulated as:

$$\min_{W_1, \dots, W_C} f(W_1, \dots, W_C) = \frac{1}{C} \sum_{c=1}^C \mathbb{E}_{(x_i^c, y_i^c) \sim \mathcal{D}_c} \mathcal{L}(W_c, x_i^c).$$

The overall optimization process proceeds through communication between clients and the server for multiple rounds, as shown in Fig. 1. In each round, every client c first performs local training on W_c using its private data $\mathcal{D}_c$. Next, the client estimates the ResFIM $\Delta \mathbf{F}_{W_c}$ for its local model and generates a binary mask $\mathbf{M}_c$ by selecting the top $\delta\%$ of parameters with the largest values in $\Delta \mathbf{F}_{W_c}$. The local models $\{W_c\}_{c=1}^C$ and their corresponding masks $\{\mathbf{M}_c\}_{c=1}^C$ are then transmitted to the server. The server proceeds to construct a personalized model $\bar{W}_c$ for each client via the following aggregation:

$$\bar{W}_c = \mathbf{M}_c * W_c + (1 - \mathbf{M}_c) \sum_{i=1}^C W_i. \quad (2)$$

Here, the first term on the right side preserves the client's domain-specific knowledge by retaining its sensitive parameters locally. The second term achieves knowledge sharing across the federation by aggregating the domain-invariant parameters of all clients.

Estimation of ResFIM in FL As shown in Eq. (1), computing the ResFIM $\Delta \mathbf{F}_{W_c}$ for a client c requires data from all other clients $\{\mathcal{D}_j\}_{j=1}^{C \setminus c}$. Direct data sharing is infeasible due to strict privacy constraints. To simulate the data of other clients with privacy leakage, we leverage a spectral transfer based method inspired by [12, 13, 14]. Specifically, for an image $x_i^c \in \mathbb{R}^{H \times W \times D}$ of the client c , we first apply the Fast Fourier Transform (FFT) to project it into the frequency domain, obtaining the signal $\mathcal{F}(x_i^c)$: $\mathcal{F}(x_i^c)(u, v, d) = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} x_i^c(h, w, d) e^{-j2\pi(\frac{h}{H}u + \frac{w}{W}v)}$. $\mathcal{F}(x_i^c)$ can be further decomposed into an amplitude spectrum $\mathcal{A}_i^c$ and a phase spectrum $\mathcal{P}_i^c$, where the low-frequency components of $\mathcal{A}_i^c$ primarily encapsulate the domain style. Before federated training begins, all clients upload their amplitude spectrum sets $\{\mathcal{A}_i^c\}_{i=1}^{N_c} \in \mathbb{R}^{N \times H \times W \times D}$ to the server in a one-time initialization step, where N is the total sample size of clients and $\mathcal{A}^c = \{\mathcal{A}_i^c\}_{i=1}^{N_c}$. The server then distributes the set $\{\mathcal{A}^j\}_{j=1}^{C \setminus c}$ to each client c . During local training, for each image x_i^c , the client c randomly samples an amplitude spectrum from each $\mathcal{A}^j$ in $\{\mathcal{A}^j\}_{j=1}^{C \setminus c}$. It then replaces the low-frequency component of its own $\mathcal{A}_i^c$ with the low-frequency component from the sampled spectra. These modified amplitude spectra are combined with the original phase spectrum $\mathcal{P}_i^c$ and transformed back via inverse FFT, generating a set

Table 1. The Dice performance of different methods on the vessel dataset.

Methods	ChaseDB	DR-Hagis	DRIVE	HRF	LES-AV	ORVS	Average
FedAvg [11]	68.24 $\pm$ 0.74	69.80 $\pm$ 0.10	76.06 $\pm$ 0.07	72.51 $\pm$ 0.09	73.30 $\pm$ 0.06	71.56 $\pm$ 0.02	71.91 $\pm$ 2.49
FedBN [7]	69.67 $\pm$ 0.83	70.80 $\pm$ 0.02	76.16 $\pm$ 0.08	71.35 $\pm$ 0.17	72.83 $\pm$ 0.38	71.60 $\pm$ 0.07	72.07 $\pm$ 2.05
SioBN [8]	68.01 $\pm$ 0.21	70.34 $\pm$ 0.15	75.55 $\pm$ 0.10	70.89 $\pm$ 0.16	73.49 $\pm$ 0.69	71.64 $\pm$ 0.09	71.65 $\pm$ 2.38
FedPer [9]	70.80 $\pm$ 0.17	69.45 $\pm$ 0.18	75.90 $\pm$ 0.02	73.00 $\pm$ 0.01	73.61 $\pm$ 0.26	71.14 $\pm$ 0.16	72.32 $\pm$ 2.11
pFL-ResFIM	72.74 $\pm$ 0.25	71.63 $\pm$ 0.17	76.20 $\pm$ 0.13	73.31 $\pm$ 0.06	74.58 $\pm$ 0.28	71.50 $\pm$ 0.13	73.33 $\pm$ 1.65

Table 2. Comparison of various methods on gland dataset.

