

# When to Adapt? Adapting the Model or Data in Federated Medical Imaging

Chamani Shiranthika<sup>[0000–0002–1608–8196]</sup> and Parvaneh  
Saeedi<sup>[0000–0002–7507–9986]</sup>

School of Engineering Science, Simon Fraser University, Burnaby, BC, V5A 1S6  
(csj5, psaeedi)@sfu.ca

**Abstract.** Federated learning enables collaborative model training across medical institutions without sharing raw data, but its performance is often limited by domain heterogeneity across clients. Existing approaches to address this challenge fall into two main paradigms: model-side personalization, which adapts model parameters to each client, and data-side harmonization, which reduces inter-client variation at the input level. Despite their widespread use, these strategies have not been systematically compared. In this work, we conduct a comprehensive study across six medical imaging settings—colon polyp, skin lesion, and breast tumor segmentation, and tuberculosis CXR, brain tumor, and breast tumor classification—covering diverse types of domain shift. We evaluate a broad set of state-of-the-art harmonization and personalization methods under a unified framework. Our results reveal a conditional trade-off driven by the nature of heterogeneity: harmonization is more effective when variation is primarily appearance-based (e.g., CXR classification), while personalization performs better when differences are structural (e.g., colon polyp segmentation). When inter-client variation is limited, both strategies perform similarly. These findings demonstrate that the effectiveness of adaptation in federated medical imaging depends on the type and magnitude of domain shift rather than the strategy alone. We provide practical guidelines for selecting between harmonization and personalization, and highlight directions for future hybrid approaches that combine both paradigms. Code is available at: <https://github.com/ChamaniS/WhenToAdapt>.

**Keywords:** data harmonization · model personalization · federated learning

## 1 Introduction

Federated learning (FL) [1] has emerged as a key paradigm for privacy-preserving medical imaging collaborations, enabling multiple institutions to jointly train models without sharing patients’ data. However, its effectiveness is often limited by inter-client heterogeneity, differences in data acquisition protocols, patient

demographics, and annotation practices. Such non-independent and identically distributed (IID) data distributions cause severe model bias and poor cross-site generalization. Therefore, managing data heterogeneity has become one of the most critical open challenges in federated medical learning [2], [3].

Two broad research directions have evolved to address such heterogeneity challenge. (1): *Model-side personalization* focuses on adapting the model to each client through modularized layers, local fine-tuning, or adaptive optimization. Early personalized federated learning methods rely on local fine-tuning or partial model sharing (FedPer [4], FedRep [5], FedBN [6]), and on meta-learning or attention mechanisms (FedMeta [7], pFedMe [8], Per-FedAvg [9], FedAMP [10]). Some research treats clients as related tasks (MOCHA [11], VIRTUAL [12]), clusters of similar clients (IFCA [13], FedCluster [14], FedGroup [15]), or uses knowledge distillation (FedMD [16], FedDF [17]). Some other widely used methods include FedAdam [18], FedProx [19], SCAFFOLD [20], pFedMe [8], FedMeta [7], FedAMP [10], and Ditto [21]. (2): *Data-side harmonization* reduces appearance discrepancies caused by scanner hardware, acquisition protocols, or imaging settings while preserving clinically relevant anatomy and is commonly studied as protocol-level preprocessing (resampling, intensity/colour normalization) or image-level harmonization using learned cross-domain mappings [22]. State-of-the-art (SoTA) data-side harmonization methods include GANs and related generative models (CycleGAN [23], CUT [24], StarGAN, DeepHarmony [25], ImUnity [26], BlindHarmony [27]), attention/transformer-based approaches (AdaIN [28], CoMoGAN [29], QS-Attn [30], Harmonization Transformer [31]), and newer diffusion-based models, which align input distributions across sites while maintaining diagnostic content.

Despite substantial progress, prior work has largely studied personalization and harmonization in isolation, without analyzing when each is preferable under different domain shifts. We address this gap with a systematic comparison across six medical imaging tasks: colon polyp, skin lesion, and breast tumor segmentation, and tuberculosis CXR, brain tumor, and breast tumor classification. These tasks span diverse heterogeneity types, from strong structural variation (e.g., colon polyp segmentation) to predominantly appearance-driven shifts (e.g., tuberculosis CXR classification), as well as moderate or mixed variations (e.g., skin, brain, and breast tumor tasks).

Using a unified experimental protocol, we evaluate 18 SoTA harmonization and personalization methods under identical federated settings. Our results reveal a conditional trade-off: harmonization is more effective for appearance-driven variation, while personalization performs better under strong structural heterogeneity; when variation is limited or mixed, both approaches perform similarly. We contribute: (1) a comprehensive comparison of harmonization and personalization in federated medical imaging, (2) an empirical analysis of how the type and magnitude of domain shift affect their effectiveness, and (3) a modular code-base supporting both paradigms. Section 2 describes the methodology, Section 3 presents the results, and Section 4 concludes the paper.

