



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

FedIA: Federated Medical Image Segmentation with Heterogeneous Annotation Completeness

Yangyang Xiang^{*1}, Nannan Wu^{*1}, Li Yu¹, Xin Yang¹,
Kwang-Ting Cheng², and Zengqiang Yan^{1(✉)}

¹ School of Electronic Information and Communications, Huazhong University of Science and Technology

{xyy_2001, wnn2000, hustlyu, xinyang2014, z_yan}@hust.edu.cn

² School of Engineering, Hong Kong University of Science and Technology
timcheng@ust.hk

Abstract. Federated learning has emerged as a compelling paradigm for medical image segmentation, particularly in light of increasing privacy concerns. However, most of the existing research relies on relatively stringent assumptions regarding the uniformity and completeness of annotations across clients. Contrary to this, this paper highlights a prevalent challenge in medical practice: incomplete annotations. Such annotations can introduce incorrectly labeled pixels, potentially undermining the performance of neural networks in supervised learning. To tackle this issue, we introduce a novel solution, named FedIA. Our insight is to conceptualize incomplete annotations as noisy data (*i.e.*, low-quality data), with a focus on mitigating their adverse effects. We begin by evaluating the completeness of annotations at the client level using a designed indicator. Subsequently, we enhance the influence of clients with more comprehensive annotations and implement corrections for incomplete ones, thereby ensuring that models are trained on accurate data. Our method’s effectiveness is validated through its superior performance on two extensively used medical image segmentation datasets, outperforming existing solutions. The code is available at <https://github.com/HUSTxyy/FedIA>.

Keywords: Federated learning · Incomplete annotation · Noisy label learning · Segmentation.

1 Introduction

Recent progress in federated learning (FL) [13] has facilitated the collaborative training of unified models across multiple decentralized entities in a privacy-preserving manner [2,6,19,22]. In medical domains, FL has seen extensive application in training segmentation models for distinct lesions and organs [18,8,20]. Nevertheless, an essential limitation in current research is the insufficient consideration of the diversity in *annotation completeness* among clients.

* Equal contribution.




Fig. 1. Heterogeneity in annotation completeness among clients. Red and blue solid lines represent the boundaries of marked lesions and unmarked lesions, respectively.

This issue primarily stems from the varying standards of annotation adopted by various clients. As depicted in Fig. 1, certain clients (*i.e.*, client k) provide complete annotations for comprehensive diagnosis and analysis. Conversely, others (*i.e.*, client i and j) may possess incomplete annotations where only partial regions are marked, to minimize labeling costs, which are adequate only for basic image-level assessments (*e.g.*, rapid screening).

Given this heterogeneity in annotation completeness, training a segmentation model via FL poses significant challenges. The inclusion of clients with incomplete annotations creates a situation where these clients are considered to be of lower quality since partial positive regions are mislabeled as background. Such imperfect annotations can negatively affect the overall performance of the model due to the memory effect of neural networks [12,11]. To tackle this, in this paper, we focus on the important yet under-explored problem: ***How to pursue better FL under heterogeneity in annotation completeness?***

Within the realm of FL, there has been some work focusing on data heterogeneity [4,5,7], but the heterogeneity in annotation completeness has been often overlooked. As for strategies to diminish the negative impact of clients with low-quality labels, these solutions predominantly focus on the classification task [23,21], which is suboptimal when applied to the segmentation task. Although FedA^{3I} [20] has recently addressed the heterogeneity in annotation quality specific to segmentation, its underlying assumption, where mislabeled pixels mainly distribute near objects’ boundaries, renders it ineffective against the challenge of incomplete annotations. Consequently, developing an effective approach to address this critical issue remains an area in need of further exploration and insight.

In this study, we tackle the pressing problem of heterogeneity in annotation completeness by introducing **FedIA**, a FL solution that is cognizant of and adaptively corrects for client annotation completeness. Our foundational insight is to perceive incomplete annotations as akin to noisy data. We commence by developing an early model robust against the noise associated with incomplete annotations, which then serves as a basis for evaluating each client’s level of annotation completeness. Subsequently, our aggregation process prioritizes clients with higher annotation completeness, and clients undertake annotation corrections before local model updating supervised by incomplete annotations. Our

approach has been tested on two real-world medical datasets: a brain multiple sclerosis MRI dataset and a COVID-19 lesion CT dataset. Results show that FedIA outperforms other SOTA methods designed to address noisy/imperfect annotations.

