

FedCLAM: Client Adaptive Momentum with Foreground Intensity Matching for Federated Medical Image Segmentation

Vasilis Siomos¹[0009–0003–0985–2672], Jonathan Passerat-Palmbach^{1,2}[0000–0003–3178–9502], and Giacomo Tarroni^{1,2}[0000–0002–0341–6138]

¹ CitAI Research Centre, Department of Computer Science City St. George's, University of London

² BioMedIA, Department of Computing, Imperial College London
vasilis.siomos@city.ac.uk

Abstract. Federated learning is a decentralized training approach that keeps data under stakeholder control while achieving superior performance over isolated training. While inter-institutional feature discrepancies pose a challenge in all federated settings, medical imaging is particularly affected due to diverse imaging devices and population variances, which can diminish the global model's effectiveness. Existing aggregation methods generally fail to adapt across varied circumstances. To address this, we propose FedCLAM, which integrates *client-adaptive momentum* terms derived from each client's loss reduction during local training, as well as a *personalized dampening factor* to curb overfitting. We further introduce a novel *intensity alignment* loss that matches predicted and ground-truth foreground distributions to handle heterogeneous image intensity profiles across institutions and devices. Extensive evaluations on two datasets show that FedCLAM surpasses eight cutting-edge methods in medical segmentation tasks, underscoring its efficacy. The code is available at <https://github.com/siomvas/FedCLAM>.

Keywords: Federated Learning · Data Heterogeneity · Segmentation

1 Introduction

Federated Learning (FL) facilitates collaborative model training across multiple institutions without the exchange of patient data, thus maintaining privacy and enhancing performance [19]. The seminal FedAvg algorithm [12] underperforms when local datasets are non-independent and identically distributed (non-IID) [8], which is typically the case in medical imaging due to different acquisition protocols, scanner types, and patient demographics. Multiple strategies have tried to address the challenge of achieving consistent performance across all clients by modifying the aggregation rule (by, e.g., incorporating momentum or adaptive optimization [5,16], or adding regularization terms to handle client drift [18,1]). Although these methods mitigate some of the destructive effects




Fig. 1. Diagram of the proposed **FedCLAM** approach: Training dynamics are fed into the CLient Adaptive Momentum (CLAM) module to produce the new global model. Local training is aided by our proposed Foreground Intensity Matching (FIM) Loss which alleviates intensity differences between client imaging devices.

of heterogeneity, they often fail to adapt to each institution’s unique learning trajectory or imaging characteristics.

In this paper, we propose FedCLAM (Client Adaptive Momentum with Foreground Intensity Matching), a novel FL algorithm tailored to medical image segmentation. FedCLAM uses a momentum-based update rule with client-adaptive momentum and dampening factors derived from each client’s local training progress, ensuring that clients with greater validation improvements guide the global model more prominently. In tandem, our Foreground Intensity Matching (FIM) loss addresses site-specific intensity differences by aligning the intensities of predicted and ground-truth foreground distributions. This design choice is motivated by the fact that even models that excel at delineating anatomical structures might still learn site-specific brightness or contrast biases. As shown by our ablation study, FIM can also be integrated into other FL methods, and FedCLAM remains robust even without FIM. Furthermore, our proposed approach is easy to tune, with consistent performance across a wide range of hyperparameters. We demonstrate FedCLAM’s effectiveness through extensive experiments on two multi-center segmentation tasks: retinal fundus images and prostate MRI. Notably, FedCLAM achieves higher Dice performance and better fairness across clients than state-of-the-art FL methods.

Our key contributions are summarized as follows: (i) We introduce a novel client-adaptive, momentum-based aggregation method that derives per-client momentum and dampening terms from the local loss reduction and overfitting ratio. (ii) We propose a Foreground Intensity Matching loss component addressing data variability in medical imaging contexts. (iii) We present extensive validations

on two real-world multi-centric medical datasets, demonstrating state-of-the-art performance with minimal to no hyper-parameter tuning requirements.