## 2 Methodology

### 2.1 Problem Formulation

We consider a cross-silo FL setting with  $k$  medical institutions (clients), each holding a local dataset:  $D_k = \{(x_i^{(k)}, y_i^{(k)})\}_{i=1}^{n_k}$  drawn from an unknown distribution  $\mathcal{P}_k$  over images  $x$  and labels  $y$ . A standard FL objective learns a single global model  $f_\theta$  by minimizing the weighted empirical risk:

$$\min_{\theta} \sum_{k=1}^K p_k \mathbb{E}_{(x,y) \sim \mathcal{P}_k} [\ell(f_\theta(x), y)], \quad (1)$$

where  $p_k = n_k / \sum_j n_j$  and  $\ell$  is a task-appropriate loss. In practice, the client distributions  $\{\mathcal{P}_k\}$  are heterogeneous. In medical imaging collaborations, this heterogeneity arises from two distinct, qualitatively different sources:

1. **Input-level/style-level shift**  $\Delta_{\text{style}}$ : variations in scanner hardware, acquisition protocols, intensity scaling, contrast, noise level, padding, and background artifacts. These changes alter the global appearance statistics while largely preserving the anatomical layout and pathology of the data.

2. **Content-level/structural shift**  $\Delta_{\text{content}}$ : variations in anatomy, pathology morphology and prevalence, viewpoint, endoscope/camera pose, and scene context. These differences alter the spatial structure and semantic content that the model must segment or classify.

Let  $\Delta_k = (\Delta_{\text{style},k}, \Delta_{\text{content},k})$  denote the domain shift of client  $k$ . A single global model  $f_\theta$  may underperform when  $\Delta_k$  is large and heterogeneous across clients.

We empirically compare these two families of heterogeneity acting differently on:

- **Data harmonization**: learn or apply client-specific transformations  $T_k$  in image or feature space such that the transformed data  $T_k(x)$  are statistically aligned to a reference domain, and then train a *shared* model  $f_\theta$  on harmonized data, and
- **Model personalization**: directly adapt the *shared* model  $f_\theta$  per client, yielding client-specific parameters  $\theta_k$ , while leaving the raw images unchanged.

Formally, a harmonization-based system learns

$$\min_{\theta, \{T_k\}} \sum_{k=1}^K p_k \mathbb{E}_{(x,y) \sim \mathcal{P}_k} [\ell(f_\theta(T_k(x)), y)] \quad (2)$$

A personalization-based system instead learns client-specific models

$$\min_{\{\theta_k\}} \sum_{k=1}^K p_k \mathbb{E}_{(x,y) \sim \mathcal{P}_k} [\ell(f_{\theta_k}(x), y)] \quad (3)$$

Accordingly, our research question is:

*Under which types of domain shift does it pay to adapt the **data** (via harmonization), and under which types does it pay to adapt the **model** (via personalization)?*

To answer this, we design two federated medical imaging benchmarks that are deliberately dominated by different heterogeneity types: (i) colon polyp segmentation with large structural/contextual variation across four endoscopy datasets and (ii) Tuberculosis CXR classification with largely preserved anatomy but strong scanner/intensity variation across four CXR datasets.

## 2.2 Model Personalization and data harmonization approaches employed

We consider representative model personalization and data harmonization strategies commonly used in federated learning.

Model personalization methods adapt the global model to client-specific distributions. We group them into three categories: (1) layer-wise personalization, which shares a common backbone while keeping client-specific components local (e.g., FedPer [4], FedRep [5], FedBN [6]); (2) optimization-based approaches, which modify the training objective or update rules to handle heterogeneity (e.g., FedProx [19], SCAFFOLD [20], FedAdam [18]); and (3) bi-model and meta-learning approaches, which explicitly learn personalized client models alongside a global model (e.g., pFedMe [32], Ditto [21]).

Data harmonization methods aim to reduce inter-client appearance variation while preserving underlying structure. We consider: (1) basic preprocessing and augmentation to standardize resolution and intensity; (2) histogram-based methods that align intensity distributions across clients; (3) frequency-domain adaptation (FDA) to transfer low-frequency style information; (4) MixStyle-based approaches that interpolate feature or input statistics to simulate domain shifts; and (5) generative translation methods (e.g., CycleGAN [33], CUT [24], AdaIN [28], CoMoGAN [29]) that learn mappings between domains.

