

FedOAP: Organ Agnostic Tumor Segmentation Using Personalized Federated Learning

Ishmam Tashdeed*

Md. Atiqur Rahman*

Sabrina Islam

Md. Azam Hossain

Islamic University of Technology

*Equal Contribution

Abstract

Personalized federated learning (PFL) possesses the unique capability of preserving data confidentiality among clients while tackling the data heterogeneity problem of non-independent and identically distributed (Non-IID) data. Its advantages have led to widespread adoption in domains such as medical image segmentation. However, the existing approaches mostly overlook the potential benefits of leveraging shared features across clients, where each client contains segmentation data of different organs. In this work, we introduce a novel personalized federated approach for organ agnostic tumor segmentation (FedOAP), that utilizes cross-attention to model long-range dependencies among the shared features of different clients and a boundary-aware loss to improve segmentation consistency. FedOAP employs a decoupled cross-attention (DCA), which enables each client to retain local queries while attending to globally shared key-value pairs aggregated from all clients, thereby capturing long-range inter-organ feature dependencies. Additionally, we introduce perturbed boundary loss (PBL) which focuses on the inconsistencies of the predicted mask’s boundary for each client, forcing the model to localize the margins more precisely. We evaluate FedOAP on diverse tumor segmentation tasks spanning different organs. Extensive experiments demonstrate that FedOAP consistently outperforms existing state-of-the-art federated and personalized segmentation methods.




Figure 1: Motivation for FedOAP. FedOAP aims to learn a shared feature representation across clients by aggregating informative tokens from different organs while preserving client-specific data privacy. This enables improved tumor segmentation by leveraging inter-organ knowledge and producing personalized predictions guided by cross-client feature dependencies.

1 Introduction

Federated Learning (FL) is a decentralized machine learning paradigm that enables collective model training across multiple data sources without the need to share raw data (Li et al., 2020a). By ensuring that data remains local to its source, FL inherently adheres to privacy regulations and data protection standards (Voigt and Von dem Bussche, 2017). This makes FL particularly well-suited for sensitive domains such as healthcare (Horst et al., 2025; Kaissis et al., 2020; Kwak and Bai, 2023). However, non-independent and identically distributed (Non-IID) data or heterogeneous data poses as a great hurdle for most FL methods (Zhao et al., 2018). A universally shared global model is usually insufficient to tackle the diverse types of data encountered across different clients. Heavily heterogeneous data may introduce client drift which severely degrades the model’s performance (Karimireddy et al., 2020). Specially in the domain of medical image analysis, it is common to have inter-client incon-

sistencies as individual clients often address distinct clinical tasks, patient populations, cultural contexts, and imaging equipment (Guan et al., 2024).

Personalized federated learning (PFL) tries to address this problem by maintaining distinct local models for each client that conforms to its unique distribution while keeping the data localized (Tan et al., 2022). This is achieved by a variety of methods such as regularization of local and global weight divergence (Karimireddy et al., 2020; Li et al., 2020b), meta-learning (Fallah, Mokhtari, and Ozdaglar, 2020; T Dinh, Tran, and Nguyen, 2020), transfer-learning (Afzali and Shamsinejadbabaki, 2025; Tan et al., 2023), decoupling parameters into shared and private segments (Arivazhagan et al., 2019), knowledge distillation (Li and Wang, 2019; Lin et al., 2020), clustering clients based on similar tasks or data distributions (Briggs, Fan, and Andras, 2020; Sattler, Müller, and Samek, 2020), among others. In medical image segmentation, a field often bound by strict data privacy policies, PFL is heavily relied upon. The most pervasive strategy is to share the feature extractor while training a separate mask prediction head for each client to accommodate local data distributions (Wang et al., 2023; Xie et al., 2024; Wang, Jin, and Wang, 2022; Jiang et al., 2023; Liu et al., 2025).

Despite recent strides made in tumor segmentation using PFL, existing approaches have largely overlooked the potential benefits of leveraging cross-organ information to enhance segmentation performance. In a federated setting, clients often possess segmentation data regarding anatomically different organs such as brain, liver, or breast. We hypothesize that a shared feature extractor with appropriate personalization can greatly benefit the segmentation process across all the clients. However, most existing methods are not well suited to accommodate data originating from different organs and often from different modalities (Madni, Umer, and Foresti, 2023; Babar, Qureshi, and Koubaa, 2024). This can potentially be handled by modeling long-range dependencies via cross-attention mechanisms, enabling the aggregation of global contextual information across the feature space of the clients. Furthermore, introducing personalization during feature extraction across clients not only aids in the adaptation of the local models to the heterogeneous data distribution but also preserves data privacy (Wang et al., 2023). For effective client personalization local fine-tuning is typically required, where each client further trains its model for a few epochs on its private dataset. During this step, inconsistencies in the predicted masks greatly degrades the segmentation performance (Wang et al., 2023; Wang, Jin, and Wang, 2022). Moreover, the few methods that attempt to

resolve this often introduce substantial computational or memory overhead.

In this work, we propose FedOAP, a novel personalized federated learning framework for organ-agnostic tumor segmentation, which integrates cross-organ aggregated feature extraction with a novel perturbed boundary loss to enhance both generalization and boundary precision across heterogeneous clients. To handle distinct heterogeneous organ data from different clients, FedOAP utilizes a shared encoder-decoder while introducing decoupled cross-attention (DCA) which decouples query-key embedding layers and employs cross-attention on aggregated key-value pairs from all clients. Inspired by (Wang et al., 2023), keeping the query layer local to each client, we ensure data privacy as query embeddings generally represent the features of the local data itself. Aggregating key-value embedding pairs across multiple clients enables the local query to attend to supportive features from different organs. Additionally, FedOAP adheres to established practices in PFL by incorporating a client-specific spatial adapter, while introducing a novel perturbed boundary loss (PBL) to enhance segmentation accuracy along proposed tumor boundaries in the fine-tuning process. PBL quantifies pixel-level inconsistencies between local predictions and ground truth and injects stochastic noise into the erroneous boundary regions. By formulating a loss based on these perturbed discrepancies, we direct greater supervision focus toward critical boundary pixels, thereby promoting the learning of more robust and precise boundary representations.

In general, our contributions are as follows:

- We propose a novel PFL framework for tumor segmentation specifically designed to handle segmentation tasks across heterogeneous organ datasets named FedOAP.
- We introduce DCA, a mechanism that decouples the query and key-value embedding layers. The query remains local to each client, while the aggregated key-value pairs from all clients are used to enable cross-organ feature fusion through cross-attention.
- In addition, we propose PBL to improve proposed tumor mask accuracy. It identifies inconsistencies between predicted and ground truth boundaries while injecting stochastic noise into these regions, and formulates a targeted loss to enhance supervision on boundary pixels, thereby improving segmentation robustness and precision during client side fine-tuning.

We conduct comprehensive experiments on three tu-

mor segmentation tasks from different organs: brain, liver, and breast. FedOAP achieves superior performance compared to existing state-of-the-art PFL methods, demonstrating its strong generalization capabilities and its effectiveness in handling inter-organ variability. We provide related work in Section A of the supplementary document.