2 Related Works

Federated learning has been shown to provide benefits compared to local training in medical imaging problems such as segmentation [20]. However, the seminal FedAvg [12] simply averages local model weights, assuming independent and identically distributed (IID) data across clients. In practice, medical images vary by scanner, protocol, and patient demographics, violating this assumption, and in such scenarios the performance of FedAvg suffers [5,8,18].

To address non-IID data distributions, researchers have proposed modified aggregation schemes. A popular approach is to modify the aggregation weights: FedCostWAvg [11] leverages the decrease of the cost function, IDA [23] increases the weights of models that lie close to the global one in parameter space, while FedEvi [3] uses an epistemic uncertainty model. Momentum, widely recognized as a technique in centralized training, was introduced into FL by FedAvgM [5]. Momentum helps stabilize training in non-IID settings [4], but FedAvgM uses a *uniform* momentum factor, which cannot capture distinct client learning trajectories caused by non-IID data. Other adaptive approaches, (e.g. FedOpt [16]), only consider progress on the training set, thus being susceptible to the effects of overfitting. Overfitting to local data is particularly problematic in medical FL where each site might have limited or biased samples. Different approaches have been proposed to combat this. FedProx [18] introduces a proximal term to limit local-global divergence, while FedDyn [1] employs historical regularizers to mitigate client drift. FedSAM [15] and HarmoFL [6] minimize loss sharpness to improve generalization, but neither adapts to local dynamics nor penalizes overfitting directly. Furthermore, previous works often rely on histogram matching or normalization steps as preprocessing, the effectiveness of which is limited by the principle of no data sharing between clients. Feature distribution alignment methods like FedFA [24] and FedDG [9] can also tackle cross-site variability but require architecture modifications or significant overhead.

FedCLAM bridges a crucial gap between previously proposed methods by simultaneously (i) deriving client-specific momentum buffers from local validation progress, offering finer control over aggregation, (ii) directly controlling for overfitting during aggregation by measuring the ratio of training loss to validation loss and using it to dampen the speed vector on a per-client basis, limiting the influence of clients with small or site-specific datasets, (iii) a Foreground Intensity Matching (FIM) module that shifts feature alignment to the loss function, while being lightweight and not requiring specialized layers.

3 Methodology

3.1 Preliminaries and Overview

We consider a cross-silo [7] scenario like the one depicted in Figure 1, where a set of clients $\mathcal{S}$ (each with a local training $\mathcal{D}_{train,i}$ and validation $\mathcal{D}_{valid,i}$ dataset) engage in Federated Learning in order to train a global model w_g . At every round r of training, the FedCLAM process can be separated into the following components: i) local training with an optional novel Foreground Intensity Matching (FIM) loss, ii) derivation of personalized speed vectors v_i using momentum β_i and dampening τ_i terms adapted to each client’s dynamics, and iii) calculation of the new global model based on the average speed vector. The pseudocode for FedCLAM is presented in Algorithm 1.

3.2 Client Adaptive Momentum and Dampening

The use of momentum when aggregating client models has recently been theoretically shown to help alleviate the non-IIDness of the data [4]. However, applying a uniform momentum factor to all clients fails to account for local training dynamics, leading to suboptimal performance. We hypothesize that using local training dynamics to derive per-client momentum and dampening terms can result in a more performant momentum-based aggregation algorithm. Specifically, in every round r , we create a momentum signal β_i (1), adapted to the local learning dynamics of each client, by measuring the sigmoid σ of the relative decrease in the validation loss before and after local training $\frac{L_{val,i}^{init} - L_{val,i}}{L_{val,i}}$. Simultaneously, we form personalized dampening factors τ_i (2), which slow down clients that overfit locally, by measuring the ratio of the training loss and validation loss $\frac{L_{train,i}}{L_{val,i}}$. Finally, β_i and τ_i are used to derive per-client *speed vectors* as in the typical momentum formulation:

$$\text{momentum: } \beta_i = \sigma \left(k \cdot \frac{L_{val,i}^{init} - L_{val,i}}{L_{val,i}} \right) \quad (1)$$

$$\text{dampening: } \tau_i = 1 - \left(\frac{L_{train,i}}{L_{val,i}} \right)^\alpha \quad (2)$$

$$\text{Speed vectors: } v_i^r = \beta_i \cdot v_i^{r-1} + (1 - \tau_i) \cdot \Delta^r \quad (3)$$