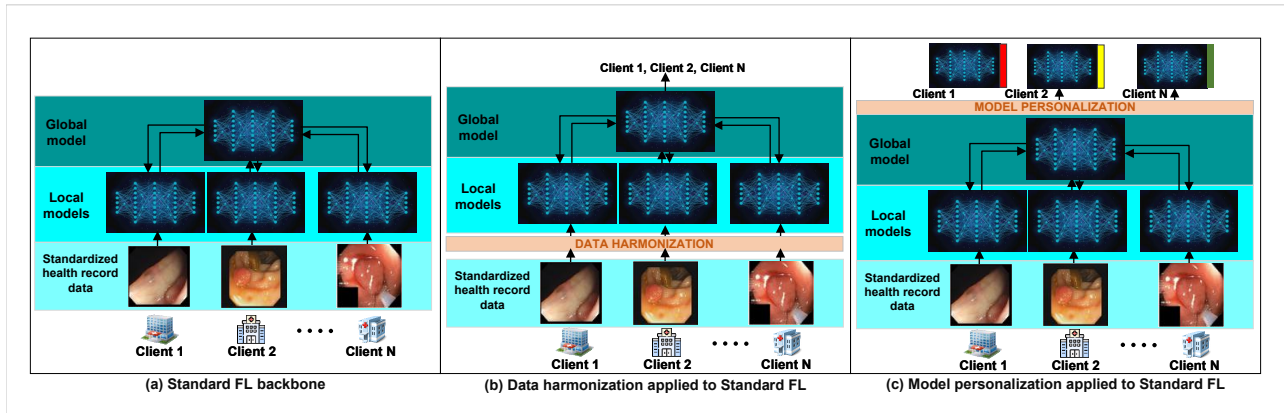


Fig. 1: Experimental methodology of the study.

Table 1: Dataset composition across clients for the 6 tasks. Each client corresponds to a dataset with its sample size and fixed test split. Remaining data is split into 85% training and 15% validation.

(a) Segmentation tasks					(b) Classification tasks				
C#	Dataset	Modality	Total #	Test #	C#	Dataset	Modality	Total #	Test #
<i>Task 1: Colon polyp segmentation</i>					<i>Task 4: Tuberculosis CXR classification</i>				
C1	Kvasir-SEG [35]	Endoscopy	1000	100	C1	Shenzhen Hospital CXR [47]	X-ray	662	65
C2	ETIS-Larib [36]	Endoscopy	196	19	C2	Montgomery County CXR [47]	X-ray	138	13
C3	CVC-ColonDB [37]	Endoscopy	380	38	C3	TBX11K [48]	X-ray	1600	160
C4	CVC-ClinicDB [38]	Endoscopy	612	60	C4	Pakistan (local)	X-ray	3008	300
<i>Task 2: Skin lesion segmentation</i>					<i>Task 5: Brain tumor classification</i>				
C1	HAM10K [39]	Dermoscopy	10015	1001	C1	Sartajbhuvaji (Kaggle) [49]	MRI	3160	314
C2	PH2 [40]	Dermoscopy	200	20	C2	RM1000 (Kaggle) [50]	MRI	7023	701
C3	ISIC2017 [41]	Dermoscopy	2000	200	C3	thomasdubail (Kaggle) [51]	MRI	3096	308
C4	ISIC2018 [42]	Dermoscopy	2594	259	C4	Figshare [52]	MRI	7200	720
<i>Task 3: Breast tumor segmentation</i>					<i>Task 6: Breast tumor classification</i>				
C1	BUSBRA [43]	Ultrasound	1875	81	C1	BUSBRA [43]	Ultrasound	1875	186
C2	BUS_UC [44]	Ultrasound	811	187	C2	BUS_UC [44]	Ultrasound	811	80
C3	BUSI [45]	Ultrasound	437	43	C3	BUSI [45]	Ultrasound	647	64
C4	UDIAT [46]	Ultrasound	163	16	C4	UDIAT [46]	Ultrasound	163	15

### 2.3 Comparative Framework

As shown in Fig. 1, we instantiate model personalization and data harmonization strategies on top of a standard FL backbone (Flower [34]-based cross-silo FL with four clients per task). Each method uses the same base architecture for each task: a U-Net for colon polyp, skin lesion, and breast tumor segmentation; a DenseNet-based classifier for tuberculosis CXR; and an EfficientNet-based classifier for brain and breast tumor classification. All methods are trained with the same schedule and communication budget, ensuring that performance differences arise solely from the applied personalization or harmonization strategy.

*Tasks.* We design six FL settings using medical imaging tasks shown in Table 1:

*Baselines.* We include five non-personalized and non-harmonized baselines:

- **Locally-centralized:** each client trains its model solely on local data.
- **Centralized-globally tested:** a central model is trained on all clients’ data and tested using all clients’ pooled data.
- **Centralized-locally tested:** a central model is trained on all clients’ data and tested using clients’ local data separately in their local spaces.
- **Vanilla FedAvg-globally tested:** a single global model trained by standard FedAvg and tested using all clients’ pooled data.
- **Vanilla FedAvg-locally tested:** a single global model trained by standard FedAvg and tested using clients’ local data separately in their local spaces.