2 Methodology

2.1 Problem Definition

Image segmentation is defined as a dense pixel-wise classification problem where given an input image, the model classifies each pixel to a predefined semantic class label. In the context of our binary tumor segmentation task, each pixel is classified as either belonging to a tumor region or the background. The trained segmentation model $f_\theta(\mathcal{X}) \rightarrow \mathcal{Y}$, parameterized by θ , takes an image $x_i \in \mathbb{R}^{H \times W \times C}$ and outputs a segmentation mask $y_i \in \{0, 1\}^{H \times W}$. Here, H , W , and C are the height, width, and number of channels respectively. Thus, the segmentation task can be formulated as:

$$f_\theta(x_i) = \operatorname{argmax}_{y_i \in \mathcal{Y}} \mathbb{P}(y_i | x_i; \theta) \quad (1)$$

We consider a federated setting with $\mathcal{K}$ clients, each possessing a data distribution $\mathcal{D}_{k \in \mathcal{K}}$ corresponding to a distinct organ. For the k -th client, a set of N_k image-mask pairs, denoted as $\{x_{i,k}, y_{i,k}\}_{i=1}^{N_k}$ constitute the non-iid dataset such that $\{x_{i,k}, y_{i,k}\} \sim \mathcal{D}_k$. In classical FL scenarios such as (McMahan et al., 2017), the goal is to learn a shared model θ in $\mathcal{T}$ rounds by initially training on a client's local distribution $\mathcal{D}_k$ and aggregating them each round as:

$$\theta^{t+1} = \theta^t - \alpha \mathbb{E}_k \mathbb{E}_{(x_{i,k}, y_{i,k}) \sim \mathcal{D}_k} \nabla \mathcal{L}_k(f_\theta(x_{i,k}), y_{i,k}) \quad (2)$$

Here, α is the learning rate, $t \in \{0, 1, \dots, \mathcal{T} - 1\}$ denotes the communication round, and $\mathcal{L}_k$ is the objective function for the k -th client. However, a single shared model θ yields suboptimal performance due to highly heterogeneous data distribution across clients. PFL tackles this problem by learning k separate models denoted as $\{\theta_k\}_{k=1}^{\mathcal{K}}$, where each θ_k aims to fit the client's distribution $\mathcal{D}_k$.

2.2 Framework Overview

Inspired by insights discussed in (Bengio, Courville, and Vincent, 2013), we aim to utilize a cross-organ shared feature representation to enhance the performance of tumor segmentation in heterogeneous federated settings. Our proposed framework, as shown in Figure 2, is composed of two sequential stages:

Algorithm 1 FedOAP Algorithm.

Input: Local datasets $\{\mathcal{X}_k, \mathcal{Y}_k\}_{k=1}^{\mathcal{K}} \sim \mathcal{D}_k$, total number of rounds $\mathcal{T}$, number of clients $\mathcal{K}$, number of fine-tuning epochs N_ϵ .

Output: Local parameters θ_k^* .

```

1: Initialize personalized parameters for each client  $\rho_k^{t=0}$ 
2: Initialize shared parameters for all clients  $\gamma^{t=0}$ 
3: for  $t = 0, 1, \dots, \mathcal{T} - 1$  do
4:   for  $k = 1, 2, \dots, \mathcal{K}$  do
5:      $\gamma_k^t, \rho_k^{t+1} \leftarrow \text{DCA}(\{\mathcal{X}_k, \mathcal{Y}_k\}, \gamma^t, \rho_k^t)$ 
6:   end for
7:    $\gamma^{t+1} \leftarrow \frac{1}{\mathcal{K}} \sum_{i=1}^{\mathcal{K}} \gamma_i^t$ 
8: end for
9: for  $k = 1, 2, \dots, \mathcal{K}$  do
10:   $\hat{\theta}_k \leftarrow \gamma^{\mathcal{T}} \cup \rho_k^{\mathcal{T}}$ 
11: end for
12: for  $k = 1, 2, \dots, \mathcal{K}$  do
13:   for  $i = 0, 1, \dots, N_\epsilon$  do
14:     $\theta_k^* \leftarrow \text{PBL}(\{\mathcal{X}_k, \mathcal{Y}_k\}, \hat{\theta}_k)$ 
15:   end for
16: end for
17: return  $\{\theta_k^*\}_{k=1}^{\mathcal{K}}$ 

```

federated representation alignment and *client-specific personalization*, where the proposed decoupled cross-attention (DCA) and perturbed boundary loss (PBL) are performed respectively. During federated representation alignment, each client k maintains a local model θ_k , which is trained on its private dataset $\mathcal{D}_k$ over $\mathcal{T}$ communication rounds. The model parameters are partitioned into two disjoint subsets: shared parameters γ_k and personalized parameters ρ_k such that $\theta_k = \gamma_k \cup \rho_k$. At each round t , the shared parameters are aggregated through averaging: $\gamma^{t+1} = \frac{1}{\mathcal{K}} \sum_{k=1}^{\mathcal{K}} \gamma_k^t$ following (McMahan et al., 2017). These aggregated parameters are broadcast back to the clients and merged with the personalized parameters to create a set of preliminary models $\hat{\theta}_k^{t+1}$ such that: $\hat{\theta}_k^{t+1} = \gamma^{t+1} \cup \rho_k^t$. The client-specific personalization phase, fine-tunes each local model $\hat{\theta}_k$ on its dataset without further parameter sharing by minimizing a composite loss. This process yields the final personalized set of models $\{\theta_k^*\}_{k=1}^{\mathcal{K}}$, each tailored to its local data distribution while benefiting from cross-client feature alignment.

2.3 Federated Representation Alignment

To effectively capture long-range dependencies across heterogeneous client representations, we slightly modify the self-attention mechanism originally introduced in (Vaswani et al., 2017). Prior work, such as (Wang et al., 2023), has demonstrated the efficacy of self-attention in modeling long-range feature interactions




Figure 2: (a) Overview of the proposed FedOAP framework. Each client $k \in \mathcal{K}$ trains a local model θ_k (composed of an encoder-decoder architecture with a client-specific spatial adapter ϕ_k) on its private dataset $\mathcal{D}_k$ for $\mathcal{T}$ rounds. During this phase, DCA separates the shared parameters γ_k for global aggregation. Finally, each client fine-tunes its model using PBL, resulting in the personalized parameters θ_k^* . (b) The proposed Decoupled Cross-Attention (DCA) mechanism. The private query embeddings $\mathbf{q}_k$ attend to the concatenated key-value $(\mathbf{k}_k, \mathbf{v}_k)$ pairs, capturing cross-organ, long-range dependencies without exposing private client representations. (c) The Perturbed Boundary Loss (PBL) mechanism. The inconsistencies between predicted masks $\hat{\mathcal{Y}}_i$ and ground truth $\mathcal{Y}_i$ is used as a guide to inject noise, reinforcing supervision on uncertain regions to improve boundary precision.

