

# Divide, Conquer and Unite: Hierarchical Style-Recalibrated Prototype Alignment for Federated Medical Segmentation

Xingyue Zhao<sup>1\*</sup>, Wenke Huang<sup>1\*</sup>, Xingguang Wang<sup>2\*</sup>, Haoyu Zhao<sup>1</sup>, Linghao Zhuang<sup>3</sup>, Anwen Jiang<sup>4</sup>, Guancheng Wan<sup>1</sup>, Mang Ye<sup>1†</sup>

<sup>1</sup> School of Computer Science, Wuhan University, China

<sup>2</sup> School of Software Engineering, Xi'an Jiaotong University, China

<sup>3</sup> School of Software Engineering, Xinjiang University, China

<sup>4</sup> College of Mathematics and System Science, Xinjiang University, China

## Abstract

Federated learning enables multiple medical institutions to train a global model without sharing data, yet feature heterogeneity from diverse scanners or protocols remains a major challenge. Many existing works attempt to address this issue by leveraging model representations (e.g., mean feature vectors) to correct local training; however, they often face two key limitations: 1) **Incomplete Contextual Representation Learning**: Current approaches primarily focus on final-layer features, overlooking critical multi-level cues and thus diluting essential context for accurate segmentation. 2) **Layerwise Style Bias Accumulation**: Although utilizing representations can partially align global features, these methods neglect domain-specific biases within intermediate layers, allowing style discrepancies to build up and reduce model robustness. To address these challenges, we propose FedBCS to bridge feature representation gaps via domain-invariant contextual prototypes alignment. Specifically, we introduce a frequency-domain adaptive style recalibration into prototype construction that not only decouples content-style representations but also learns optimal style parameters, enabling more robust domain-invariant prototypes. Furthermore, we design a context-aware dual-level prototype alignment method that extracts domain-invariant prototypes from different layers of both encoder and decoder and fuses them with contextual information for finer-grained representation alignment. Extensive experiments on two public datasets demonstrate that our method exhibits remarkable performance.

**Code** — <https://github.com/zxy1234321/FedBCS>

## 1 Introduction

Medical image segmentation is crucial for computer-assisted diagnosis, and most deep learning-based techniques require large amounts of data for effective training (Wang et al. 2024b; Zhao et al. 2024d,c). However, sharing such data among hospitals is often restricted by privacy concerns and regulatory constraints. Federated learning has emerged as a promising strategy for multi-site collaborative training of a

\*These authors contributed equally.

†Corresponding author: yemang@whu.edu.cn

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

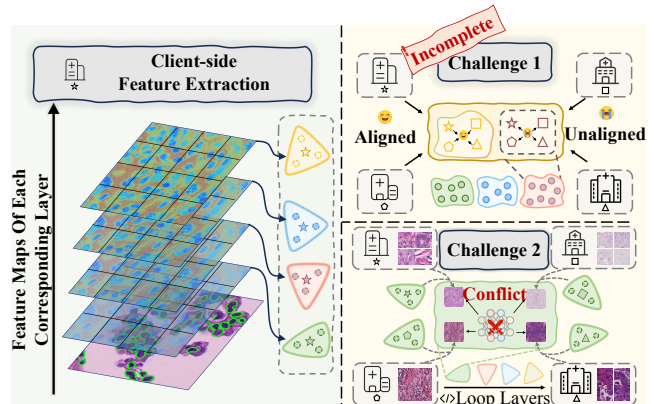


Figure 1: Problem illustration of existing methods under domain skew setting. The top row highlights two major challenges—**incomplete contextual representation learning** and **layerwise style bias accumulation**.

global model without compromising patient privacy. A foundational method in this domain, FedAvg (McMahan et al. 2017), aggregates model parameters from different clients to form a global model, which is then redistributed to each client for further training. Despite these advances, data heterogeneity remains a significant challenge in real-world scenarios (Huang et al. 2024; Zhao et al. 2024a). Datasets from different medical sites are often non-IID due to variations in scanners and imaging protocols, leading to different feature distributions across clients (Jiang, Wang, and Dou 2022; Wang et al. 2025; Zhao et al. 2024b; Zhuang et al. 2025). This diversity can hinder model convergence, reduce generalization capabilities, and result in suboptimal performance.

To address data heterogeneity, researchers use model representations to correct local training and improve generalization, employing mean feature vectors for privacy (Li, He, and Song 2021; Li et al. 2020). Subsequent work then introduces prototypes (class-specific mean feature vectors) and federated prototype learning, where client prototypes minimize distances to global or same-class prototypes while maximizing separation from different-class ones (Tan et al. 2022; Huang et al. 2023b). Many researchers have also in-

roduced specialized strategies to further improve training in these settings. For example, one method clusters client features to generate multiple prototypes per domain instead of a single average, and refines feature-level distance by elevating it to an  $\alpha$  power, addressing unequal learning due to domain diversity (Wang et al. 2024a). Another study leverages Gaussian Mixture Models to generate prototypes and soft predictions, then weights these outputs based on both quality and quantity (Zhang, Shan, and Han 2024).

In federated learning across multiple medical institutions, variations in scanning protocols, staining methods, and patient populations often amplify feature distribution differences (Jiang et al. 2023; Chen et al. 2024; Li et al. 2025). Despite recent advances, many existing approaches still depend on a single-layer or simplified representation, making it challenging to fully capture the multi-scale contextual complexities and style variations in medical image segmentation, where distribution shifts can be more pronounced. We group these limitations into two key challenges and highlight them in Figure 1:

#### **Layerwise Style Bias Accumulation:**

**Style discrepancies from different medical protocols accumulate through intermediate layers, degrading cross-domain feature alignment.** Many existing methods only normalize style at input level or align final-layer features (Huang et al. 2023b), leaving intermediate-layer style biases unaddressed and hindering model generalization.

#### **Incomplete Contextual Representation Learning:**

**Single-layer feature alignment fails to capture the complete semantic hierarchy needed for medical image segmentation.** While lower layers encode crucial local details and higher layers capture abstract structures (Li et al. 2023), relying solely on final-layer representations leads to incomplete semantic understanding and compromised performance (Sung et al. 2024).