## 3 Experimental Results

Table 2 shows one representative colon polyp sample per dataset after harmonization. Harmonized samples of other datasets are included in our code repository and are omitted here for space. We include the amplified difference, defined

as the absolute per-pixel RGB difference between the original and harmonized images, scaled to highlight subtle color and contrast changes. Segmentation tasks use Dice as the primary metric (with IoU, pixel accuracy, precision, recall, and specificity), while classification tasks use Cohen’s  $\kappa$  (with accuracy, precision, recall, F1, and specificity). We compare harmonization and personalization methods and analyze trends by heterogeneity type (style vs. content). Results are reported in Tables 3 and 4.

Table 2: Examples of harmonized samples and amplified differences for Colon polyp (top), and skin lesion (bottom) datasets used for segmentation. H. = harmonized image; D. = amplified difference; Aug. = augmentation; Hist.mat. = histogram matching; ARI = average representative image; SRI = single representative image; FDA = Fourier domain adaptation.

Image	Aug.	Hist. mat. (SRI)	Hist. mat. (ARI)	FDA	Input level mixstyle	Feature level mixstyle	CycleGAN	CUT	AdaIN	CoMoGAN	
Colon polyp datasets											
C1	H.										
	D.										
C2	H.										
	D.										
C3	H.										
	D.										
C4	H.										
	D.										

For colon polyp segmentation, where inter-client variation is large and primarily structural (e.g., shape and geometry), personalization methods consistently outperform harmonization, highlighting the need to adapt models to client-specific structures. In contrast, for skin lesion segmentation, where variation is more limited and largely appearance-driven, harmonization methods (e.g., MixStyle, SRI, ComoGAN) perform comparably to personalization. Breast tumor segmentation follows a similar pattern, with both strategies achieving competitive performance, suggesting that when structural differences are moderate, neither approach has a clear advantage.

A similar trend is observed in classification tasks. For brain tumor classification, where inter-site variation is minimal, all methods reach near-ceiling performance,

Table 3: Data harmonization vs. model personalization across segmentation tasks. ARI = average representative image, SRI = single representative image. Best results are in **bold**.

Method	Colon Polyp				Skin Lesion				Breast Tumor			
	C1	C2	C3	C4	C1	C2	C3	C4	C1	C2	C3	C4
<i>Baseline methods</i>												
Locally centralized	0.8297	0.7804	0.8857	0.8500	0.8920	0.8073	0.7692	0.8071	0.8367	0.9135	0.8049	0.7325
Centralized – globally tested		0.8499				0.9250				0.9008		
Centralized – locally tested	0.8024	0.7892	0.9121	0.8428	0.9178	0.8928	0.9374	0.9434	0.8998	0.9307	0.8419	0.9153
Vanilla FedAvg – globally tested		0.7286				0.9355				0.8933		
Vanilla FedAvg – locally tested	0.7559	0.4430	0.7508	0.8258	0.9283	0.9009	0.9448	0.9476	0.8964	0.9316	0.8027	0.9129
<i>Harmonization methods</i>												
Augmentations	0.6698	0.3014	0.4714	0.6804	0.9217	0.8800	0.9322	0.9394	0.9351	0.9178	0.7424	0.9220
Histogram matching (ARI)	0.7149	0.3901	0.5892	0.7144	0.9153	0.8909	0.9322	0.9379	0.9145	0.8936	0.7997	0.8875
Histogram matching (SRI)	0.7749	0.6626	0.7336	0.7733	0.9141	0.9162	0.9307	0.9374	0.9193	0.8989	0.7747	0.8937
FDA (SRI)	0.7498	0.3888	0.6928	0.7233	0.8940	0.8465	0.9176	0.9216	0.8592	0.8807	0.7842	0.8477
Input level MixStyle	0.8296	0.6821	0.8270	0.8428	0.9114	0.8672	0.9324	0.9374	0.9173	0.8990	0.8073	0.9126
Feature level MixStyle	0.8336	0.6418	0.8273	0.8324	0.9237	0.8894	0.9375	0.9428	0.9237	0.8992	0.8022	0.9026
Cycle-GAN	0.8129	0.5506	0.6435	0.6685	0.9310	0.8926	0.9176	0.9355	0.8656	0.9174	0.7836	0.8390
CUT	0.6918	0.6069	0.8222	0.7056	0.9078	0.5536	0.4683	0.7853	0.5758	0.8600	0.3547	0.0421
AdaIN	0.7062	0.6282	0.8676	0.7679	0.7666	0.6694	0.7564	0.7484	0.9163	0.9025	0.8156	0.9062
ComoGAN	0.7774	0.5653	0.8498	0.8418	0.9204	0.9086	0.9362	0.9420	0.9185	0.9033	0.8014	0.9273
<i>Personalization methods</i>												
Local Finetuning	0.8329	0.6914	0.8678	0.8311	0.9290	0.9100	0.9513	0.9501	0.9195	0.9087	0.8014	0.8875
FedPer	0.8085	0.4808	0.8212	0.8633	0.9206	0.9239	0.9427	0.9449	0.9233	0.9178	0.8270	0.8935
FedRep	0.8141	0.5503	0.8297	0.8775	0.9179	0.8984	0.9389	0.9423	0.9345	0.9073	0.8118	0.9093
FedBN	0.8200	0.6550	0.8287	0.8663	0.9229	0.8973	0.9366	0.9418	0.9341	0.9116	0.7887	0.9032
FedProx	0.8248	0.6829	0.7959	0.8173	0.8780	0.9307	0.9334	0.9156	0.9270	0.8955	0.8021	0.9138
FedAdam	<b>0.8569</b>	<b>0.8332</b>	<b>0.9391</b>	<b>0.8906</b>	0.9287	0.9124	0.9514	0.9511	0.9247	0.8975	0.8124	0.9043
SCAFFOLD	0.7368	0.3799	0.7083	0.7025	0.8715	0.8333	0.8940	0.9055	0.7006	0.7429	0.5701	0.5576
pFedMe	0.4564	0.1201	0.2354	0.4054	0.7233	0.8516	0.7843	0.7732	0.4158	0.1823	0.1654	0.1621
Ditto	<b>0.8769</b>	<b>1.0000</b>	<b>1.0000</b>	<b>1.0000</b>	0.9210	0.9232	0.9343	0.9416	0.9338	0.9026	0.8420	0.9089