Methods	LFSC	MGD-1K	Average
FedAvg [11]	74.52 $\pm$ 1.86	81.86 $\pm$ 0.96	78.19 $\pm$ 3.67
FedBN [7]	74.18 $\pm$ 3.26	81.86 $\pm$ 1.68	78.02 $\pm$ 3.99
SioBN [8]	63.42 $\pm$ 7.04	81.49 $\pm$ 0.76	72.45 $\pm$ 9.03
FedPer [9]	72.67 $\pm$ 3.08	81.72 $\pm$ 2.00	77.20 $\pm$ 4.52
pFL-ResFIM	76.58 $\pm$ 1.75	82.32 $\pm$ 1.01	79.45 $\pm$ 2.86

of simulated images $\{\bar{x}_{i,k}^c\}_{k=1}^{C-1}$. This process creates a simulated dataset $\bar{\mathcal{D}}_c = \{(\bar{x}_i^c, y_i^c)\}_{i=1}^{N_c \times (C-1)}$ that preserves the semantic content of $\mathcal{D}_c$ while exhibiting the style variations of other clients. The ResFIM $\Delta \mathbf{F}_{W_c}$ can then be reliably estimated using the local data $\mathcal{D}_c$ and the simulated data $\bar{\mathcal{D}}_c$.

3. EXPERIMENTS


3.1. Datasets and Implementation Details

Datasets. (1) VESSEL SEGMENTATION. Retina images are collected from ChaseDB [15], DR-Hagis [16], DRIVE [17], HRF [18], LES-AV [19], ORVS [20], forming six clients. Based on original partition, their training and sets contain 20 and 8, 34 and 6, 20 and 20, 15 and 30, 17 and 5, 41 and 8 images, respectively. (2) MEIBOMIAN GLAND SEGMENTATION. The infrared images of this dataset are from two clients. Their data are collected from the Lo Fong Shiu Po Eye Centre (LFSC) in Hong Kong and a public dataset MGD-1K [21], respectively.

Implementation. The proposed pFL-ResFIM and all baseline methods are implemented with PyTorch library. We adopt UNet as the backbone network of all methods. For all datasets, we use the Adam optimizer with a batch size of 4. The default numbers of local epochs and communication rounds are 4 and 100. The vessel and gland datasets use a initial learning rate of 1×10^{-3} with a decaying rate of 0.99 per round. Model performance is evaluated using the Dice similarity coefficient. For each dataset, we conduct three trials and present the mean and the standard deviation.

3.2. Comparison with State-of-the-art Methods

We compare pFL-ResFIM with existing methods on three datasets, including FedAvg [11], FedBN [7], SioBN [8], and FedPer [9]. As shown in Tables 1–2, baseline meth-

**Fig. 2.** Masking rates of different layers of two clients.

ods [7, 8, 9] that designate certain layers as domain-sensitive often exhibit inconsistent performance across clients. For example, in Table 1, SioBN yields top performance on ORVS yet yields the lowest result on ChaseDB. These observations indicate that personalizing same layers fails to generalize across all clients. In contrast, pFL-ResFIM consistently achieves the best performance on nearly every client across all three datasets, demonstrating the effectiveness of its client-adaptive personalization strategy.

3.3. Ablation Study

To evaluate the effectiveness of ResFIM on client-adaptive personalization, Fig. 2 plots the masking rates of different layers of two client models during training on the vessel dataset. It can be observed that the input and classifier layers of two clients exhibit significantly different masking rates at the same epoch, whereas their intermediate layers (enc4.0) maintain relatively lower masking rates. This observation aligns with the current empirical consensus that both shallow and prediction layers are highly sensitive to domain shift. These results confirm the necessity of client-adaptive personalization and the effectiveness of ResFIM.

4. CONCLUSION

In this work, we introduce a new pFL framework called pFL-ResFIM that uses a Residual Fisher Information Matrix (ResFIM) based metric to measure the sensitivity of model parameters to domain shift. Based on the estimated ResFIM, we separate the parameters of a client model into domain-sensitive and domain-invariant parts and only aggregating the latter at the server for model personalization. Extensive experiments on public datasets demonstrate that pFL-ResFIM consistently outperforms state-of-the-art methods, validating its effectiveness for client-adaptive personalization.