Considering these issues, we propose a novel method called FedBCS, specifically designed for feature heterogeneous federated medical image segmentation. For **Layerwise Style Bias Accumulation** challenge, we propose a Frequency-domain Style Recalibration (FSR) module, specifically designed for constructing domain-invariant prototypes in federated learning. We explore methods for decoupling and recalibrating domain-specific style variations during federation by learning adaptive style parameters from cross-client feature representations. Given that style variations propagate through network layers and affect prototype quality differently, the FSR module modulates feature styles in the frequency domain for robust prototype construction. First, we decompose intermediate features into content and style components via Fourier Transform, where phase spectra preserve semantic content while amplitude spectra encode style information (Lee, Bae, and Kim 2023; Chen et al. 2021). Building on this, we introduce learnable style parameters to adaptively recalibrate the normalized amplitude

during federation, enabling optimal style modulation while maintaining semantic consistency. This strategy effectively eliminates domain-specific style biases while ensuring robust prototype construction across different domains.

**Once robust domain-invariant prototypes are constructed, effectively leveraging multi-level semantic information becomes crucial for accurate feature alignment.** Building upon these domain-invariant prototypes, we further address the challenge of **Incomplete Contextual Representation Learning**. Current methods typically construct prototypes from single-layer features, leading to incomplete semantic representations. To tackle this limitation, we propose a context-aware dual-level prototype alignment (CDPA) strategy that leverages both encoder and decoder features. Specifically, we extract domain-invariant prototypes from multiple network layers and design a hierarchical fusion mechanism to integrate multi-scale contextual information. This approach enables more comprehensive semantic alignment by capturing both low-level details and high-level semantic patterns, ultimately achieving finer-grained feature representation. Our main contributions are summarized below.

- **Revealing Federated Learning Limitations in Heterogeneous FL.** We identify two critical limitations in Heterogeneous federated medical segmentation: layerwise style bias accumulation and incomplete contextual representation learning, which significantly impact cross-client feature alignment.
- **Novel Federated Prototype Learning Framework.** We propose FedBCS that integrates frequency-domain style recalibration for robust prototype construction with context-aware dual-level alignment to achieve comprehensive semantic representation learning.
- **Theoretical and Empirical Validation.** We provide theoretical convergence guarantees and demonstrate superior empirical performance through extensive experiments on two medical segmentation tasks.

## **2 Related Work**

### **2.1 Data Heterogeneous Federated Learning**

Federated learning (FL) provides a privacy-preserving framework for collaborative model training without data centralization (Konečný et al. 2016). The foundational method FedAvg (McMahan et al. 2017) iteratively builds a global model by aggregating locally trained parameters. While effective in homogeneous settings, FedAvg often struggles with data heterogeneity, which frequently arises in medical contexts where varying imaging protocols, staining methods, or patient populations lead to significant divergence across institutions (Wang, Jin, and Wang 2022). Based on FedAvg, numerous methods have attempted to address the data heterogeneity problem from different perspectives (Huang et al. 2024). Some approaches focus on training a single global model, aiming to refine the global parameters for improved cross-domain performance (Huang et al. 2023b; Huang, Ye, and Du 2022; Ma et al. 2025), while others emphasize fairness, ensuring balanced outcomes across

heterogeneous clients. Among these, individual/group fairness (Ezzeldin et al. 2023), collaboration fairness (Jiang et al. 2023), and performance fairness (Chen, Huang, and Ye 2024; Li et al. 2021a) are three widely studied notions. Meanwhile, personalization methods (Li et al. 2021b; Huang et al. 2023a) learn individualized models tailored to each client’s domain, and domain generalization approaches (Liu et al. 2021; Zhang et al. 2023) seek robust performance on unseen distributions. In this paper, we address the challenge of data heterogeneity across different medical institutions and concentrate on training a single global model that can robustly perform across these diverse domains.

## 2.2 Model Representation Alignment in Federated Learning

In Federated Learning (FL), many methods rely on model representations to enhance local training in heterogeneous data settings (Wan, Huang, and Ye 2024; Liao et al. 2024; Ye et al. 2023). However, sharing raw features can pose privacy concerns. To address this, mean feature vectors are used to retain essential information while minimizing data leakage. Mean feature vectors computed from same-class samples are called prototypes. MOON (Li, He, and Song 2021) and FedProx (Li et al. 2020) utilize contrastive learning to maximize the agreement between local and global model mean feature vectors, thereby aligning local representations with the global model. FedProto (Tan et al. 2022) enhances the alignment of client-side training by minimizing the distance between client-side class prototypes and global class prototypes. However, by emphasizing single-domain performance under label skew, these approaches often neglect domain shift, thereby limiting their generalization across diverse domains. Thus, many methods adopt prototype-based approaches to address domain skew. For instance, FPL (Huang et al. 2023b) employs server-side, cluster-based prototypes combined with contrastive learning to align same-class representations. However, it relies on a single representation for each client, leading to limited representational capacity and causing unequal learning in heterogeneous domains. By contrast, FedContrast-GPA (Zhou and Zheng 2023) and FedPLVM (Wang et al. 2024a) cluster features on the client side to produce multiple prototypes per class, mitigating the drawbacks of single-representation strategies. FedGMKD (Zhang, Shan, and Han 2024), on the other hand, generates prototypes using Gaussian Mixture Models. Despite their effectiveness, these methods rely solely on final-layer encoder features and overlook the influence of domain-specific style differences on feature representations. In this work, we focus on finer-grained feature alignment and explicitly handling style-induced variations to enhance federated segmentation across diverse medical domains.

## 3 Methodology

### 3.1 Preliminaries

Following typical Federated Learning setup (McMahan et al. 2017; Li et al. 2020), we assume there are  $M$  clients. Each client  $m$  holds a private dataset  $\{(x_i, y_i)\}_{i=1}^{N_m}$ , where  $N_m$

denotes the size of client  $m$ ’s local dataset. The objective of federated learning is to minimize the weighted average of local objectives:

$$\min_{\omega} F(\omega) = \sum_{m=1}^M w_m \left( \frac{1}{N_m} \sum_{i=1}^{N_m} L(x_i, y_i; \omega) \right), \quad (1)$$

where  $w_m$  is the weight of client  $m$ , and  $L(x_i, y_i; \omega)$  is the loss for data point  $(x_i, y_i)$  under model parameters  $\omega$ . Each client updates its model locally and sends the updates to a central server for aggregation.

**Data Heterogeneity.** In federated medical image segmentation, while the label distribution  $P(y)$  remains consistent across clients, the feature distribution  $P(x | y)$  varies due to different scanning devices and protocols, leading to significant feature heterogeneity across medical institutions.