and neither harmonization nor personalization provides a clear benefit. Breast tumor classification shows comparable behavior, with both approaches yielding similar results across clients. In contrast, for tuberculosis CXR classification, where variation is dominated by appearance differences, harmonization consistently outperforms personalization by reducing input variability. Overall, when inter-client differences are small or moderate, both strategies perform similarly regardless of task. When differences are large, the dominant type of variation determines the more effective approach: personalization is better suited for structurally heterogeneous segmentation tasks, whereas harmonization is more effective for appearance-driven classification tasks.

## 4 Conclusions & Future works

The results reveal a conditional trade-off rather than a strictly task-dependent pattern. When inter-client differences are large and primarily structural, as in colon polyp segmentation, personalization outperforms harmonization by adapting to client-specific geometry and context. In contrast, when variation is dominated by appearance and intensity differences, as in tuberculosis CXR classification, harmonization is more effective by reducing input variability and supporting a shared global model. However, this trend does not hold universally.

Table 4: Data harmonization vs. model personalization across classification tasks. ARI = average representative image, SRI = single representative image. Best results are in **bold**.

Method	Tuberculosis CXR				Brain Tumor				Breast Tumor			
	C1	C2	C3	C4	C1	C2	C3	C4	C1	C2	C3	C4
<i>Baseline methods</i>												
Locally centralized	0.8462	0.6486	0.9750	0.9883	0.9563	0.9866	0.9690	0.9722	0.6353	0.6770	0.7764	0.7000
Centralized – globally tested		0.9656				0.9941				0.5872		
Centralized – locally tested	0.8768	0.6486	0.9875	0.9881	1.0000	0.9943	0.9911	0.9926	0.6398	0.4413	0.6513	0.5263
Vanilla FedAvg – globally tested		0.7763				0.9974				0.5466		
Vanilla FedAvg – locally tested	0.0000	0.0000	0.8500	0.9883	1.0000	0.9962	0.9956	0.9981	0.5645	0.3976	0.6074	0.6667
<i>Harmonization methods</i>												
Augmentations	0.0901	0.0679	0.4631	0.6385	0.9956	0.9962	0.9956	0.9944	0.7251	0.4158	0.5320	0.6667
Histogram matching (ARI)	0.4937	0.4990	0.8625	1.0000	0.9956	0.9962	0.9823	0.9944	0.7451	0.4483	0.5412	0.6667
Histogram matching (SRI)	0.2536	0.3158	0.9000	1.0000	0.9956	0.9962	0.9603	0.9926	0.6385	0.3750	0.5412	0.6667
FDA (SRI)	0.1317	0.0000	0.8875	0.9542	0.9782	0.9790	0.9294	0.9778	0.3329	0.1832	0.3621	0.6667
Input level MixStyle	0.1898	0.0000	0.7750	1.0000	1.0000	1.0000	0.9734	0.9981	0.6596	0.3750	0.5218	0.6667
Feature level MixStyle	0.0634	0.0000	0.9250	1.0000	1.0000	1.0000	0.9823	0.9963	0.7359	0.3415	0.6177	0.6667
Cycle-GAN	<b>0.5317</b>	<b>0.8158</b>	<b>0.7750</b>	<b>0.9883</b>	0.9825	0.9867	0.9735	0.9648	0.6750	0.3077	0.5053	0.1429
CUT	0.0000	0.0000	0.7000	0.9883	1.0000	0.9962	0.9647	0.9926	0.6494	0.3333	0.2101	0.4706
AdaIN	<b>0.5534</b>	<b>0.8434</b>	<b>0.8625</b>	<b>0.8606</b>	0.9956	0.9962	0.9425	0.9889	0.4615	0.3671	0.3081	0.4706
ComGAN	0.0000	0.0000	0.5250	1.0000	1.0000	0.9962	0.9734	0.9944	0.3654	0.2751	0.3932	0.7000
<i>Personalization methods</i>												
Local Finetuning	0.0329	0.0914	0.8678	0.8217	0.9854	0.9760	0.9614	0.9340	0.5874	0.4476	0.5249	0.5682
FedPer	0.0634	0.0000	0.9125	0.9883	0.9956	0.9943	0.9911	0.9963	0.5994	0.4551	0.6074	0.6667
FedRep	0.0959	0.2041	0.9500	0.9883	0.9956	0.9943	0.9956	0.9926	0.6972	0.5280	0.6275	0.4000
FedBN	0.0000	0.0000	0.0000	0.0000	0.9826	0.9522	0.9297	0.8796	0.4639	0.5190	0.6456	0.4545
FedProx	0.3223	0.0990	0.9375	0.9883	0.9956	0.9981	0.9735	0.9963	0.3871	0.2298	0.4188	0.5263
FedAdam	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
SCAFFOLD	0.0000	0.0000	0.0000	0.0000	0.9869	0.9790	0.9691	0.9778	0.0000	0.0000	0.0000	0.0000
pFedMe	0.2547	0.1420	0.2847	0.4875	0.0890	0.0890	0.0305	0.1130	0.0000	0.0000	0.0000	0.0000
Ditto	0.7841	0.6486	1.0000	0.9268	0.9956	0.9943	0.9956	0.9926	0.6446	0.3191	0.5849	0.6667