Here, Δ^r is the pseudo-gradient, i.e. the average of local updates $(w_g - w_i)$, and we initialize the speed vectors as $v_i = \Delta^0$. The use of the sigmoid function is motivated by the need for the momentum factor to be a smooth but informative signal lying in the region of $[0, 1]$. Eq. 1 results in an adaptive signal which at every round assigns higher momentum to those clients which decreased their validation loss more. An optional steepness factor k in 1 accommodates different loss scales in different tasks, while an optional exponent α in 2 controls the dampening strength.

Algorithm 1. FedCLAM

```

1: Input: Global Model  $w_g^0$ , Rounds  $R$ , Local Epochs  $E$ , weight  $\lambda_{FIM}$ , dampening
   exponent  $\alpha$ , steepness  $k$ , aggregator learning rate  $lr$ , local learning rate  $\eta_l$ 
2: for  $r = 0, \dots, R - 1$  do
3:   for each client  $i \in \mathcal{S}$  in parallel do
4:      $w_{i,0}^r = w_g^r$ 
5:     Compute initial validation loss  $L_{val,i}^{init}$ 
6:     for  $e = 0, \dots, E - 1$  do
7:       Compute an unbiased estimate  $g_{i,e}^r$  of  $\nabla F_i(w_{i,e}^r)$  using Eq. 5
8:        $w_{i,e+1}^r = w_{i,e}^r - \eta_l \cdot g_{i,e}^r$ 
9:       Store train loss  $L_{train,i}$ 
10:    Compute validation loss  $L_{val,i}$ 
11:     $\Delta_i^r = w_{i,E}^r - w_g^r$  ▷ Local Update
12:     $\Delta^r = 1/|\mathcal{S}| \sum_{i \in \mathcal{S}} \Delta_i^r$  ▷ Pseudo-Gradient
13:    for each client  $i \in \mathcal{S}$  do
14:      if  $r=0$  then
15:         $v_i^r = \Delta_i^r$ 
16:      else
17:         $\beta_i = \sigma(k \cdot \frac{L_{val,i}^{init} - L_{val,i}}{L_{val,i}})$ ,  $\tau_i = 1 - \left(\frac{L_{train,i}}{L_{val,i}}\right)^\alpha$ 
18:         $v_i^r = \beta_i \cdot v_i^{r-1} + (1 - \tau_i) \cdot \Delta_i^r$ 
19:       $v_{avg} = 1 / |\mathcal{S}| \cdot \sum_{i \in \mathcal{S}} v_i^r$ 
20:       $w_g^{r+1} = w_g^r - lr \cdot v_{avg}$ 

```

3.3 Foreground Intensity Matching Loss

Heterogeneity in medical images often stems from differences in how scanners capture intensity values. Such mismatches complicate pixel-level tasks such as segmentation because the model might erroneously learn site-specific intensity cues instead of general anatomical representations. To mitigate this, we propose a novel *Foreground Intensity Matching* (FIM) loss that penalizes discrepancies between the predicted and ground-truth foreground regions. In standard segmentation, losses such as Dice or cross-entropy focus primarily on overlap or pixel-wise label correctness. Our additional FIM component explicitly enforces the distribution alignment for intensities in the predicted mask versus the actual annotated regions. Concretely, let F_G represent the set of pixel intensities belonging to the ground-truth foreground, and $\widehat{F}_G$ be the intensities of the predicted segmentation region. We consider these sets as 1-D samples from distributions we wish to align, and compute the 2-Wasserstein distance between them. In practice, for each pixel in the predicted mask, we compare its intensity with the intensity of a corresponding ground-truth pixel, after sorting the vectors. A soft weighting by the predicted probability (rather than a hard mask) preserves the original image dimensions so that both the predicted and ground-truth foreground intensity sets are flattened into 1-D vectors of equal length, enabling a

direct comparison of their sorted intensity distributions.