We adopt a standard segmentation network architecture shared across all clients, consisting of an encoder  $F_{\text{enc}} : X \rightarrow Z$  that maps an input image  $x$  to a  $d$ -dimensional feature map  $\mathbf{z}_{\text{enc}} = F_{\text{enc}}(x) \in \mathbb{R}^d$ , and a decoder  $F_{\text{dec}} : Z \rightarrow O$  that produces a segmentation mask  $\mathbf{o}_{\text{dec}} = F_{\text{dec}}(\mathbf{z}_{\text{enc}}) \in \mathbb{R}^c$ , where  $c$  denotes the number of classes.

### 3.2 Frequency-domain Style Recalibration for Prototype Construction

**Federated prototype learning.** Recent federated learning approaches leverage class prototypes alongside network weights for cross-client knowledge sharing. A class prototype represents the mean feature representation of that class on each client:

$$p_m^c = \frac{1}{|N_m^c|} \sum_{x_i \in N_m^c} F_{\text{enc}}^m(x_i), \quad (2)$$

where  $F_{\text{enc}}^m(\cdot)$  represents client  $m$ ’s local encoder (architecturally identical across clients but with parameters synchronized only during FedAvg rounds), and  $N_m^c$  denotes the set of pixels/patches labeled as class  $c$  on that client. The server aggregates these local prototypes to guide subsequent local training phases, enabling efficient cross-client semantic alignment while maintaining privacy and minimizing communication costs.

**Motivation.** Existing prototype-based methods assume features from the same class follow similar distributions across clients, suggesting simple averaging should suffice for prototype construction. However, in medical imaging, diverse scanning protocols introduce significant style variations that propagate through network layers, manifesting in the amplitude spectra of feature representations. When constructing prototypes through direct feature averaging, this leads to:

*Layerwise Style Bias Accumulation:* Style discrepancies compound through intermediate layers, causing features of the same anatomical structure to appear increasingly different across institutions and undermining the effectiveness of prototype learning.

To establish reliable semantic anchors for cross-client feature alignment, prototypes must be constructed from style-invariant features. A promising direction is to leverage frequency domain analysis, where content and style information separate into phase and amplitude components, enabling effective style-invariant prototype construction.

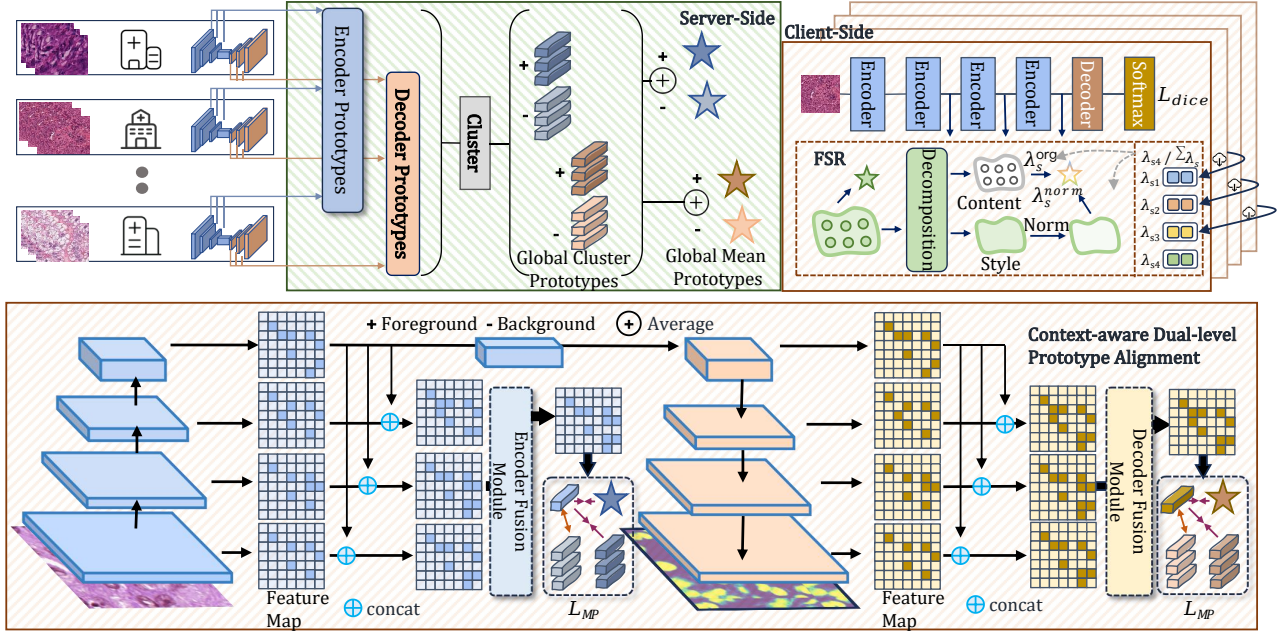


Figure 2: Architecture illustration of FedBCS. To address layerwise style bias accumulation, we implement frequency-domain style recalibration during local prototype construction. To capture complete contextual representations, we extract and align multi-level prototypes from both encoder and decoder pathways. These prototypes are aggregated at the server to derive global prototypes for guiding local training. Zoom in for details.

**Style-invariant Prototype Generation.** To address the aforementioned challenges in prototype-based federated learning for medical image segmentation, we propose Frequency-domain Style Recalibration that removes domain-specific style variations while preserving semantic information during prototype construction. Our key insight is that *style and content information naturally separate in the frequency domain, allowing for targeted style modulation before prototype generation*. Given the encoder feature map  $\mathbf{z}_{\text{enc}} \in \mathbb{R}^{C \times H \times W}$ , conventional prototype-based methods construct class-specific prototypes through direct averaging:

$$p_m^c = \frac{1}{|N_m^c|} \sum_{x_i \in N_m^c} F_{\text{enc}}(x_i), \quad (3)$$

where  $N_m^c$  denotes samples of class  $c$  in client  $m$ 's local dataset, and  $F_{\text{enc}}(\cdot)$  represents the encoder network that extracts feature representations. However, these prototypes directly inherit domain-specific style variations from different medical institutions, compromising their effectiveness as semantic anchors for cross-client feature alignment. To address this limitation, we introduce FSR that processes features in the frequency domain before prototype construction. Specifically, given a feature map  $\mathbf{z}_{\text{enc}}$ , we first decompose it into amplitude and phase components via 2D Fourier transform:

$$Z(\alpha, \beta) = \frac{1}{HW} \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} z(h, w) \exp\left(-i2\pi\left(\frac{\alpha h}{H} + \frac{\beta w}{W}\right)\right), \quad (4)$$

where the amplitude spectrum  $\chi$  and phase spectrum  $\gamma$  are

computed as:

$$\begin{aligned} \chi(\alpha, \beta) &= \sqrt{\mathcal{R}(\alpha, \beta)^2 + \mathcal{I}(\alpha, \beta)^2}, \\ \gamma(\alpha, \beta) &= \arctan\left(\frac{\mathcal{I}(\alpha, \beta)}{\mathcal{R}(\alpha, \beta)}\right). \end{aligned} \quad (5)$$