When inter-client differences are small or moderate, both strategies perform similarly, as observed in skin lesion and breast tumor segmentation, as well as brain and breast tumor classification, where neither approach offers a clear advantage. Overall, the effectiveness of harmonization versus personalization is governed by the type and magnitude of dataset heterogeneity rather than the task alone. In practice, federated medical imaging systems should assess cross-site variation before selecting a strategy, rather than assuming a fixed preference. A key limitation is the assumption that inter-client heterogeneity is known a priori, whereas in practice it is unknown and dynamic. Future work should develop metrics to characterize heterogeneity directly from data and enable adaptive frameworks that combine harmonization and personalization at multiple levels. Extending this analysis to non-stationary settings and incorporating causal modeling are important directions. Standardized benchmarks and evaluation of performance–privacy–communication trade-offs are also needed.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., Aguera y Arcas, B.: Communication-Efficient Learning of Deep Networks from Decentralized Data. In: Proc. of AISTATS, pp.

- 1273–1282. PMLR, Fort Lauderdale, FL, USA (2017)
2. Shiranthika, C., Saeedi, P., Bajić, I.V.: Decentralized Learning in Healthcare: A Review of Emerging Techniques. *IEEE Access* **11**, 54188–54209 (2023). <https://doi.org/10.1109/ACCESS.2023.3281832>
  3. Shiranthika, C., Hadizadeh, H., Saeedi, P., Bajić, I.V.: Adaptive Asynchronous Split Federated Learning for Medical Image Segmentation. *IEEE Access* **12**, 182496–182515 (2024). <https://doi.org/10.1109/ACCESS.2024.3511430>
  4. Arivazhagan, M.G., Aggarwal, V., Singh, A.K., Choudhary, S.: Federated Learning with Personalization Layers. arXiv preprint arXiv:1912.00818 (2019)
  5. Collins, L., Hassani, H., Mokhtari, A., Shakkottai, S.: Exploiting Shared Representations for Personalized Federated Learning. In: Proc. of ICML, Proc. Mach. Learn. Res., vol. 139, pp. 2089–2099. PMLR (2021)
  6. Li, X., Jiang, M., Zhang, X., Kamp, M., Dou, Q.: FedBN: Federated Learning on Non-IID Features via Local Batch Normalization. arXiv preprint arXiv:2102.07623 (2021)
  7. Chen, F., Luo, M., Dong, Z., Li, Z., He, X.: Federated Meta-Learning with Fast Convergence and Efficient Communication. arXiv preprint arXiv:1802.07876 (2018)
  8. Dinh, C.T., Tran, N., Nguyen, J.: Personalized Federated Learning with Moreau Envelopes. In: Proc. of NeurIPS, 33, 21394–21405 (2020)
  9. Fallah, A., Mokhtari, A., Ozdaglar, A.: Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach. In: Proc. of NeurIPS, 33, 3557–3568 (2020)
  10. Huang, Y., Chu, L., Zhou, Z., Wang, L., Liu, J., Pei, J., Zhang, Y.: Personalized Cross-Silo Federated Learning on Non-IID Data. In: Proc. of AAAI 35(9), 7865–7873 (2021)
  11. Smith, V., Chiang, C.-K., Sanjabi, M., Talwalkar, A.S.: Federated Multi-Task Learning. In: Proc. of NeurIPS, 30 (2017)