$$W_2(FG, \widehat{FG}) = \left(\frac{1}{n} \sum_{i=1}^n \|X_i - Y_i\|^2 \right)^{0.5} \quad (4)$$

where $X_i \in FG$ and $Y_i \in \widehat{FG}$ and n is the number of pixels in the image. We incorporate this term into our overall loss:

$$L_{total} = L_{seg} + \lambda_{FIM} \cdot L_{FIM} \quad (5)$$

Here, L_{seg} can be any standard segmentation loss (Dice, cross-entropy, or a combination) and λ_{FIM} is a weighing factor controlling how strongly we penalize intensity mismatch relative to standard segmentation errors. By forcing the network to match the intensity distributions in foreground regions, we reduce reliance on device-specific brightness or contrast cues. Unlike a fixed preprocessing step, the model dynamically learns to correct mismatches as it trains, adjusting the feature representations to align more naturally across centers. The FIM term requires minimal overhead beyond computing pixel intensities within the predicted mask. A visual demonstration of FIM is presented in Fig. 1.

4 Experiments

4.1 Experimental Settings

Datasets and Evaluation Metrics. We evaluate FedCLAM on two well-known medical image segmentation tasks derived from real-world multi-center data: 1) a retinal fundus dataset consisting of 4 clients [21,13,2] with 400, 400, 159, 101 images, where the task is joint segmentation of the optic disc and optic cup, and 2) a prostate MRI dataset sourced from 6 centers [10] containing 30, 30, 19, 12, 12 volumes. We treat each center as a separate client, mimicking real-world conditions. The images are pre-processed to a size of 384x384. For the prostate MRI dataset, we stack adjacent slices together to form inputs with 3 channels. We use the test-set Dice score (the average of disc and cup scores for fundus) on each client, and the average across clients, to evaluate the segmentation results. We also calculate the standard deviation of test performance between clients to gauge *fairness* [18].

Implementation Details. We use a U-Net [17] with Instance Norm [22] as the backbone network. Following [10] we use random flips to augment the prostate dataset. The models are trained with Adam with a learning rate of 10^{-4} for the prostate and 10^{-3} for fundus and a weight decay of 10^{-4} . We conduct 200 rounds for fundus and 500 rounds for prostate of 1 local epoch each, as is common practice. The loss function is the sum of Dice, cross-entropy, and our FIM loss for the prostate dataset and the sum of Dice loss and FIM for the fundus dataset. For the RGB fundus images we apply the FIM loss to the lightness channel of the

Table 1. Performance comparison for fundus segmentation, using the average Dice score. FedCLAM outperforms baselines in average test performance and fairness (measured using the standard deviation across clients).

Method	Client 1	Client 2	Client 3	Client 4	Average $\uparrow$	Std $\downarrow$
FedAvg [12]	74.53	90.83	90.80	82.71	84.72	7.79
FedAvgM [5]	81.58	90.52	90.55	80.81	85.86	5.40
FedEvi [3]	85.48	91.07	92.20	82.97	87.93	4.42
FedFA [24]	79.56	91.77	91.69	82.10	86.28	6.37
FedProx [18]	84.24	91.62	92.21	82.71	87.69	4.91
FedSAM [15]	81.44	91.42	90.69	80.64	86.05	5.80
FedDyn [1]	85.78	90.28	90.87	83.05	87.49	3.73
HarmoFL [6]	80.55	91.52	91.13	80.82	86.00	6.14
FedCLAM	88.14	90.30	91.66	85.17	88.82	2.83

Table 2. Performance comparison for MRI prostate segmentation, using the Dice score on the test set. The proposed FedCLAM outperforms baselines in average test performance and provides very competitive fairness.