The amplitude spectrum encodes style information while the phase spectrum preserves semantic content. To achieve adaptive style control, we introduce learnable parameters that balance normalized and original style information:

$$\lambda_s^{\text{norm}}, \lambda_s^{\text{org}} = \sigma(W_s[\text{GAP}(\chi_{\text{norm}}); \text{GAP}(\chi)]), \quad (6)$$

where GAP denotes global average pooling,  $W_s$  are learnable weights, and  $\chi_{\text{norm}}$  is the instance-normalized amplitude. The style-recalibrated features are then obtained through inverse Fourier transform:

$$\hat{z}_{\text{enc}} = \text{IFT}(\lambda_s^{\text{norm}} \chi_{\text{norm}} + \lambda_s^{\text{org}} \chi, \gamma). \quad (7)$$

By applying FSR to the encoder features before averaging, we construct style-invariant prototypes that better preserve class-specific semantic information:

$$\hat{p}_m^c = \frac{1}{|N_m^c|} \sum_{x_i \in N_m^c} \hat{z}_{\text{enc}}, \quad (8)$$

where  $\hat{z}_{\text{enc}}$  is the style-recalibrated feature obtained through our FSR module. Compared to conventional prototypes that directly inherit domain-specific styles, our FSR-based prototypes better capture the underlying semantic representations while being more robust to cross-domain variations.

### 3.3 Context-aware Dual-Level Prototype Alignment

**Motivation.** While our FSR module effectively addresses style variations, existing prototype-based methods (Tan et al.

2022; Huang et al. 2023b) typically only utilize final-layer features for prototype construction and alignment. This single-layer approach overlooks the rich hierarchical information distributed across network layers, where both local details and global structures are crucial for medical image segmentation. When constructing prototypes for cross-client alignment, this leads to:

*Incomplete Contextual Representation Learning:* Single-layer feature alignment fails to capture the complete semantic hierarchy needed for medical image segmentation, where both fine-grained tissue details and global anatomical structures must be precisely delineated.

This motivates our development of a context-aware prototype alignment strategy that specifically addresses the multi-scale nature of medical image segmentation.

**Multi-level Feature Integration.** Building upon our FSR module, which generates style-invariant features  $\hat{z}_{enc}$  through frequency-domain recalibration, we now extend prototype construction to multiple network layers. Given the encoder-decoder architecture with  $K$  layers, for layer  $k \in \{1, \dots, K\}$ , we compute class-specific prototypes from the style-recalibrated features:

$$P_{enc}^{k,c} = \frac{\sum_{v \in \hat{z}_{enc}^k} v \mathbb{I}[g(v) = c]}{|N^c|}, P_{dec}^{k,c} = \frac{\sum_{v \in \hat{z}_{dec}^k} v \mathbb{I}[g(v) = c]}{|N^c|}, \quad (9)$$

where  $g(v)$  maps the feature vector to its ground-truth class label, and  $\mathbb{I}[\cdot]$  is the indicator function. The resulting layer-wise prototypes are denoted as  $P_{enc}^k = \{P_{enc}^{k,1}, \dots, P_{enc}^{k,c}\}$  and  $P_{dec}^k = \{P_{dec}^{k,1}, \dots, P_{dec}^{k,c}\}$ .

In medical image segmentation, different network layers capture complementary anatomical information. Shallow layers encode local tissue textures and edge patterns crucial for precise boundary delineation, while deeper layers represent high-level anatomical structures and spatial relationships between organs. To effectively combine these multi-scale features, we design a hierarchical fusion strategy. Specifically, we concatenate shallow-layer prototypes with deeper ones:

$$\hat{P}_{enc}^k = \text{Concat}(P_{enc}^k, P_{enc}^{k+1}), \quad (10)$$

This design allows high-level semantic guidance to enhance local feature discrimination - for instance, knowing the overall organ shape (deep features) helps better determine tissue boundaries (shallow features) in ambiguous regions.

To mitigate communication overhead while preserving multi-scale information, we introduce a lightweight fusion module  $F : \hat{P} \rightarrow Q$  that adaptively integrates features across layers:

$$Q_{enc} = F(\hat{P}_{enc}), Q_{dec} = F(\hat{P}_{dec}), \quad (11)$$

The fusion module learns to emphasize the most discriminative features at each scale through 1x1 convolutions, resulting in compact yet comprehensive prototypes for each client  $m$ :

$$Q_{enc}^m = \{q_{enc}^{m,1}, \dots, q_{enc}^{m,c}\}, Q_{dec}^m = \{q_{dec}^{m,1}, \dots, q_{dec}^{m,c}\}. \quad (12)$$

**Server-side Prototype Alignment.** In federated medical image segmentation, different institutions may focus on different anatomical regions or pathological conditions, leading to varied feature distributions. To achieve consistent cross-

client alignment while preserving privacy, we introduce a clustering-based alignment strategy. Using FINCH (Sarfranz, Sharma, and Stiefelhagen 2019), we group similar prototypes from different clients:

$$\text{Cluster} \left( \{q_{enc}^{m,c}\}_{m=1}^M \right) \rightarrow \{q_{enc}^{m,c}\}_{m=1}^S \in \mathbb{R}^{S \times d} \quad (13)$$

$$\text{Cluster} \left( \{q_{dec}^{m,c}\}_{m=1}^M \right) \rightarrow \{q_{dec}^{m,c}\}_{m=1}^S \in \mathbb{R}^{S \times d}.$$

To stabilize optimization, we further compute mean prototypes for each class:

$$\bar{q}_{enc}^c = \frac{1}{S} \sum_{q^c \in Q_{enc}^c} q^c, \bar{q}_{dec}^c = \frac{1}{S} \sum_{q^c \in Q_{dec}^c} q^c, \quad (14)$$

where  $S$  denotes the number of clusters.

**Learning Objectives.** For effective feature alignment, we design a dual-stream learning objective. Given a medical image sample  $(x_i, y_i)$ , we extract feature embeddings  $e^i = \{e_{enc}^i, e_{dec}^i\}$ . The contrastive term encourages alignment with same-class prototypes while maintaining separation from different anatomical regions:

$$L_{contra} = -\log \frac{\sum_{q \in Q^c} e^{\text{Sim}(e^i, q)}}{\sum_{q \in Q^c} e^{\text{Sim}(e^i, q)} + \sum_{q \in Q^{c'}} e^{\text{Sim}(e^i, q)}}, \quad (15)$$

where  $Q^c$  contains class-specific prototypes and  $Q^{c'}$  contains prototypes from other classes.