  12. Corinzia, L., Beuret, A., Buhmann, J.M.: Variational Federated Multi-Task Learning. arXiv preprint arXiv:1906.06268 (2019)
  13. Ghosh, A., Chung, J., Yin, D., Ramchandran, K.: An Efficient Framework for Clustered Federated Learning. In: Proc. of NeurIPS, 33, 19586–19597 (2020)
  14. Chen, C., Chen, Z., Zhou, Y., Kailkhura, B.: FedCluster: Boosting the Convergence of Federated Learning via Cluster-Cycling. *IEEE BigData 2020*, pp. 5017–5026. IEEE (2020)
  15. Duan, M., Liu, D., Ji, X., Liu, R., Liang, L., Chen, X., Tan, Y.: FedGroup: Efficient Federated Learning via Decomposed Similarity-Based Clustering. In: 2021 IEEE International Conference on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCLOUD/SocialCom/SustainCom), pp. 228–237. IEEE (2021)
  16. Li, D., Wang, J.: FedMD: Heterogeneous Federated Learning via Model Distillation. arXiv preprint arXiv:1910.03581 (2019)
  17. Zhang, H., Zhu, Y., Wu, T., Cheng, S., Liu, J.: Model-Heterogeneous Federated Learning With Bidirectional Knowledge Distillation. *IEEE Trans. Mob. Comput.* (2025). <https://doi.org/10.1109/TMC.2025.3599315>
  18. Reddi, S., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., McMahan, H.B.: Adaptive Federated Optimization. In: Proc. of ICLR (2021)
  19. Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., Smith, V.: Federated Optimization in Heterogeneous Networks. In: Proc. of MLSys 2, 429–450 (2020)
  20. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S., Stich, S., Suresh, A.T.: SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In: Proc. of the 37th International Conference on Machine Learning (ICML 2020), Proc. Mach. Learn. Res., vol. 119, pp. 5132–5143. PMLR (2020)
  21. Li, T., Hu, S., Beirami, A., Smith, V.: Ditto: Fair and Robust Federated Learning Through Personalization. In: Proc. of ICML 2021, Proc. Mach. Learn. Res., vol. 139, pp. 6357–6368. PMLR (2021)
  22. Mali, S.A., Ibrahim, A., Woodruff, H.C., Andrearczyk, V., Müller, H., Primakov, S., Salahuddin, Z., Chatterjee, A., Lambin, P.: Making radiomics more reproducible across scanner and imaging protocol variations: A review of harmonization methods. *J. Pers. Med.* **11**(9), 842 (2021)
  23. Liu, M., Zhu, A.H., Maiti, P., Thomopoulos, S.I., Gadewar, S., Chai, Y., Kim, H., Jahanshad, N., Alzheimer’s Disease Neuroimaging Initiative: Style transfer generative adversarial networks to harmonize multisite MRI to a single reference image to avoid overcorrection. *Hum. Brain Mapp.* **44**(14), 4875–4892 (2023)
  24. Park, T., Efros, A.A., Zhang, R., Zhu, J.-Y.: Contrastive Learning for Unpaired Image-to-Image Translation. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) *Computer Vision – ECCV 2020*, LNCS, vol. 12354, pp. 319–345. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58545-7\\_19](https://doi.org/10.1007/978-3-030-58545-7_19)
  25. Dewey, B.E., Zhao, C., Reinhold, J.C., Carass, A., Fitzgerald, K.C., Sotirchos, E.S., et al.: DeepHarmony: A deep learning approach to contrast harmonization across scanner changes. *Magn. Reson. Imaging*, **64**, 160–170 (2019). <https://doi.org/10.1016/j.mri.2019.05.041>