Method	Client 1	Client 2	Client 3	Client 4	Client 5	Client 6	Average $\uparrow$	Std $\downarrow$
FedAvg [12]	91.46	93.06	92.61	95.4	93.84	90.24	92.77	1.80
FedAvgM [5]	89.08	93.46	94.31	96.00	94.18	89.44	92.74	2.82
FedEvi [3]	92.02	93.36	93.45	95.71	93.60	90.66	93.13	1.69
FedFA [24]	91.42	93.29	93.26	95.28	93.87	90.74	92.98	1.65
FedProx [18]	91.35	93.57	92.07	95.36	93.70	90.20	92.71	1.86
FedSAM [15]	92.44	93.23	93.02	95.57	94.30	90.41	93.16	1.74
FedDyn [1]	91.90	93.44	92.98	95.69	93.82	89.97	92.97	1.92
HarmoFL [6]	92.61	93.87	94.6	95.96	94.23	90.4	93.61	1.91
FedCLAM	92.85	94.61	94.77	95.71	94.27	90.21	93.74	1.96

CIELAB colorspace. We set the FIM weighing $\lambda_{FIM} = 10^{-2}$ for both datasets. The sigmoid steepness k and the dampening exponent α are both set to 1.0.

Compared Methods. We compare FedCLAM with vanilla FedAvg [12], FedAvgM [5], which uses a non-personalized momentum update, FedProx [18], which uses a proximal loss term to regularize the global model, FedDyn [1], which dynamically regularizes the model, [24], which uses feature alignment at the model level, FedSAM [15], which performs sharpness-aware minimization, HarmoFL [6], which uses a weight-perturbation optimizer similar to FedSAM, and FedEvi [3], which re-weights client contributions based on an evidential model for uncertainty.

4.2 Main Results and Discussion

Tables 1 and 2 present the segmentation results on the Fundus and Prostate MRI datasets, respectively, including performance on each client, average performance, and standard deviation across clients. For Fundus, all methods perform significantly better than FedAvg in terms of both average Dice score and




Fig. 2. Sensitivity analysis of k and λ_{FIM} on the Fundus dataset.

Table 3. Ablation study on the components of FedCLAM. Each component increases performance on its own, and their combination achieves the best result.

Components	Prostate Fundus	
None	92.77	84.72
FIM	93.07	86.93
CLAM	93.50	88.27
FIM + CLAM	93.74	88.82

standard deviation. Our method notably improves upon the best baseline in average performance, FedEvi, by 0.89% Dice, and the most fair baseline (lowest deviation), FedDyn, by 0.9%, offering an overall increase of 4.1% compared to FedAvg. For Prostate MRI, baseline performance is higher, hence the results are closer, but most methods still outperform FedAvg. FedCLAM outperforms the best baseline, HarmoFL, in average performance, and results in the best performance for 3 out of 6 clients overall, offering an increase of 0.97% over FedAvg.

4.3 Additional Experiments

Hyper-parameter sensitivity analysis. FedCLAM introduces 2 main hyper-parameters, namely the sigmoid steepness factor k in Eq. 1 and the weighting of the FIM loss λ_{FIM} in Eq. 5. We vary the values of k in the set $\{1, 2, 5, 10, 20\}$, and the values of λ_{FIM} in the set $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$ and re-train on the Fundus dataset to study the level of sensitivity of FedCLAM with respect to them and present the results in Figure 2. We observe that FedCLAM is very robust to values of k , but very high values result in a decrease compared to the optimum as β_i becomes too sensitive to loss changes. Similarly, for λ_{FIM} values in the range $[10^{-4}, 10^{-2}]$ performance remains consistent, but significantly decreases for higher values as the scale of the loss becomes too high. Overall, this analysis reinforces that it is easy to tune FedCLAM, and based on it we recommend default values of $k = 1, \lambda_{FIM} = 10^{-2}$.