To ensure consistent feature learning across institutions, we introduce a mean prototype consistency regularization:

$$L_{consis} = \sum_{v=1}^d \left( \left( e_{enc}^{i,v} - \bar{q}_{enc}^{c,v} \right)^2 + \left( e_{dec}^{i,v} - \bar{q}_{dec}^{c,v} \right)^2 \right). \quad (16)$$

The final loss combines our multi-level prototype loss ( $L_{MP} = L_{contra} + L_{consis}$ ) with the standard Dice loss:

$$L_{total} = L_{MP} + L_{dice}. \quad (17)$$

### 3.4 Convergence Analysis

We establish the theoretical foundations of our method under the following standard assumptions.

**Assumption 1 (Lipschitz Smoothness of Loss Functions):** The loss function  $\mathcal{F}(\Theta_e)$  for any client  $e$  satisfies  $L_{sm}$ -smoothness, implying its gradient exhibits Lipschitz continuity with a positive constant  $L_{sm}$ . Mathematically, the following inequality holds:

$$F(\Theta) = \frac{1}{N} \sum_{(x,y)} \left[ L_{dice}(x, y; \Theta) + L_{contra}(x, y; \Theta) + L_{consis}(x, y; \Theta) \right]. \quad (18)$$

**Assumption 2 (Unbiased Stochastic Gradients with Bounded Dispersion):** Each client  $e$  computes stochastic gradients that are unbiased estimators of the global gradient, with their variance constrained by  $\sigma^2$ :

$$\mathbb{E}[g_e] = \nabla \mathcal{F}, \mathbb{E}\|g_e - \nabla \mathcal{F}\|^2 \leq \sigma^2. \quad (19)$$

**Assumption 3 (Bounded Prototype Norm):** The uploaded encoder/decoder prototypes  $p_{dec}^k$  from each client satisfy  $\|p_{dec}^k\| \leq G$ , where  $G > 0$  is a predefined constant ensuring stable optimization.

**Assumption 4 (Lipschitz Continuity of InfoNCE Loss):** For unit-norm anchor vectors, the InfoNCE loss  $L_{contra}$  with respect to the positive sample  $r$  satisfies the Lipschitz gradi-

ent condition:

$$\left| \frac{\partial L_{\text{contra}}}{\partial r} \right| \leq \frac{1}{\tau}. \quad (20)$$

**Theorem 1 (One-Round Convergence Bound).** Under Assumptions 1-4, for any client participating in the Federated Learning process, the expected objective function after each communication round satisfies:

$$\begin{aligned} \mathbb{E}[\mathcal{F}_{t+1}] &\leq \mathcal{F}_t - \alpha(\eta) \sum_{e=0}^{E-1} \|\nabla \mathcal{F}(\Theta_e)\|^2 \\ &\quad + \frac{L_{sm}\eta^2 E}{2} \sigma^2 + \lambda_c \frac{E\eta G}{\tau}. \end{aligned} \quad (21)$$

This result establishes an upper bound on the expected objective function after each communication round, with convergence guaranteed under proper choice of  $\eta$  and  $\lambda_c$ :

$$\eta_{e'} < \frac{2(\sum_{e=0}^{e'} \|\nabla \mathcal{F}(\Theta_e)\|^2 - \frac{\lambda_c EG}{\tau})}{L_{sm}(\sum_{e=0}^{e'} \|\nabla \mathcal{F}(\Theta_e)\|^2 + E\sigma^2)}, \quad (22)$$

where  $e' = 1/2, 1, \dots, E-1$ , and

$$\lambda_t < \frac{\tau \|\nabla \mathcal{F}(\Theta_e)\|^2}{EG}. \quad (23)$$

**Theorem 2 (Non-convex convergence rate).** Under Assumptions 1-4, and  $\Delta = \mathcal{F}_0 - \mathcal{F}^*$  where  $\mathcal{F}^*$  refers to the local optimum. For an arbitrary client, given any  $\epsilon > 0$ , after

$$T = \frac{2\Delta}{E\epsilon(2\eta - L_{sm}\eta^2) - E\eta(L_{sm}\eta\sigma^2 + \frac{2\lambda_c G}{\tau})} \quad (24)$$

communication rounds, we have

$$\frac{1}{TE} \sum_{t=0}^{T-1} \sum_{e=0}^{E-1} \mathbb{E}[\|\nabla \mathcal{F}(\Theta_e)\|^2] < \epsilon, \quad (25)$$

if

$$\eta < \frac{2(\epsilon - \frac{\lambda_c G}{\tau})}{L_{sm}(\epsilon + \sigma^2)} \quad \text{and} \quad \lambda_c < \frac{\tau\epsilon}{G}. \quad (26)$$

The result shows that with proper hyperparameters ( $\eta, \lambda_c$ ), the expected norm of gradients can be bounded within any given precision level  $\epsilon > 0$ . **Please refer to the appendix for detailed evidence and analysis.**

## 4 Experiments

### 4.1 Experiment Settings

**Datasets.** We evaluate methods on two segmentation tasks.

- **Histology nuclei segmentation:** We assembled four public datasets (TCIA (Hou et al. 2020), CRC (Awan et al. 2017), KIRC (Irshad et al. 2014), TNBC (Naylor et al. 2018)) for binary segmentation of cell nuclei in histology images.
- **Prostate MRI segmentation:** We used a multi-site dataset from (Liu et al. 2020), which comprises six sources from three public datasets (Lemaître et al. 2015; Litjens et al. 2014; Nicholas et al. 2015) for binary segmentation of the prostate.

We treat each data source as a separate client. Due to variations in pathological staining and MRI scanning protocols, images from each client differ visually, yet they maintain a consistent label distribution. The example cases in each domain are presented in Fig. 1. Following (Wang et al. 2023; Liu et al. 2020), we adopt the same preprocessing




Figure 3: Performance analysis. (a) Convergence comparison of different methods across communication rounds. (b) Effectiveness of FSR in addressing layerwise style bias compared to input-level normalization. See details in section 4.2.

methods and data splits, and directly evaluate the results on the test sets. For more dataset details, please refer to Appendix. For both segmentation tasks, we use commonly used UNet (Ronneberger, Fischer, and Brox 2015).