while simultaneously accommodating inter-client variability in federated settings. Concurrently, methods including (Li et al., 2021; Arivazhagan et al., 2019; Dai et al., 2024) have emphasized the importance of personalizing specific model components to better adapt to client-specific data distributions. Let $z \in \mathbb{R}^{n \times d}$ denote the input to a self-attention block. Here, n is the sequence length and d is the feature embedding dimension. The *query* ($\mathbf{q} \in \mathbb{R}^{n \times d_q}$), *key* ($\mathbf{k} \in \mathbb{R}^{n \times d_k}$), and *value* ($\mathbf{v} \in \mathbb{R}^{n \times d_v}$) are calculated as: $\mathbf{q} = zW_{query}$, $\mathbf{k} = zW_{key}$, $\mathbf{v} = zW_{value}$, where $W_{query}, W_{key}, W_{value} \in \mathbb{R}^{d \times d}$ are learnable projection matrices. For simplicity, we assume $d_q = d_k = d_v = d$. The scaled dot-product attention $A(z)$ is calculated as:

$$A(\mathbf{q}, \mathbf{k}, \mathbf{v}) = \text{softmax}\left(\frac{\mathbf{q}\mathbf{k}^T}{\sqrt{d}}\right) \mathbf{v} \quad (3)$$

Following (Wang et al., 2023), we propose *Decoupled Cross-Attention* (DCA). In this formulation, each client k maintains a personalized query embedding $\mathbf{q}_k$, while attending to globally shared key-value embeddings $(\mathbf{k}_{shared}, \mathbf{v}_{shared})$ concatenated from all clients as $\mathbf{k}_{shared} = \text{concat}(\mathbf{k}_1, \mathbf{k}_2, \dots, \mathbf{k}_K) \in \mathbb{R}^{N \times d}$ and $\mathbf{v}_{shared} = \text{concat}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K) \in \mathbb{R}^{N \times d}$, where N is the total number of aggregated tokens. Additionally, to help our models adapt to the local distributions, we append a personalized spatial adapter $\phi_k \in \mathcal{K}$ to the decoder of each client. Following established notation, the personalized parameters for client k are $\rho_k = \{\mathbf{q}_k, \phi_k\}$, while the shared parameters are γ_k , such that $\gamma_k \cap \rho_k = \emptyset$. During training, client-specific

parameters ρ_k are updated as:

$$\rho_k^{t+1} = \rho_k^t - \alpha \mathbb{E}_{(x_{i,k}, y_{i,k}) \sim \mathcal{D}_k} \nabla \mathcal{L}_k(f_{\theta_k}(x_{i,k}), y_{i,k}) \quad (4)$$

In parallel, the shared parameters γ_k are collaboratively optimized across all clients as:

$$\gamma_k^{t+1} = \gamma_k^t - \alpha \mathbb{E}_k \mathbb{E}_{(x_{i,k}, y_{i,k}) \sim \mathcal{D}_k} \nabla \mathcal{L}_k(f_{\theta_k}(x_{i,k}), y_{i,k}) \quad (5)$$

After completing $\mathcal{T}$ communication rounds, each client obtains an intermediate parameters $\hat{\theta}_k$, which are further updated in the subsequent phase.

2.4 Perturbation Guided Calibration

Different clients possess varying modalities and sources of data and the absence of a unified imaging protocol being enforced across all clients introduces domain shift (Zhu and Tang, 2024; Pan et al., 2025). This introduces minor inconsistencies between the predicted mask and the ground truth which greatly affects a model’s performance. Degradation caused by this is often more noticeable when the clients contain remarkably heterogeneous data such as from different organs. Additionally, inconsistencies can often guide the model towards improving segmentation results (Ji et al., 2021; Wang et al., 2023; Wang, Jin, and Wang, 2022). Focusing on these discrepancies during the local fine-tuning phase, our models can greatly acclimate to the client’s unique distribution. Emboldened by this, we propose a *Perturbed Boundary Loss* (PBL), which aims to adjust the model’s prediction guided by the differences between the predicted and the actual mask in a robust manner.

Given intermediate parameters $\hat{\theta}_k$, the model infers each sample $(x_{i,k}, y_{i,k}) \sim \mathcal{D}_k$ to obtain predicted masks $\hat{y}_{i,k}$. The inconsistency mask $\delta_{i,k} \in \{0, 1\}^{H \times W}$ is computed using:

$$\delta_{i,k} = \begin{cases} 1 & \text{if } |y_{i,k} - \sigma(\hat{y}_{i,k})| > \tau \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Here, we pass $\hat{y}_{i,k}$ through the sigmoid function σ and $\tau \in (0, 1)$ is a cut-off threshold value. We introduce additive noise $\epsilon_{i,k} \sim \mathcal{N}$ to the regions indicated as inconsistencies from the inconsistency mask $\delta_{i,k}$ to get the noisy mask: $\tilde{y}_{i,k} = \hat{y}_{i,k} + \delta_{i,k} \cdot \epsilon_{i,k}$. Here, $\mathcal{N}$ is the gaussian distribution with mean of 0 and variance of 0.1. We simultaneously compute two losses: Segmentation loss $\mathcal{L}_s$ using the clean predicted masks $\{\hat{y}_{i,k}\}_{k=1}^K$ and Perturbed loss $\mathcal{L}_p$ with the noisy masks $\{\tilde{y}_{i,k}\}_{k=1}^K$. The composite objective function is defined as $\mathcal{L} = (1 - \lambda)\mathcal{L}_s + \lambda\mathcal{L}_p$, where $\lambda \in [0, 1]$ is the trade-off parameter between the two objectives. The complete algorithm regarding the FedOAP framework is detailed in Algorithm 1.

3 Experiments and Results

3.1 Datasets

We evaluate FedOAP on three publicly available datasets encompassing diverse organs and imaging modalities. Specifically, we conduct tumor segmentation experiments on (i) dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) of the breast, (ii) structural MRI scans of the brain, and (iii) computed tomography (CT) images of the liver. To assess the cross-organ generalization capability of our framework, we further conduct evaluation on a dataset consisting of CT images of the lung, which is not included in the training phase.

All input images and their corresponding ground truth segmentation masks are uniformly resized to a spatial resolution of 128×128 . 10% of each dataset is reserved for testing, while 10% of the remaining data is allocated for validation. The rest is used for training the model. The detailed characteristics of each dataset are as follows:

- **BreastDM** dataset (Zhao et al., 2023) contains 232 DCE-MRI scans of breast tissue, annotated for tumor presence. Among these, 147 samples are malignant cases, while 85 are benign. The DCE-MRI modality offers dynamic contrast information, which poses a distinct challenge for tumor boundary delineation due to its temporal nature. For our experiments, we ignore the benign cases and only work with the malignant ones.