The contributions of this paper are three-fold: (1) A new FL problem concentrating on heterogeneity in annotation completeness; (2) A novel solution named FedIA to tackle incomplete annotations; (3) Extensive evaluation to demonstrate the superiority of the proposed solution.

2 Methodology

2.1 Preliminaries and Overview

This paper focus on a a single-class multi-lesion¹ segmentation problem in a federated scenario. Given K clients, each client possesses its private dataset $D_k = \{(x_i \in \mathcal{X} \subseteq \mathbb{R}^{H \times W \times C}, \tilde{y}_i \in \mathcal{Y} = \{0, 1\}^{H \times W})\}_{i=1}^{n_k}$, where n_k is the size of D_k and (x_i, \tilde{y}_i) represents the image-annotation pair characterized by dimensions: height H , width W , and channel C . Contrary to an ideal situation, the annotations in our case are considered imperfect due to incompleteness, with not every lesion being marked. The completeness ratio $a_k = c_{k,i}^n / c_{k,i}^g$, indicating the proportion of marked lesions to the total actual lesions within D_k , which remains identical among samples in D_k but differs across clients.

Our objective is to devise a robust algorithm capable of diminishing the negative effects of incomplete annotations on the global model’s accuracy. The cornerstone of our approach involves deriving an initial model that is minimally affected by noise through the utilization of extensive noisy data, followed by assessing each client’s annotation completeness ratio based on this initial model. The strategy prioritizes learning from clients with higher completeness rates (*i.e.*, higher-quality data), thereby enhancing the model’s performance. Furthermore, a mechanism is incorporated to correct incomplete annotations at a certain stage of the learning process, using a specially designed metric based on Intersection over Union (IoU). The overview of FedIA is illustrated in Fig. 2.

2.2 Annotation Completeness Estimation

Assessing the level of annotation completeness across clients is imperative, as it directly influences the tailored handling of each client’s data. To accomplish this, obtaining a model that is unaffected by imperfect annotations becomes essential. Our approach begins by interpreting incomplete annotations as noisy labels, with unmarked lesions contributing noise by altering pixel-level labels from 1 to 0. Drawing inspiration from the early learning phenomenon in noisy label learning [11,10,12], which posits that neural networks initially adapt to clean labels in the early stages before progressively accommodating noisy labels, we can place confidence in the training process despite the prevalence of

¹ This means an image can contain multiple lesions, each forming a connected region.

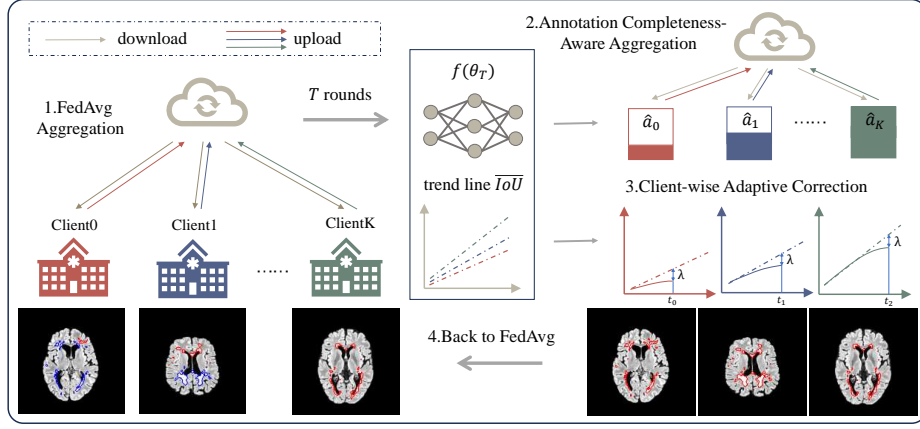


Fig. 2. Overview of the proposed FedIA. The first stage is the early learning phase, the global model is updated by FedAvg [13]. The second is the modification stage, re-weighting each client by calculating its annotation completeness rate and correcting incomplete annotations synchronously. In the last stage, local models are trained with the corrected labels and aggregated for federated updating through FedAvg [13].