**Comparison Methods.** We compare our method against several SOTA federated learning approaches designed to construct a shared global model, including FedAvg (McMahan et al. 2017), FedProx (Li et al. 2020), MOON (Li, He, and Song 2021), FedProto (Tan et al. 2022), HarmoFL (Jiang, Wang, and Dou 2022), FPL (Huang et al. 2023b), FedContrast (Zhou and Zheng 2023), Fed-PLVM (Wang et al. 2024a) and FedUV (Son et al. 2024).

**Implementation Details.** All approaches are implemented under identical configurations, using 400 communication rounds and 1 local epoch. For histology nuclei segmentation, we adopt the SGD optimizer with a learning rate of 0.01, while for prostate MRI segmentation we use the Adam optimizer with a learning rate of 1e-4. The optimizer parameters are set with a weight decay of 1e-4, and the training batch size is 6.

**Evaluation Metric.** Following (Wang et al. 2023), we adopt the widely used Dice score as the evaluation metric for both segmentation tasks. Our evaluation focuses solely on the foreground Dice score. To ensure reproducibility, we fix the random seed for all experiments.

### 4.2 Diagnostic Analysis

**Hyper-parameter Study.** Figure 4 illustrates the impact of hyperparameter  $\tau$  (eq. (20)) on performance. For histology nuclei segmentation, optimal performance is achieved at  $\tau=0.005$ , while for prostate MRI segmentation the best results occur at  $\tau=0.4$ . The results, depicted in Fig. 4, demonstrate that our method consistently surpasses the second-place method across a range of temperatures (refer to Tab. 1).

**Ablation Study.** To comprehensively analyze the effectiveness of the two components, we conducted an ablation study on both datasets, as shown in Table 2. The quantitative findings confirm that each component individually enhances performance, and their combined application leads to the optimal outcomes.

**Communication Cost Analysis.** We compare our method’s communication overhead with other federated prototype approaches in Tab. 3 (prostate MRI segmentation) and Fig. 4

Methods	Histology nuclei segmentation						Prostate MRI segmentation							
	TCIA	CRC	KIRC	TNBC	AVG	$\Delta$	BIDMC	HK	HCRUDB	RUNMC	BMC	UCL	AVG	$\Delta$
FedAvg [ASTAT17]	72.92	73.74	73.60	57.72	69.50	-	68.81	79.86	86.85	80.39	82.23	74.90	78.80	-
FedProx [arXiv20]	72.53	74.00	74.07	55.60	69.00	-0.50	68.44	82.06	86.01	80.38	81.53	70.44	78.10	-0.70
MOON [CVPR21]	72.66	73.81	73.53	57.25	69.30	-0.20	70.76	84.39	85.89	81.32	82.86	71.51	79.50	+0.70
FedProto [AAAI22]	71.66	71.79	73.31	56.22	68.20	-1.30	67.06	79.30	85.06	78.99	80.97	71.22	77.10	-1.70
HarmoFL [AAAI22]	76.77	76.02	74.72	60.09	<u>71.90</u>	<u>+2.40</u>	79.66	80.81	87.88	78.34	83.51	76.70	<u>81.20</u>	<u>+2.40</u>
FPL [CVPR23]	74.02	75.12	73.36	64.16	71.66	+2.16	67.01	83.96	83.08	80.51	80.82	69.14	77.40	-1.40
FedContrast [MICCAI23]	75.98	73.69	72.86	59.38	70.50	+1.00	65.43	83.87	83.73	80.20	79.74	66.29	76.50	-2.30
FedPLVM [NeurIPS24]	73.46	75.58	72.23	62.60	71.00	+1.50	67.41	80.70	84.73	79.59	81.28	71.17	77.50	-1.30
FedUV [CVPR24]	72.71	74.74	73.01	59.48	70.00	+0.50	72.40	82.12	86.04	78.56	82.23	69.59	78.50	-0.30
<b>FedBCS</b>	<b>76.21</b>	<b>78.70</b>	<b>74.42</b>	<b>67.00</b>	<b>74.10</b>	<b>+4.60</b>	<b>77.56</b>	<b>84.56</b>	<b>87.22</b>	<b>83.31</b>	<b>85.01</b>	<b>77.93</b>	<b>82.60</b>	<b>+3.80</b>

Table 1: Comparison with the SOTA methods on Histology nuclei segmentation and Prostate MRI segmentation tasks. AVG denotes average accuracy calculated on all domains. Best in bold and second with underline. These notes are the same to others.

CDPA	FSR	Histology nuclei segmentation						
		TCIA	CRC	KIRC	TNBC	AVG		
		72.92	73.74	73.60	57.72	69.50		
✓		75.77	73.92	73.74	65.38	72.20		
	✓	74.80	76.45	74.65	63.14	72.30		
✓	✓	76.21	78.70	74.42	67.00	<b>74.10</b>		
CDPA	FSR	Prostate MRI segmentation						
		BIDMC	HK	HCRUDB	RUNMC	BMC	UCL	AVG
		68.81	79.86	86.85	80.39	82.23	74.90	78.80
✓		72.73	86.46	88.54	83.94	83.78	75.82	81.90
	✓	74.53	83.58	86.59	79.29	82.33	74.13	80.10
✓	✓	77.56	84.56	87.22	83.31	85.01	77.93	<b>82.60</b>

Table 2: Ablation study of key components of our method in Histology nuclei segmentation and Prostate MRI segmentation task. Please see section 4.2 for details.

Methods	Avg # of Prototypes Uploaded						AVG
	BIDMC	HK	HCRUDB	RUNMC	BMC	UCL	
FedProto	2.00	2.00	2.00	2.00	2.00	2.00	77.10
FPL	2.00	2.00	2.00	2.00	2.00	2.00	77.40
FedContrast	9.00	9.00	9.00	9.00	9.00	9.00	76.50
FedPLVM	212.10	128.44	365.09	323.39	306.84	140.14	77.50
<b>FedBCS</b>	4.00	4.00	4.00	4.00	4.00	4.00	<b>82.60</b>

Table 3: Comparison of communication costs with other federated prototype learning methods in prostate MRI segmentation task. **Avg # of prototypes uploaded** denotes the average number of prototypes each client sends to the server per training epoch.

(a). As shown in Fig. 4(a), our method achieves better performance while maintaining stable convergence. In terms of communication efficiency (Tab. 3), FedProto (Tan et al. 2022) and FPL (Huang et al. 2023b) generate two prototypes by averaging each class’s features, while FedContrast (Zhou and Zheng 2023) and FedPLVM (Wang et al. 2024a) require more prototypes due to local clustering. Although our method leverages multi-level features to address incomplete contextual representation, we maintain efficiency by uploading only two prototypes per class. The results demonstrate that our approach achieves superior segmentation perfor-




Figure 4: Analysis of FedBCS with different temperature (eq. (20)). “Base” denotes FedAvg.

mance without substantial communication overhead.