- **BraTS 2020** dataset (Menze et al., 2014; Bakas et al., 2017, 2018) consists of 369 multi-contrast brain MRI scans annotated for glioma tumor sub-regions. Although the original BraTS dataset includes four imaging modalities: T1, T1-contrast (T1c), T2, and FLAIR, we aggregate the sequences and construct a single-channel input focused on whole tumor segmentation. This setup allows us to simplify modality dependency and emphasize anatomical variance across clients.
- **LiTS** dataset (Bilic et al., 2023) includes 201 contrast-enhanced abdominal CT volumes with annotated liver tumors. For consistency with the other datasets and computational efficiency, each volume is preprocessed into two-dimensional slices and normalized as single-channel grayscale images.
- **Lung Cancer Segmentation dataset** (Nam, Panina, and Pak, 2024) consists of 972 lung CT slices and associated tumor segmentation masks. This dataset is used exclusively for out-of-distribution testing. It serves as a benchmark to evaluate the cross-organ generalizability of the proposed framework under domain and modality shifts.

3.2 Evaluation Metrics

To evaluate the performance of our proposed FedOAP framework on tumor segmentation tasks, we use the mean of the Dice Similarity Coefficient (mDSC) (Milletari, Navab, and Ahmadi, 2016) of the clients as the primary evaluation metric. Given a predicted mask $\hat{y} \in \{0, 1\}^{H \times W}$ and the corresponding ground truth $y \in \{0, 1\}^{H \times W}$, the Dice score is defined as:

$$\text{Dice}(\hat{y}, y) = \frac{2 \cdot |\hat{y} \cap y|}{|\hat{y}| + |y|} \quad (7)$$

Here, $|\hat{y} \cap y|$ represents the correctly predicted tumor pixels, while $|\hat{y}|$ denote the total predicted tumor pixels and $|y|$ depict total true tumor pixels. Higher values of the dice score indicate a good overlap between predicted and true tumor regions.

3.3 Implementation Details

Our framework is implemented in Python 3.11.13 with PyTorch 2.7.1, CUDA 12.6, and the Flower 1.4.0 framework (Beutel et al., 2020) for client-server communication. Experiments are conducted on a single NVIDIA RTX A4000 GPU. We set the number of communication rounds to $\mathcal{T} = 5$ with $\mathcal{K} = 3$ clients, each training for one local epoch per round, followed by additional fine-tuning epochs $N_\epsilon = 2$. The best-performing parameters are retained. Training uses

Table 1: Comparison of FedOAP with state-of-the-art methods across multiple datasets. Reported values denote the mean and standard deviation over five runs for each client, along with the average across clients. The best performing results are highlighted in bold.

Method	BreastDM		BraTS		LiTS		Average(%) $\uparrow$
	Dice(%) $\uparrow$	Variance	Dice(%) $\uparrow$	Variance	Dice(%) $\uparrow$	Variance	
FedAvg	50.91	± 3.62	1.28	± 1.17	9.71	± 0.95	20.63
FedAvgM	43.29	± 4.66	0.96	± 1.48	10.66	± 1.1	18.30
FedAdagrad	4.51	± 0.27	35.16	± 6.01	3.63	± 0.19	14.43
FedPer	59.91	± 23.47	19.15	± 23.28	42.06	± 24.12	40.37
FedRep	85.92	± 1.99	75.56	± 19.09	49.75	± 2.94	70.41
FedDP	91.47	± 0.72	94.04	± 0.56	83.88	± 3.27	89.80
FedOAP (Ours)	94.39	± 0.62	95.31	± 0.41	87.22	± 0.57	92.31

AdamW (Loshchilov and Hutter, 2017) with learning rate $\alpha = 10^{-4}$, cosine annealing down to 10^{-6} , and weight decay 10^{-5} . Each minibatch has 16 samples, and all random seeds are fixed to 42 for reproducibility. For PBL, the threshold value is set to $\tau = 0.75$, and the composite loss trade-off parameter λ defaults to 0.1. Justification about $\mathcal{K}$ included in Appendix C.

3.4 Comparison with State-of-the-Art Methods

We evaluate the performance of our proposed FedOAP framework against a suite of established federated learning (FL) and personalized federated learning (PFL) baselines. Among the FL baselines, we include the canonical FedAvg algorithm (McMahan et al., 2017) along with its notable extensions FedAvgM (Hsu, Qi, and Brown, 2019) and FedAdagrad (Reddi et al., 2020). For PFL comparisons, we consider state-of-the-art methods such as FedPer (Ari-vazhagan et al., 2019), FedRep (Collins et al., 2021), and FedDP (Wang et al., 2023), which have demonstrated strong performance in heterogeneous data environments. To keep our comparisons fair, we evaluate each method under the same conditions and using the same hyperparameters.

Table 1 presents the quantitative evaluation of all competing methods, demonstrating that the proposed FedOAP framework consistently outperforms both standard FL and PFL baselines across all datasets. It is evident that regular FL methods are not able to capture the cross-client and cross-domain feature space, resulting in their poor performance. We can see FedAvg typically outperforming FedAvgM and FedAdagrad. These results highlight the inherent limitations of global model aggregation in the presence of strong data heterogeneity. Notably, while most FL methods perform poorly on the BraTS and LiTS datasets, FedAdagrad exhibits relatively higher accuracy on BraTS. Looking at most of the results, FedAvg and

FedAvgM mostly fail to identify even a single tumor pixel whereas FedAdagrad barely identifies a basic outlined shape. This could be attributed to the presence of larger but structurally complex tumors, where simple global updates may suffice to capture coarse tumor characteristics without requiring precise boundary delineation.

The PFL methods are comparatively better as they are specifically tailored to personalize every client to its data distribution. FedPer and FedRep both achieve significant gains when compared to FL methods. But their high inter-client variance points to an inability to handle cross-client knowledge transfer. Both FedDP and FedOAP (ours) exhibit superior results across different clients. This further supports the ability of the attention mechanism to learn from a shared representation. Furthermore, our method outperforms FedDP because of our focus in perfecting the prediction boundaries using noisy supervision. The benefits of this is particularly evident in the BraTS dataset, where complex tumor morphologies demand fine-grained spatial awareness. Additionally, our proposed personalized cross-attention mechanism enables clients to integrate context-aware features from different organ domains while preserving privacy through local query retention.

To qualitatively analyze the performance, we visualize the segmentation results of the methods in Figure 3. For a randomly selected instance from each dataset, we display the input image, corresponding ground truth mask, and the predicted segmentation masks generated by the competing methods. Nearly all the methods have a difficult time delineating complex tumor boundaries, particularly in anatomically diverse scenarios. In contrast, the proposed FedOAP framework demonstrates a notable ability to adapt to intricate spatial patterns and generate precise segmentation masks. The comparatively tiny tumor regions in the BreastDM dataset is difficult to capture for the




Figure 3: Qualitative comparison of FedOAP with existing FL and PFL methods across different organ datasets. Each row depicts a random sample from a client, illustrating segmentation performance across methods.

Table 2: Generalization results comparing FedOAP with FedDP on the Lung Cancer Segmentation dataset. Best scores are in bold.