Component Ablation. We study the effect of each of the main components of FedCLAM, by ablating them in turn. When neither the client adaptive momentums nor the FIM loss are used, the method reverts to FedAvg. As shown in Table 3, both components contribute to increasing the model’s performance, with client adaptive momentum offering the bigger performance boost: 0.73 and 3.55 Dice points on the prostate and fundus datasets, respectively. These results highlight FedCLAM can function independently of FIM and be successfully deployed in tasks without mismatches in the intensity of ground-truth foreground, and that our FIM loss can potentially increase the performance of other FL aggregation methods.

5 Conclusion and Discussion

We present a novel FL aggregation method aimed at medical image segmentation named FedCLAM. By incorporating personalized, adaptive momentum and dampening components, FedCLAM improves non-IID performance while avoiding overfitting and promoting performance fairness. A novel loss module aligning intensity distributions between the ground-truth and predicted foreground areas further improves performance. Comprehensive evaluations demonstrate that our method surpasses the current state-of-the-art in two standard FL medical segmentation tasks. Finally, our ablation results confirm that each component contributes significant performance gains while requiring minimal tuning, positioning FedCLAM to play a pivotal role in advancing practical, privacy-preserving collaborations across diverse clinical and research settings. Potential extensions of our work include evaluating the framework’s performance with more advanced backbone architectures and benchmarking against a centralized learning approach to quantify any performance gap. It would also be beneficial to explore the integration of formal privacy-preserving mechanisms to more accurately reflect real-world scenarios. Furthermore, to provide a more nuanced analysis under data heterogeneity, future studies could incorporate federated learning-specific evaluation metrics, such as convergence scores [14].

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

References

1. Acar, D.A.E., Zhao, Y., Navarro, R.M., Mattina, M., Whatmough, P.N., Saligrama, V.: Federated learning based on dynamic regularization. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021), <https://openreview.net/forum?id=B7v4QMR6Z9w>
2. Batista, F.J.F., Diaz-Aleman, T., Sigut, J., Alayon, S., Arnay, R., Angel-Pereira, D.: Rim-one dl: A unified retinal image database for assessing glaucoma using deep learning. *Image Analysis and Stereology* **39**(3), 161–167 (2020)
3. Chen, J., Ma, B., Cui, H., Xia, Y.: Fedevi: Improving federated medical image segmentation via evidential weight aggregation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 361–372. Springer (2024)
4. Cheng, Z., Huang, X., Wu, P., Yuan, K.: Momentum benefits non-iid federated learning simply and provably. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=TdhkAcXkRi>
5. Hsu, T.H., Qi, H., Brown, M.: Measuring the effects of non-identical data distribution for federated visual classification. *CoRR* **abs/1909.06335** (2019), <http://arxiv.org/abs/1909.06335>
6. Jiang, M., Wang, Z., Dou, Q.: Harmoff: Harmonizing local and global drifts in federated learning on heterogeneous medical images. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 1087–1095 (2022)

7. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al.: Advances and Open Problems in Federated Learning. arXiv:1912.04977 [cs, stat] (Dec 2019), <http://arxiv.org/abs/1912.04977>, arXiv: 1912.04977
8. Li, X., Huang, K., Yang, W., Wang, S., Zhang, Z.: On the convergence of fedavg on non-iid data (2020), <https://arxiv.org/abs/1907.02189>
9. Liu, Q., Chen, C., Qin, J., Dou, Q., Heng, P.: Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. CoRR **abs/2103.06030** (2021), <https://arxiv.org/abs/2103.06030>
10. Liu, Q., Dou, Q., Yu, L., Heng, P.A.: Ms-net: Multi-site network for improving prostate segmentation with heterogeneous mri data. IEEE Transactions on Medical Imaging **39**(9), 2713–2724 (Sep 2020). <https://doi.org/10.1109/tmi.2020.2974574>, <http://dx.doi.org/10.1109/TMI.2020.2974574>
11. Mächler, L., Ezhov, I., Kofler, F., Shit, S., Paetzold, J.C., Loehr, T., Wiestler, B., Menze, B.H.: Fedcostwavg: A new averaging for better federated learning. CoRR **abs/2111.08649** (2021), <https://arxiv.org/abs/2111.08649>
12. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Artificial intelligence and statistics. pp. 1273–1282. PMLR (2017)
13. Orlando, J.I., Fu, H., Breda, J.B., Van Keer, K., Bathula, D.R., Diaz-Pinto, A., Fang, R., Heng, P.A., Kim, J., Lee, J., et al.: Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. Medical image analysis **59**, 101570 (2020)
14. Pati, S., Baid, U., Zenk, M., Edwards, B., Sheller, M., Reina, G.A., Foley, P., Gruzdev, A., Martin, J., Albarqouni, S., Chen, Y., Shinohara, R.T., Reinke, A., Zimmerer, D., Freymann, J.B., Kirby, J.S., Davatzikos, C., Colen, R.R., Kotrotsou, A., Marcus, D., Milchenko, M., Nazer, A., Fathallah-Shaykh, H., Wiest, R., Jakab, A., Weber, M.A., Mahajan, A., Maier-Hein, L., Kleesiek, J., Menze, B., Maier-Hein, K., Bakas, S.: The federated tumor segmentation (fets) challenge (2021)
15. Qu, Z., Li, X., Duan, R., Liu, Y., Tang, B., Lu, Z.: Generalized federated learning via sharpness aware minimization (2022), <https://arxiv.org/abs/2206.02618>
16. Reddi, S.J., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., McMahan, H.B.: Adaptive federated optimization. CoRR **abs/2003.00295** (2020), <https://arxiv.org/abs/2003.00295>
17. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18. pp. 234–241. Springer (2015)
18. Sahu, A.K., Li, T., Sanjabi, M., Zaheer, M., Talwalkar, A., Smith, V.: On the convergence of federated optimization in heterogeneous networks. CoRR **abs/1812.06127** (2018), <http://arxiv.org/abs/1812.06127>
19. Sheller, M.J., Edwards, B., Reina, G.A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R.R., Bakas, S.: Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific Reports **10**(1), 12598 (Jul 2020). <https://doi.org/10.1038/s41598-020-69250-1>, <https://www.nature.com/articles/s41598-020-69250-1>, number: 1 Publisher: Nature Publishing Group
20. Sheller, M.J., Reina, G.A., Edwards, B., Martin, J., Bakas, S.: Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain

- tumor segmentation. In: Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part I 4. pp. 92–104. Springer (2019)
21. Sivaswamy, J., Krishnadas, S., Chakravarty, A., Joshi, G., Ujjwal: A comprehensive retinal image dataset for the assessment of glaucoma from the optic nerve head analysis. JSM Biomed Imaging Data Pap **2** (01 2015)
 22. Ulyanov, D., Vedaldi, A., Lempitsky, V.S.: Instance normalization: The missing ingredient for fast stylization. CoRR **abs/1607.08022** (2016), <http://arxiv.org/abs/1607.08022>
 23. Yeganeh, Y., Farshad, A., Navab, N., Albarqouni, S.: Inverse distance aggregation for federated learning with non-iid data. CoRR **abs/2008.07665** (2020), <https://arxiv.org/abs/2008.07665>
 24. Zhou, T., Konukoglu, E.: Fedfa: Federated feature augmentation. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023. OpenReview.net (2023), <https://openreview.net/forum?id=U9yFP90jU0>