**Effectiveness of FSR in Addressing Style Bias.** Figure 3 (b) compares our FSR method with input-level amplitude normalization based on FedAvg, evaluated on the Histology nuclei segmentation dataset with pronounced style differences. Instance normalization was consistently applied in all experiments. The results demonstrate that FSR, by addressing style bias accumulation across layers, achieves better performance than methods that only handle style variations at the input level.

### 4.3 Comparison to State-of-the-Arts

As evidenced by the comparative analysis in Table 1, our framework demonstrates significant improvements over existing approaches, validating FedBCS’s enhanced cross-domain generalization capability. Furthermore, the communication efficiency benchmarks in Table 3 reveal that FedBCS maintains superior segmentation performance while incurring only marginal communication overhead.

## 5 Conclusion

In this paper, we addressed two critical challenges in federated medical image segmentation: layerwise style bias accumulation and incomplete contextual learning. Our solution comprises frequency-domain style recalibration to address style variations, and context-aware dual-level prototype alignment to capture comprehensive semantics. Experiments on two datasets demonstrate improved performance.

## Acknowledgements

This work is supported by the National Key Research and Development Program of China (2023YFC2705700), National Natural Science Foundation of China under Grant (62361166629, 62225113, 623B2080), the Major Project of Science and Technology Innovation of Hubei Province (2024BCA003, 2025BEA002), and the Innovative Research Group Project of Hubei Province under Grants 2024AFA017. The supercomputing system at the Supercomputing Center of Wuhan University supported the numerical calculations in this paper.

## References

- Awan, R.; Sirinukunwattana, K.; Epstein, D.; Jefferyes, S.; Qidwai, U.; Aftab, Z.; Mujeeb, I.; Snead, D.; and Rajpoot, N. 2017. Glandular morphometrics for objective grading of colorectal adenocarcinoma histology images. *Scientific reports*, 7(1): 16852.
- Chen, G.; Peng, P.; Ma, L.; Li, J.; Du, L.; and Tian, Y. 2021. Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 458–467.
- Chen, J.; Ma, B.; Cui, H.; and Xia, Y. 2024. Think Twice Before Selection: Federated Evidential Active Learning for Medical Image Analysis with Domain Shifts. In *CVPR*, 11439–11449.
- Chen, Y.; Huang, W.; and Ye, M. 2024. Fair Federated Learning under Domain Skew with Local Consistency and Domain Diversity. In *CVPR*, 12077–12086.
- Ezzeldin, Y. H.; Yan, S.; He, C.; Ferrara, E.; and Avestimehr, A. S. 2023. Fairfed: Enabling group fairness in federated learning. In *AAAI*, volume 37, 7494–7502.
- Hou, L.; Gupta, R.; Van Arnem, J. S.; Zhang, Y.; Sivalenka, K.; Samaras, D.; Kurc, T. M.; and Saltz, J. H. 2020. Dataset of segmented nuclei in hematoxylin and eosin stained histopathology images of ten cancer types. *Scientific data*, 7(1): 185.
- Huang, W.; Ye, M.; and Du, B. 2022. Learn from others and be yourself in heterogeneous federated learning. In *CVPR*, 10143–10153.
- Huang, W.; Ye, M.; Shi, Z.; and Du, B. 2023a. Generalizable heterogeneous federated cross-correlation and instance similarity learning. *IEEE TPAMI*.
- Huang, W.; Ye, M.; Shi, Z.; Li, H.; and Du, B. 2023b. Rethinking Federated Learning With Domain Shift: A Prototype View. In *CVPR*, 16312–16322.
- Huang, W.; Ye, M.; Shi, Z.; Wan, G.; Li, H.; Du, B.; and Yang, Q. 2024. A Federated Learning for Generalization, Robustness, Fairness: A Survey and Benchmark. *TPAMI*.
- Irshad, H.; Montaser-Kouhsari, L.; Waltz, G.; Bucur, O.; Nowak, J.; Dong, F.; Knoblauch, N. W.; and Beck, A. H. 2014. Crowdsourcing image annotation for nucleus detection and segmentation in computational pathology: evaluating experts, automated methods, and the crowd. In *Pacific symposium on biocomputing Co-chairs*, 294–305. World Scientific.
- Jiang, M.; Roth, H. R.; Li, W.; Yang, D.; Zhao, C.; Nath, V.; Xu, D.; Dou, Q.; and Xu, Z. 2023. Fair federated medical image segmentation via client contribution estimation. In *CVPR*, 16302–16311.
- Jiang, M.; Wang, Z.; and Dou, Q. 2022. Harmoff: Harmonizing local and global drifts in federated learning on heterogeneous medical images. In *AAAI*, volume 36, 1087–1095.
- Konečný, J.; McMahan, H. B.; Ramage, D.; and Richtárik, P. 2016. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*.
- Lee, S.; Bae, J.; and Kim, H. Y. 2023. Decompose, adjust, compose: Effective normalization by playing with frequency for domain generalization. In *CVPR*, 11776–11785.
- Lemaître, G.; Martí, R.; Freixenet, J.; Vilanova, J. C.; Walker, P. M.; and Meriaudeau, F. 2015. Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric MRI: a review. *Computers in biology and medicine*, 60: 8–31.
- Li, L.; Liu, J.; Xiao, H.; Zhou, G.; Liu, Q.; and Zhang, Z. 2025. Expert guidance and partially-labeled data collaboration for multi-organ segmentation. *Neural Networks*, 187: 107396.
- Li, Q.; He, B.; and Song, D. 2021. Model-Contrastive Federated Learning. In *CVPR*, 10713–10722.
- Li, T.; Hu, S.; Beirami, A.; and Smith, V. 2021a. Ditto: Fair and robust federated learning through personalization. In *International conference on machine learning*, 6357–6368. PMLR.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated Optimization in Heterogeneous Networks. *arXiv preprint arXiv:1812.06127*.
- Li, X.; Jiang, M.; Zhang, X.; Kamp, M.; and Dou, Q. 2021b. Fed{BN}: Federated Learning on Non-{IID} Features via Local Batch Normalization. In *ICLR*.
- Li, Z.; Shang, Z.; Liu, J.; Zhen, H.; Zhu, E.; Zhong, S.; Sturgess, R. N.; Zhou, Y.; Hu, X.; Zhao, X.; et al. 2023. D-LMBmap: a fully automated deep-learning pipeline for whole-brain profiling of neural circuitry. *Nature Methods*, 20(10): 1593–1604.
- Liao, X.; Liu, W.; Chen, C.; Zhou, P.; Yu, F.; Zhu, H.; Yao, B.; Wang, T.; Zheng, X.; and Tan, Y. 2024. Rethinking the Representation in Federated Unsupervised Learning with Non-IID Data. In *CVPR*, 22841–22850.
- Litjens, G.; Toth, R.; Van De Ven, W.; Hoeks, C.; Kerkstra, S.; Van Ginneken, B.; Vincent, G.; Guillard, G.; Birbeck, N.; Zhang, J.; et al. 2014. Evaluation of prostate segmentation algorithms for MRI: the PROMISE12 challenge. *Medical image analysis*, 18(2): 359–373.
- Liu, Q.; Chen, C.; Qin, J.; Dou, Q.; and Heng, P.-A. 2021. FedDG: Federated Domain Generalization on Medical Image Segmentation via Episodic Learning in Continuous Frequency Space. In *CVPR*, 1013–1023.
- Liu, Q.; Dou, Q.; Yu, L.; and Heng, P. A. 2020. MS-Net: multi-site network for improving prostate segmentation with heterogeneous MRI data. *IEEE transactions on medical imaging*, 39(9): 2713–2724.