Method	Dice(%) $\uparrow$	
	Zero-Shot	Fine-Tuned
FedDP	1.1	65.4
FedOAP (Ours)	7.1	72.3

FL methods. Only the personalized models have localized these regions with varying levels of accuracy. This further supports the claim of how high data heterogeneity can devastate model performance. Although they may have localized the regions, most PFL models struggle with predicting the irregular boundaries. In comparison, FedOAP achieves superior boundary adherence and spatial precision. This is particularly evident in samples from both the BraTS and LiTS datasets, where tumors display substantial morphological variability. These results affirm the efficacy of our boundary-aware optimization strategy and decoupled cross-attention mechanism, which combined enable robust and context-sensitive segmentation.

3.5 Generalization on Unseen Data

As FedOAP aims to leverage a shared representation across clients and incorporate organ-specific features from to enhance its own predictions, we hypothesize that it should perform significantly better when tested on a data distribution that was unseen during training. To empirically validate our hypothesis, we compare FedOAP (ours) against FedDP, the best PFL framework in previous evaluations, on a completely unseen lung cancer segmentation dataset (Nam, Panina, and Pak, 2024). The results of our experiment

are reported in Table 2. Initially, we evaluate both FedOAP and FedDP on the new dataset in a zero-shot setting. And both of them achieve a low dice score (1.1% for FedDP and 7.1% for FedOAP) as they are not personalized to the new domain. However, FedOAP consistently outperforms FedDP, demonstrating its capacity to exploit shared representations via the proposed DCA mechanism, even without client-specific tuning. Finally, we fine-tune both methods on the new dataset for only two epochs without any inter-client parameter exchange. This simulates a realistic scenario where a new client adapts its model using only previously shared global knowledge. Post fine-tuning, FedOAP achieves a noticeable performance improvement and outperforms FedDP by a significant margin, thereby affirming the effectiveness of its design in enabling rapid adaptation through robust feature generalization.

3.6 Ablation Study

We thoroughly investigate the individual and cumulative contribution of each proposed component, namely DCA, spatial adapter, and PBL. To this end, we incrementally incorporate these modules into the base framework and report the mean Dice Similarity Coefficient (mDSC%) across all clients in Table 3. In the baseline setting, none of the proposed enhancements are included. Even in this configuration, our framework achieves competitive performance on the BreastDM dataset. This observation may be attributed to the well-defined tumor boundaries in DCE-MRI breast images, which makes the segmentation task less reliant on complex contextual modeling.

The incorporation of the proposed DCA mechanism yields a marked improvement in performance on the

Table 3: Ablation study showing mean Dice scores (%) $\uparrow$ for the proposed components: DCA, spatial adapter, and PBL across datasets. Best results are highlighted in bold.

DCA	Spatial Adapter	PBL	BreastDM	BraTS	LiTS
$\times$	$\times$	$\times$	85.1	25.9	5.3
$\checkmark$	$\times$	$\times$	88.9	52.7	8.3
$\checkmark$	$\checkmark$	$\times$	89.99	65.8	86.1
$\checkmark$	$\checkmark$	$\checkmark$	94.4	95.3	87.3

BraTS dataset. This empirical finding reinforces our hypothesis that leveraging cross-organ feature representations substantially enhances the model’s ability to capture complex structural patterns, particularly in cases involving intricate brain tumor morphology. In contrast, the performance gains observed on the BreastDM and LiTS datasets are relatively modest, suggesting that the benefits of DCA are most pronounced where the target structures exhibit high variability and spatial complexity.

The inclusion of the spatial adapter significantly improves the performance of the challenging LiTS dataset (from 8.6% to 86.1%). This underscores the critical importance of client-specific adaptation in effectively handling the low-contrast, structurally cluttered, and highly heterogeneous nature of abdominal CT scans. Additionally, all the other clients possess data from a different modality (MRI) when compared to the LiTS dataset, which contains CT images, further adding to its heterogeneity. A moderate performance gain is also observed on the BraTS dataset, suggesting that personalized modeling facilitates improved localization in structurally variable brain tumor regions. In contrast, the performance on the BreastDM dataset remains relatively unchanged, indicating a potential performance plateau wherein further personalization yields diminishing returns.

Finally, we include PBL to complete the set of proposed enhancements. This yields a marked improvement in segmentation performance on the BraTS dataset, attributed to the noisy supervision mechanism which encourages the model to better delineate complex tumor boundaries seen on brain MRI scans. Similarly, the BreastDM dataset benefits from enhanced boundary precision, resulting in more accurate segmentation masks for the comparatively small and well-defined tumor regions. However, the performance gain on the LiTS dataset remains marginal, attributed to its heavy reliance on client-specific tuning.

3.7 Efficiency Analysis

Table 4 depicts the transmission cost between the clients and the server for each communication round.

Table 4: Comparison of transmission overhead (in MB) per communication round between a client and server for different federated learning methods.

Method	Transmission Size (MB)
FedAvg	~ 118.37
FedAvgM	~ 118.37
FedAdagrad	~ 118.37
FedPer	~ 71.88
FedRep	~ 71.88
FedDP	~ 130.44
FedOAP (Ours)	~ 130.44

FedOAP introduces a modest increase in transmission overhead relative to FedPer and FedRep while delivering superior performance. The lower transmission cost of these methods arises from their design choice of transmitting only the encoder parameters to the server while retaining the decoder locally on each client. In contrast, FedOAP personalizes only the query vectors $\mathbf{q}_k$ and the spatial adapter ϕ_k , with all remaining parameters shared with the server. When compared to FedAvg, FedAvgM, and FedAdagrad, FedOAP incurs slightly higher transmission cost, as it requires each client to communicate the concatenated key–value representations aggregated from all participants, whereas the baselines only transmit parameters from a single client. The key observation is that FedOAP achieves transmission overheads on par with FedDP while consistently outperforming it in model performance. Additional experiments about efficiency analysis are provided in section B of the supplementary document.

4 Conclusion

We introduced FedOAP, a PFL framework for organ-agnostic tumor segmentation that combines decoupled cross-attention with a boundary-aware loss. FedOAP consistently outperforms existing FL and PFL methods in accuracy and generalization across heterogeneous and unseen domains. Future work will explore its scalability with larger datasets and under irregular federated conditions.