26. Cackowski, S., Barbier, E.L., Dojat, M., Christen, T.: Imunity: A generalizable VAE-GAN solution for multicenter MR image harmonization. *MIA*, 88, 102799 (2023). <https://doi.org/10.1016/j.media.2023.102799>
27. Jeong, H., Byun, H., Kang, D.U., Lee, J.: BlindHarmony: "Blind" Harmonization for MR Images via Flow Model. In: Proc. of ICCV, pp. 21129–21139.(2023)
28. Huang, X., Belongie, S.: Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. In: Proc. of ICCV, pp. 1501–1510. IEEE (2017)
29. Pizzati, F., Cerri, P., De Charette, F.: CoMoGAN: Continuous Model-Guided Image-to-Image Translation. In: Proc. of CVPR, pp. 14288–14298. IEEE/CVF (2021)
30. Hu, X., Zhou, X., Huang, Q., Shi, Z., Sun, L., Li, Q.: QS-Attn: Query-Selected Attention for Contrastive Learning in Image-to-Image Translation. In: Proc. of CVPR, pp. 18291–18300. IEEE/CVF (2022)
31. Guo, Z., Guo, D., Zheng, H., Gu, Z., Zheng, B., Dong, J.: Image Harmonization with Transformer. In: Proc. of ICCV, pp. 14870–14879. (2021)
32. Dinh, C.T., Tran, N., Nguyen, J.: Personalized Federated Learning with Moreau Envelopes. In: Proc. of NeurIPS, 33, 21394–21405 (2020)
33. Zhu, J.-Y., Park, T., Isola, P., Efros, A.A.: Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In: Proc. of ICCV, pp. 2223–2232. IEEE (2017)
34. 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., Lane, N.D.: Flower: A friendly federated learning research framework. arXiv preprint arXiv:2007.14390 (2020)
35. Jha, D., Smedsrud, P.H., Riegler, M.A., Halvorsen, P., De Lange, T., Johansen, D., Johansen, H.D.: Kvasir-SEG: A segmented polyp dataset. In: Proc. of MMM, LNCS, pp. 451–462. Springer, Cham (2019)
36. Silva, J., Histace, A., Romain, O., Dray, X., Granado, B.: Toward embedded detection of polyps in WCE images for early diagnosis of colorectal cancer. *IJCARS* 9(2), 283–293 (2014)
37. Bernal, J., Sánchez, J., Vilarino, F.: Towards Automatic Polyp Detection with a Polyp Appearance Model. *Pattern Recognit* 45(9), 3166–3182 (2012)
38. Bernal, J., Sánchez, F.J., Fernández-Esparrach, G., Gil, D., Rodríguez, C., Vilariño, F.: WMDOVA maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Comput. Med. Imaging Graph* 43, 99–111 (2015)
39. Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Sci. Data* 5(1), 180161 (2018)
40. Mendonça, T., Ferreira, P.M., Marques, J., Marcal, A.R.S., Rozeira, J.: PH<sup>2</sup> - A dermoscopic image database for research and benchmarking. In: Proc. of IEEE EMBC, Osaka, Japan (2013)
41. Codella, N.C.F., Nguyen, Q.-B., Pankanti, S., Gutman, D.A., Helba, B., Halpern, A.C., Smith, J.R.: Skin lesion analysis toward melanoma detection: 2017 ISBI challenge. arXiv preprint arXiv:1703.00523 (2017)
42. Codella, N.C.F., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S.W., Gutman, D.A., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M.A., Kittler, H., Halpern, A.C.: Skin lesion analysis toward melanoma detection 2018: 2017 ISIC challenge. arXiv preprint arXiv:1807.07391 (2018)
43. Gómez-Flores, W., Gregorio-Calas, M.J., Coelho de Albuquerque Pereira, W.: BUS-BRA: A breast ultrasound dataset for assessing computer-aided diagnosis systems. *Medical Physics* 51(4), 3110–3123 (2024)
44. Rodrigues, P.S., et al.: BUS\_UC: Breast ultrasound dataset. Mendeley Data, V1 (2017). <https://data.mendeley.com/datasets/3ksd7w7jkk/1>
45. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. *Data in Brief* 28, 104863 (2020)
46. Yap, M.H., Pons, G., Martí, J., Ganau, S., Sentís, M., Zwiggelaar, R., Davison, A.K., Martí, R.: Automated breast ultrasound lesions detection using convolutional neural networks. *IEEE JBHI* 22(4), 1218–1226 (2018)
47. Jaeger, S., Candemir, S., Antani, S., Wáng, Y.-X.J., Lu, P.-X., Thoma, G.: Two public chest X-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative Imaging in Medicine and Surgery* 4, 475–477 (2014)
48. Liu, Y., Wu, Y.-H., Ban, Y., Wang, H., Cheng, M.-M.: Rethinking Computer-Aided Tuberculosis Diagnosis. In: Proc. of CVPR, IEEE/CVF (2020) In: Proc. of CVPR, pp. 2646–2655. IEEE/CVF (2020)
49. Bhuvaji, S.: Brain tumor classification (MRI). Kaggle dataset (2020). <https://www.kaggle.com/datasets/sartajbhuvaji/brain-tumor-classification-mri>
50. rm1000: Brain tumor MRI scans. Kaggle dataset. <https://www.kaggle.com/datasets/rm1000/brain-tumor-mri-scans>
51. Dubail, T.: Brain tumors (256x256). Kaggle dataset (2020). <https://www.kaggle.com/datasets/thomasdubail/brain-tumors-256x256>
52. Cheng, J., Huang, W., Cao, S., Yang, R., Yang, W., Yun, Z., Wang, Z., Feng, Q.: Enhanced performance of brain tumor classification via tumor region augmentation and partition. *PLoS ONE* 10(10), e0140381 (2015)