- Ma, Y.; Dai, W.; Huang, W.; and Chen, J. 2025. Geometric Knowledge-Guided Localized Global Distribution Alignment for Federated Learning. In *CVPR*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *AISTATS*, 1273–1282.
- Naylor, P.; Laé, M.; Reyat, F.; and Walter, T. 2018. Segmentation of nuclei in histopathology images by deep regression of the distance map. *IEEE TMI*, 38(2): 448–459.
- Nicholas, B.; Anant, M.; Henkjan, H.; John, F.; Justin, K.; et al. 2015. Nci-proc. ieee-isbi conf. 2013 challenge: Automated segmentation of prostate structures. The Cancer Imaging Archive.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Sarfraz, S.; Sharma, V.; and Stiefelwagen, R. 2019. Efficient parameter-free clustering using first neighbor relations. In *CVPR*, 8934–8943.
- Son, H. M.; Kim, M.-H.; Chung, T.-M.; Huang, C.; and Liu, X. 2024. FedUV: Uniformity and Variance for Heterogeneous Federated Learning. In *CVPR*, 5863–5872.
- Sung, C.; Kim, W.; An, J.; Lee, W.; Lim, H.; and Myung, H. 2024. Contextrast: Contextual Contrastive Learning for Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3732–3742.
- Tan, Y.; Long, G.; Liu, L.; Zhou, T.; Lu, Q.; Jiang, J.; and Zhang, C. 2022. FedProto: Federated Prototype Learning across Heterogeneous Clients. In *AAAI*.
- Wan, G.; Huang, W.; and Ye, M. 2024. Federated graph learning under domain shift with generalizable prototypes. In *AAAI*, volume 38, 15429–15437.
- Wang, J.; Jin, Y.; and Wang, L. 2022. Personalizing federated medical image segmentation via local calibration. In *ECCV*, 456–472. Springer.
- Wang, L.; Bian, J.; Zhang, L.; Chen, C.; and Xu, J. 2024a. Taming Cross-Domain Representation Variance in Federated Prototype Learning with Heterogeneous Data Domains. In *NeurIPS*.
- Wang, T.; Zhao, X.; Zhuang, L.; Zhao, H.; Yin, J.; He, Y.; Yu, G.; and Lin, B. 2025. pFedSAM: Personalized Federated Learning of Segment Anything Model for Medical Image Segmentation. *arXiv preprint arXiv:2509.15638*.
- Wang, X.; Li, Z.; Luo, X.; Wan, J.; Zhu, J.; Yang, Z.; Yang, M.; and Xu, C. 2023. Black-box Domain Adaptive Cell Segmentation via Multi-source Distillation. In *MICCAI*, 749–758. Springer.
- Wang, Z.; Yi, R.; Wen, X.; Zhu, C.; and Xu, K. 2024b. Cardiovascular medical image and analysis based on 3d vision: A comprehensive survey. *Meta-Radiology*, 2(4): 100102.
- Ye, M.; Fang, X.; Du, B.; Yuen, P. C.; and Tao, D. 2023. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3): 1–44.
- Zhang, J.; Shan, C.; and Han, J. 2024. FedGMKD: An Efficient Prototype Federated Learning Framework through Knowledge Distillation and Discrepancy-Aware Aggregation. In *NeurIPS*.
- Zhang, R.; Xu, Q.; Yao, J.; Zhang, Y.; Tian, Q.; and Wang, Y. 2023. Federated domain generalization with generalization adjustment. In *CVPR*, 3954–3963.
- Zhao, H.; Dong, W.; Yu, R.; Zhao, Z.; Du, B.; and Xu, Y. 2024a. Morestyle: relax low-frequency constraint of fourier-based image reconstruction in generalizable medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 434–444. Springer.
- Zhao, H.; Gu, Y.; Zhao, Z.; Du, B.; Xu, Y.; and Yu, R. 2024b. WIA-LD2ND: Wavelet-based image alignment for self-supervised low-dose CT denoising. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 764–774. Springer.
- Zhao, X.; Li, P.; Luo, X.; Yang, M.; Chang, S.; and Li, Z. 2024c. Sam-Driven Weakly Supervised Nodule Segmentation with Uncertainty-Aware Cross Teaching. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Zhao, X.; Li, Z.; Luo, X.; Li, P.; Huang, P.; Zhu, J.; Liu, Y.; Zhu, J.; Yang, M.; Chang, S.; et al. 2024d. Ultrasound Nodule Segmentation Using Asymmetric Learning with Simple Clinical Annotation. *IEEE TCSVT*.
- Zhou, Q.; and Zheng, G. 2023. FedContrast-GPA: Heterogeneous Federated Optimization via Local Contrastive Learning and Global Process-Aware Aggregation. In *MICCAI*, 660–670. Cham: Springer Nature Switzerland.
- Zhuang, L.; Zhang, Y.; Yuan, G.; Zhao, X.; and Jiang, Z. 2025. ActiveSSF: An Active-Learning-Guided Self-Supervised Framework for Long-Tailed Megakaryocyte Classification. *arXiv preprint arXiv:2502.08200*.