References

- Afzali, A.; and Shamsinejadbabaki, P. 2025. PHiFL-TL: Personalized hierarchical federated learning using transfer learning. *Future Generation Computer Systems*, 166: 107672.
- Arivazhagan, M. G.; Aggarwal, V.; Singh, A. K.; and Choudhary, S. 2019. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*.
- Babar, M.; Qureshi, B.; and Koubaa, A. 2024. Investigating the impact of data heterogeneity on the performance of federated learning algorithm using medical imaging. *Plos one*, 19(5): e0302539.
- Bakas, S.; Akbari, H.; Sotiras, A.; Bilello, M.; Rozycki, M.; Kirby, J. S.; Freymann, J. B.; Farahani, K.; and Davatzikos, C. 2017. Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1): 1–13.
- Bakas, S.; Reyes, M.; Jakab, A.; Bauer, S.; Rempfler, M.; Crimi, A.; Shinohara, R. T.; Berger, C.; Ha, S. M.; Rozycki, M.; et al. 2018. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *arXiv preprint arXiv:1811.02629*.
- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828.
- Beutel, D. J.; Topal, T.; Mathur, A.; Qiu, X.; Fernandez-Marques, J.; Gao, Y.; Sani, L.; Li, K. H.; Parcollet, T.; De Gusmão, P. P. B.; et al. 2020. Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390*.
- Bilic, P.; Christ, P.; Li, H. B.; Vorontsov, E.; Ben-Cohen, A.; Kaissis, G.; Szeskin, A.; Jacobs, C.; Mamani, G. E. H.; Chartrand, G.; et al. 2023. The liver tumor segmentation benchmark (lits). *Medical image analysis*, 84: 102680.
- Briggs, C.; Fan, Z.; and Andras, P. 2020. Federated learning with hierarchical clustering of local updates to improve training on non-IID data. In *2020 international joint conference on neural networks (IJCNN)*, 1–9. IEEE.
- Chang, Q.; Qu, H.; Zhang, Y.; Sabuncu, M.; Chen, C.; Zhang, T.; and Metaxas, D. N. 2020. Synthetic learning: Learn from distributed asynchronous discriminator gan without sharing medical image data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13856–13866.
- Collins, L.; Hassani, H.; Mokhtari, A.; and Shakkottai, S. 2021. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, 2089–2099. PMLR.
- Dai, Q.; Wei, D.; Liu, H.; Sun, J.; Wang, L.; and Zheng, Y. 2024. Federated modality-specific encoders and multimodal anchors for personalized brain tumor segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1445–1453.
- Dong, N.; and Voiculescu, I. 2021. Federated contrastive learning for decentralized unlabeled medical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 378–387. Springer.
- Fallah, A.; Mokhtari, A.; and Ozdaglar, A. 2020. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in neural information processing systems*, 33: 3557–3568.
- Guan, H.; Yap, P.-T.; Bozoki, A.; and Liu, M. 2024. Federated learning for medical image analysis: A survey. *Pattern recognition*, 151: 110424.
- Horst, A.; Loustalot, P.; Yoganathan, S.; Li, T.; Xu, J.; Tong, W.; Schneider, D.; Löffler-Perez, N.; Di Renzo, E.; and Renaudin, M. 2025. Federated learning: a privacy-preserving approach to data-centric regulatory cooperation. *Frontiers in Drug Safety and Regulation*, 5: 1579922.
- Hsu, T.-M. H.; Qi, H.; and Brown, M. 2019. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*.
- Ji, W.; Yu, S.; Wu, J.; Ma, K.; Bian, C.; Bi, Q.; Li, J.; Liu, H.; Cheng, L.; and Zheng, Y. 2021. Learning calibrated medical image segmentation via multi-rater agreement modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12341–12351.
- Jiang, M.; Yang, H.; Cheng, C.; and Dou, Q. 2023. IOP-FL: Inside-outside personalization for federated medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(7): 2106–2117.
- Kaissis, G. A.; Makowski, M. R.; Rückert, D.; and Braren, R. F. 2020. Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6): 305–311.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. Scaffold: Stochastic

- controlled averaging for federated learning. In *International conference on machine learning*, 5132–5143. PMLR.
- Konečný, J.; McMahan, H. B.; Yu, F. X.; Richtárik, P.; Suresh, A. T.; and Bacon, D. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.
- Kwak, L.; and Bai, H. 2023. The role of federated learning models in medical imaging. *Radiology: Artificial Intelligence*, 5(3): e230136.
- Li, D.; and Wang, J. 2019. Fedmd: Heterogenous federated learning via model distillation. *arXiv preprint arXiv:1910.03581*.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2020a. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3): 50–60.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020b. Federated Optimization in Heterogeneous Networks. In Dhillon, I.; Papailiopoulos, D.; and Sze, V., eds., *Proceedings of Machine Learning and Systems*, volume 2, 429–450.
- Li, X.; Jiang, M.; Zhang, X.; Kamp, M.; and Dou, Q. 2021. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble distillation for robust model fusion in federated learning. *Advances in neural information processing systems*, 33: 2351–2363.
- 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1013–1023.
- Liu, S.; Zhang, R.; Fang, M.; Li, H.; Xun, T.; Wang, Z.; Shang, W.; Tian, J.; and Dong, D. 2025. PCRFed: personalized federated learning with contrastive representation for non-independently and identically distributed medical image segmentation. *Visual Computing for Industry, Biomedicine, and Art*, 8(1): 6.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Madni, H. A.; Umer, R. M.; and Foresti, G. L. 2023. Federated learning for data and model heterogeneity in medical imaging. In *International conference on image analysis and processing*, 167–178. Springer.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Menze, B. H.; Jakab, A.; Bauer, S.; Kalpathy-Cramer, J.; Farahani, K.; Kirby, J.; Burren, Y.; Porz, N.; Slotboom, J.; Wiest, R.; et al. 2014. The multi-modal brain tumor image segmentation benchmark (BRATS). *IEEE transactions on medical imaging*, 34(10): 1993–2024.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, 565–571. Ieee.
- Nam, D.; Panina, A.; and Pak, A. 2024. Lung cancer segmentation dataset with Lung-RADS class. *Mendeley Data*, 1.
- Oh, J.; Kim, S.; and Yun, S.-Y. 2021. Fedbabu: Towards enhanced representation for federated image classification. *arXiv preprint arXiv:2106.06042*.
- Pan, H.; Durak, G.; Zhang, Z.; Taktak, Y.; Koles, E.; Aktas, H. E.; Medetalibeyoglu, A.; Velichko, Y.; Spampinato, C.; Schoots, I.; et al. 2025. Adaptive Aggregation Weights for Federated Segmentation of Pancreas MRI. In *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Reddi, S.; Charles, Z.; Zaheer, M.; Garrett, Z.; Rush, K.; Konečný, J.; Kumar, S.; and McMahan, H. B. 2020. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Sattler, F.; Müller, K.-R.; and Samek, W. 2020. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE transactions on neural networks and learning systems*, 32(8): 3710–3722.
- T Dinh, C.; Tran, N.; and Nguyen, J. 2020. Personalized federated learning with moreau envelopes.

Advances in neural information processing systems, 33: 21394–21405.

Tan, A. Z.; Yu, H.; Cui, L.; and Yang, Q. 2022. Towards personalized federated learning. *IEEE transactions on neural networks and learning systems*, 34(12): 9587–9603.

Tan, Y. N.; Tinh, V. P.; Lam, P. D.; Nam, N. H.; and Khoa, T. A. 2023. A transfer learning approach to breast cancer classification in a federated learning framework. *IEEE Access*, 11: 27462–27476.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Voigt, P.; and Von dem Bussche, A. 2017. The eu general data protection regulation (gdpr). *A practical guide, 1st ed., Cham: Springer International Publishing*, 10(3152676): 10–5555.

Wang, J.; Jin, Y.; Stoyanov, D.; and Wang, L. 2023. Feddp: Dual personalization in federated medical image segmentation. *IEEE Transactions on Medical Imaging*, 43(1): 297–308.

Wang, J.; Jin, Y.; and Wang, L. 2022. Personalizing federated medical image segmentation via local calibration. In *European Conference on Computer Vision*, 456–472. Springer.