5. COMPLIANCE WITH ETHICAL STANDARDS

The use of open-access datasets complied with their licensing terms and did not require separate ethical approval. As for the in-house data, its use was approved by the Ethics Committee of the data provider's institution.

6. REFERENCES

- [1] Meilu Zhu, Axiu Mao, Jun Liu, and Yixuan Yuan, "Deer: Deviation eliminating and noise regulating for privacy-preserving federated low-rank adaptation," *IEEE Trans. Med. Imaging*, 2024.
- [2] Meilu Zhu, Qiushi Yang, Zhifan Gao, Yixuan Yuan, and Jun Liu, "Fedbm: Stealing knowledge from pre-trained language models for heterogeneous federated learning," *Medical Image Analysis*, vol. 102, pp. 103524, 2025.
- [3] Axiu Mao, Endai Huang, Haiming Gan, and Kai Liu, "Fedaar: A novel federated learning framework for animal activity recognition with wearable sensors," *Animals*, vol. 12, no. 16, pp. 2142, 2022.
- [4] Meilu Zhu, Qiushi Yang, Zhifan Gao, Jun Liu, and Yixuan Yuan, "Stealing knowledge from pre-trained language models for federated classifier debiasing," in *MICCAI*. Springer, 2024, pp. 685–695.
- [5] Meilu Zhu, Jing Liao, Jun Liu, and Yixuan Yuan, "Fedoss: Federated open set recognition via inter-client discrepancy and collaboration," *IEEE Trans. Med. Imaging*, vol. 43, no. 1, pp. 190–202, 2023.
- [6] Zhen Chen, Meilu Zhu, Chen Yang, and Yixuan Yuan, "Personalized retrogress-resilient framework for real-world medical federated learning," in *MICCAI*. Springer, 2021, pp. 347–356.
- [7] Xiaoxiao Li, Meirui Jiang, et al., "Fedbn: Federated learning on non-iid features via local batch normalization," in *ICLR*, 2021.
- [8] Mathieu Andreux, Jean Ogier du Terrail, et al., "Siloed federated learning for multi-centric histopathology datasets," in *MICCAI Workshop on Distributed and Collaborative Learning*, pp. 129–139. Springer, 2020.
- [9] Manoj Ghuhari Arivazhagan, Vinay Aggarwal, et al., "Federated learning with personalization layers," *arXiv preprint arXiv:1912.00818*, 2019.
- [10] Michael S Matena and Colin A Raffel, "Merging models with fisher-weighted averaging," *NeurIPS*, vol. 35, pp. 17703–17716, 2022.
- [11] Brendan McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in *AISTATS*, 2017, pp. 1273–1282.
- [12] Yanchao Yang and Stefano Soatto, "Fda: Fourier domain adaptation for semantic segmentation," in *CVPR*, 2020, pp. 4085–4095.
- [13] Quande Liu, Cheng Chen, et al., "Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space," in *CVPR*, 2021, pp. 1013–1023.
- [14] Zhiwei Wang, Defeng He, Li Zhao, Bo Liu, Yayu Zheng, and Xiaoqin Zhang, "Difusionseg: Diffusion-driven semantic segmentation with multi-modal image fusion for enhanced perception," *Knowledge-Based Systems*, p. 114481, 2025.
- [15] Muhammad Moazam Fraz, Paolo Remagnino, et al., "An ensemble classification-based approach applied to retinal blood vessel segmentation," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 9, pp. 2538–2548, 2012.
- [16] Sven Holm, Greg Russell, Vincent Nourrit, and Niall McLoughlin, "Dr hags—a fundus image database for the automatic extraction of retinal surface vessels from diabetic patients," *Journal of Medical Imaging*, vol. 4, no. 1, pp. 014503–014503, 2017.
- [17] Joes Staal, Michael D Abramoff, et al., "Ridge-based vessel segmentation in color images of the retina," *IEEE Trans. Med. Imaging*, vol. 23, no. 4, pp. 501–509, 2004.
- [18] Attila Budai, Rüdiger Bock, Andreas Maier, et al., "Robust vessel segmentation in fundus images," *Int. J. Biomed. Imaging*, vol. 2013, 2013.
- [19] José Ignacio Orlando et al., "Towards a glaucoma risk index based on simulated hemodynamics from fundus images," in *MICCAI*, 2018, pp. 65–73.
- [20] Abdullah Sarhan, Jon Rokne, et al., "Transfer learning through weighted loss function and group normalization for vessel segmentation from retinal images," in *ICPR*, 2021, pp. 9211–9218.
- [21] Ripon Kumar Saha, AM Mahmud Chowdhury, et al., "Automated quantification of meibomian gland dropout in infrared meibography using deep learning," *The Ocular Surface*, vol. 26, pp. 283–294, 2022.