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 *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 599–610. Springer.

Zhao, X.; Liao, Y.; Xie, J.; He, X.; Zhang, S.; Wang, G.; Fang, J.; Lu, H.; and Yu, J. 2023. BreastDM: A DCE-MRI dataset for breast tumor image segmentation and classification. *Computers in Biology and Medicine*, 164: 107255.

Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; and Chandra, V. 2018. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.

Zhu, D.; and Tang, S. 2024. Adaptive Aggregate Federated Learning Image Segmentation Method Based on Euclidean Distance. In *2024 9th International Conference on Intelligent Computing and Signal Processing (ICSP)*, 612–615. IEEE.

A Related Work

Federated learning (FL) has garnered significant attention as a decentralized paradigm for collaboratively training models across distributed clients without sharing raw data (Horst et al., 2025; McMahan et al., 2017; Konečný et al., 2016; Li et al., 2020a). The foundational algorithm, FedAvg (Federated Averaging) proposed by (McMahan et al., 2017), serves as a baseline approach for federated training. However, recent works such as (Karimireddy et al., 2020; Collins et al., 2021; Li, He, and Song, 2021) show how the model convergence is effected in heterogeneous settings caused by local updates on non-IID data. FedAvg has been modified in (Hsu, Qi, and Brown, 2019; Reddi et al., 2020; Li et al., 2020a) to address the client drift issue, focused on learning a singular global model while also being generalizable enough to suit the needs of heterogeneous clients. Nonetheless, relying solely on a global model often fails to adequately capture the diverse and client-specific patterns inherent in highly non-IID data distributions (Zhao et al., 2018).

Personalized Federated Learning (PFL) addresses the limitations of global models by learning client-specific models tailored to local data distributions (Tan et al., 2022). Approaches such as (Arivazhagan et al., 2019) introduce client-specific layers, while others like (Collins et al., 2021) focus on learning a shared representation across heterogeneous clients. Parameter decoupling strategies have also been explored, including FedBN (Li et al., 2021), which personalizes batch normalization layers, and FedRep or FedBABU (Collins et al., 2021; Oh, Kim, and Yun, 2021), which personalize prediction layers. Despite these advancements, the use of cross-attention in PFL remains largely underexplored. Given its effectiveness in modeling long-range dependencies and complex feature interactions (Wang et al., 2023), attention based methods present a promising direction for enhancing both global representation sharing and local adaptation. While recent PFL approaches have shown progress in medical image segmentation (Chang et al., 2020; Dong and Voiculescu, 2021; Kaissis et al., 2020; Jiang et al., 2023; Wang, Jin, and Wang, 2022; Liu et al., 2021; Dai et al., 2024; Wang et al., 2023), few consider the benefits of leveraging cross-organ feature representations to improve local predictions.

Inconsistencies between predicted masks and ground truth offer a valuable signal for enhancing model supervision, especially for fine-grained boundary refinement. These subtle yet critical differences can impair segmentation performance, yet few existing methods leverage them for model calibration. (Wang, Jin, and Wang, 2022) propose a local calibration framework that incorporates inconsistency-aware supervision at both the feature and prediction levels. Their method combines contrastive site embeddings with channel selection and attention-based refinement to better capture anatomically ambiguous boundaries, albeit at increased inference cost. Similarly, (Wang et al., 2023) introduce an inconsistency-guided calibration strategy that computes inter-site prediction disagreements to inform additional supervision. While effective, this method incurs significant memory overhead by requiring access to multiple clients’ predictions or model weights.

B Additional Efficiency Analysis

B.1 Computational Overhead

Table 5 compares the computational complexity of FedOAP with FedDP. The total FLOPs capture the overall training cost aggregated across all clients, while the Attention Mechanism column isolates the computational load attributed to Local Query (LQ) introduced in Wang et al. (2023) and the proposed Decoupled Cross-Attention (DCA). Both mechanisms exhibit comparable FLOP requirements and scale linearly with the number of clients, indicating similar efficiency. Despite this equivalence, FedOAP achieves superior performance due to the incorporation of cross-attention, enabling richer feature exchange across distributions.

Table 5: Comparison of FLOP counts between FedOAP and FedDP across clients for the training stage.

Method	GFLOPs	
	Attention Mechanism	Total
FedDP	~ 8.73	~ 197.5
FedOAP (Ours)	~ 8.73	~ 232.6

The slightly higher total FLOP count observed for FedOAP arises from the inclusion of the spatial adapter ϕ_k , which is not present in FedDP. This adapter facilitates personalized adaptation to local data distributions

while remaining client-specific, thereby mitigating client drift during parameter aggregation (Arivazhagan et al., 2019; Collins et al., 2021). Overall, while FedOAP introduces a slight increase in total computational cost, this overhead is justified by its improved generalization and personalization efficiency across heterogeneous clients.

Table 6 presents a comparison of the memory requirements for boundary-guided edge refinement mechanisms within a single client. Wang et al. (2023) proposed the Inconsistency Guided Calibration (IGC) module, which enhances predictions by comparing the outputs of multiple client models on identical inputs to identify inconsistency regions. These inconsistencies are then used to refine each client’s prediction. In contrast, our proposed Perturbed Boundary Loss (PBL) introduces additive noise to the inconsistent regions between the prediction and the corresponding ground truth, thereby improving boundary precision without relying on inter-client information exchange.

Table 6: Comparison of memory overhead (in MB) for boundary-guided refinement mechanisms per client.

Method	Memory Required (MB)
IGC	~ 4.23
PBL (Ours)	~ 3.5

IGC exhibits substantially higher memory consumption compared to PBL. This increase arises because IGC requires loading and running inference using all client-specific models simultaneously on a single client to compute inter-client inconsistencies. Conversely, PBL only compares the local model’s output with its ground truth labels, which significantly reduces memory usage. Moreover, IGC is less effective in cross-organ scenarios where client data originate from distinct organ distributions, leading to excessively large inconsistency regions when comparing to predictions from other clients. By grounding the comparison within each client’s local domain, PBL achieves more stable boundary refinement while maintaining lower memory overhead.

B.2 Analysis of Training and Inference Overheads

Table 7: Comparison of FLOP counts between FedOAP and FedDP across clients for the training stage.

Method	Time Taken (s)		
	Federated Training	Local Fine-tuning	Inference
FedAvg	400.45	-	6.11
FedAvgM	401.32	-	6.21
FedAdagrad	393.42	-	6.29
FedPer	421.31	-	6.77
FedRep	622.95	115.41	7.19
FedDP	496.71	248.27	7.49
FedOAP (Ours)	621.77	176.5	8.16

Table 7 presents the total training, local fine-tuning and inference time for different federated learning methods for the experimental setting described in Section 3.3. FedOAP incurs higher training and inference times compared to other other learning methods. This could be attributed to having higher FLOP count as seen previously. The architecture of our method slightly increases the number of parameters *i.e.* having local spatial adapters. This can slightly affect the performance of our method. This trend is also noticeable when comparing the inference time per sample (tested with 1104 samples) as shown in Figure 4. However, FedOAP took considerably less time in the local fine-tuning step, ensuring efficient transfer of knowledge to a specific data distribution from a shared feature space.




Figure 4: Average inference time (in ms) per sample for each FL method.

C Experimental Parameters

C.1 Number of Clients

In our experimental setup, we set the number of clients to $\mathcal{K} = 3$, which differs from conventional federated learning (FL) configurations that typically involve a larger number of clients to simulate real-world network heterogeneity. However, in our study, each client represents an independent clinical institution or repository, where data is sourced from a distinct organ domain. Specifically, we utilize three datasets: BreastDM (Zhao et al., 2023), BraTS (Menze et al., 2014; Bakas et al., 2017, 2018), and LiTS (Bilic et al., 2023); corresponding to breast, brain, and liver tumor segmentation tasks, respectively. Assigning one dataset per client allows us to model an organ-distributed federated scenario, where data heterogeneity naturally arises from domain differences rather than patient-level variations. This design choice enables a controlled yet realistic evaluation of the proposed FedOAP framework’s ability to generalize across organ-specific feature distributions, while maintaining the interpretability of cross-organ knowledge aggregation.

C.2 Communication Rounds and Local Training




Figure 5: Comparison of FedOAP using $\mathcal{K} = 30$ in various training scenarios. Here, R denotes the number of communication rounds and F denotes the number of local fine-tuning epochs.

In our experiments, the total number of federated communication rounds is set to $\mathcal{T} = 5$, followed by $N_e = 2$ local fine-tuning epochs. This configuration was chosen based on the observed convergence behavior of FedOAP. Empirically, FedOAP demonstrates substantially faster convergence compared to other federated and personal-

ized federated learning baselines, achieving stable performance within only a few global rounds. The decoupled cross-attention mechanism in FedOAP facilitates more effective knowledge aggregation across clients, enabling the model to reach optimal performance without extensive communication. Furthermore, limiting the number of global rounds reduces the overall computational and communication overhead, which is desirable in realistic federated settings. The additional fine-tuning epochs allow each client to refine its personalized parameters on local data quickly, ensuring improved adaptation to organ-specific characteristics without incurring significant training cost.

To further assess the scalability and robustness of FedOAP, we extend our experiments to a larger federated setting with $\mathcal{K} = 30$ clients. The BreastDM, BraTS, and LiTS datasets are evenly partitioned into multiple subsets to accommodate the increased number of clients while maintaining balanced data distribution. We evaluate three distinct training paradigms, defined by $(\mathcal{T} = 5, N_\epsilon = 2)$, $(\mathcal{T} = 10, N_\epsilon = 4)$, and $(\mathcal{T} = 15, N_\epsilon = 6)$. The results, illustrated in Figure 5, clearly demonstrate that FedOAP continues to outperform all baseline methods across different configurations. Notably, FedOAP achieves faster convergence with fewer communication rounds, highlighting the effectiveness of its shared feature space and spatial adapter modules in mitigating client drift and improving generalization. In contrast, methods such as FedAvg exhibit increasing divergence with prolonged training, underscoring the stability and efficiency advantages of our proposed framework.

D Justification for Organ-Agnostic Claim

The term organ-agnostic in FedOAP reflects the framework’s ability to generalize and adapt across diverse organ domains relying on minimal organ-specific training. This characteristic is seen in Table 2 through the model’s superior zero-shot performance, achieving a 6% improvement over FedDP under cross-organ evaluation. Such improvement demonstrates FedOAP’s enhanced transferability and capacity to learn a shared, domain-invariant feature representation that generalizes effectively to unseen organ distributions. Furthermore, when fine-tuned locally for only two additional epochs, FedOAP exhibits a substantial performance boost, surpassing all competing methods by a notable margin. This rapid adaptation underscores the strength of its transferable representations and the efficacy of its decoupled cross-attention and spatial adaptation mechanisms. Hence, the framework’s consistent cross-domain generalization and efficient personalization together justify its designation as an organ-agnostic federated learning paradigm.

E Model

To implement our proposed method and test all other FL and PFL methods, we employ a modified U-Net backbone (Ronneberger, Fischer, and Brox, 2015) augmented with two key components: a Cross-Attention module and a Spatial Adapter. This architecture is designed to facilitate cross-organ feature interaction and domain-adaptive refinement while maintaining computational efficiency. The base architecture follows the conventional U-Net encoder-decoder structure, comprising symmetric downsampling and upsampling paths with skip connections to preserve spatial information.

The encoder begins with an initial double convolution block comprising two sequential 3×3 convolutional layers, each followed by instance normalization and ReLU activation. Feature channels are progressively expanded through three Down blocks using 2×2 max pooling and subsequent double convolutions, with channel dimensions increasing from 64 to 512. The bottleneck consists of two parallel 3×3 convolutional stages with 1024 channels: one serving as the query branch and the other as the key-value branch for cross-attention. The cross-attention module performs spatial feature alignment between the query and key-value representations using 1×1 convolutions to generate queries, keys, and values, followed by multi-head attention and normalization via group normalization. This module also allows cross-client feature exchange, enhancing generalization across organ domains.

The decoder mirrors the encoder through four Up blocks, each using a 2×2 transposed convolution for upsampling, concatenation with the corresponding encoder feature map, and a subsequent double convolution to refine spatial details. Channel dimensions are reduced progressively from 1024 back to 64. At the final stage, a spatial adapter module composed of two 3×3 convolutional layers with residual connections—further refines spatial consistency and preserves localized texture information. A final 1×1 convolution projects the 64-channel feature map to the desired segmentation mask.

F Dataset Preparation

We have already provided a description of the datasets used in our experimentation in Section 3.1. As we are dealing with data from multiple input modalities, we unify them by converting all ground truth to a single channel 2 dimensional input. 20% of the data from each dataset is used for testing. 20% of the remaining data is used for validation during training while the rest is kept for training.

G Segmentation Loss

We already discussed the perturbed loss $\mathcal{L}_p$ in Section 2.4. Our composite objective function additionally depends on a segmentation loss $\mathcal{L}_s$. To ensure accurate delineation of organ boundaries during segmentation, we employ a loss function that combines the strengths of Binary Cross-Entropy (BCE) and the Dice coefficient. The BCE component penalizes pixel-wise misclassifications, promoting accurate foreground-background discrimination, while the Dice component directly optimizes the overlap between predicted and ground truth masks, enhancing spatial consistency in regions of fine structure.

Formally, given a predicted probability map $\hat{y}_{i,k}$ and its corresponding ground truth mask $y_{i,k}$, the segmentation loss $\mathcal{L}_s$ is defined as a weighted combination of the two terms:

$$\mathcal{L}_s = 0.5 \times \mathcal{L}_{BCE}(\hat{y}_{i,k}, y_{i,k}) + \mathcal{L}_{Dice}(\hat{y}_{i,k}, y_{i,k}) \quad (8)$$

where the BCE loss operates on the logits, and the Dice term measures the degree of spatial overlap between prediction and ground truth. This joint formulation enables stable optimization across varying organ sizes and image resolutions while maintaining sensitivity to small structures.