noisy labels and utilize the early model phase to gauge annotation completeness across clients. Specifically, we develop an early-stage global model parameterized by θ_T , capable of basic segmentation, by undergoing a warm-up period for T communication rounds employing FedAvg [13]. This process involves the server aggregating client models based on their respective data contributions to formulate the global model, defined as:

$$\theta_t = \sum_{k=1}^K \frac{n_k}{n} \theta_{t,k}, \quad (1)$$

where t denotes the current training round under the constraint $1 \leq t \leq T$, and n represents the amount of data from all clients, *i.e.*, $n = \sum_{k=1}^K n_k$. In this stage, local optimization of each client is established on:

$$\min_{\theta_k} \sum_{(x_i, y_i) \in D_k} \ell_{dc}(f(x_i; \theta_{t,k}), \tilde{y}_i), \quad (2)$$

where $f: \mathcal{X} \rightarrow [0, 1]^{\text{H} \times \text{W}}$ is the neural network and $\ell_{dc}: [0, 1]^{\text{H} \times \text{W}} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is the Dice loss [14]. After obtaining the warm-up model parameterized by θ_T that is relatively unaffected by noise, the annotation completeness rate \hat{a}_k of each client k is estimated by the following formula:

$$\hat{a}_k = \frac{\sum_{i=1}^{n_k} c_{k,i}^n}{\sum_{i=1}^{n_k} c_{k,i}^p}, \quad (3)$$

where $c_{k,i}^n$ and $c_{k,i}^p$ denote the number of lesions in the noisy label \tilde{y}_i and in the predicted map $y_i^p = f(x_i; \theta_T)$, respectively.

2.3 Annotation Completeness-Aware Aggregation

Quantity-based aggregation (*i.e.*, FedAvg) is susceptible to noise caused by incomplete annotations, especially when such annotations are numerous [20]. To mitigate this, clients with higher annotation quality should dominate FL more. Therefore, we calculate a completeness-aware aggregation weight w_k^t at round t for each client k , defined as

$$w_k^t = \frac{\exp\left(\frac{\hat{a}_k}{\ell_k^t}\right)}{\sum_{j=1}^K \exp\left(\frac{\hat{a}_j}{\ell_j^t}\right)}, \quad (4)$$

where ℓ_k^t denotes the average loss of client k at round t calculated by

$$\ell_k^t = \frac{\sum_{(x_i, y_i) \in D_k} \ell_{dc}(f(x_i; \theta_{t,k}), \tilde{y}_i)}{n_k}. \quad (5)$$

Generally, the observed loss is lower when annotation completeness is elevated. Consequently, the server prioritizes clients exhibiting lower losses, effectively reducing the negative effects of imprecise estimation of a_k on the weighting process, potentially arising from inappropriate selection of T .

2.4 Client-wise Adaptive Correction

The volume of data significantly influences the performance of neural networks, and datasets characterized by low annotation completeness represent valuable resources that should not be overlooked. Hence, rectifying incomplete annotations to acquire cleaner data for further training of the model is essential. Given that different clients pose datasets with varying levels of annotation completeness, the onset of noise impact and their robustness to noise vary accordingly.

To capture this information, in the early learning phase (*i.e.*, $1 \leq t \leq T$), we compute IoU values every round for each client k and fit it with the first-order polynomial function:

$$\overline{\text{IoU}}_k(t) = l_k \cdot t + b_k, \quad (6)$$

where l_k and b_k are two parameters of the polynomial function. After early learning and in sync with the completeness-aware aggregation, we monitor the change of IoU_k every round. Any client satisfying the following formula will correct its annotations in the next round:

$$\overline{\text{IoU}}_k(t) - \text{IoU}_k(t) > \lambda, \quad (7)$$

where $\text{IoU}_k(t)$ denotes the actual IoU value of client k at round t . λ is an adjustable hyperparameter, set to 0.03 by default. In addition, the client only corrects annotations for which its model output predicted probability has confidence above a certain threshold setting of 0.8. It is worth noting that we only correct the pixels with a value of 0 because only false negative lesions and no false positive lesions are presented in this setting.

3 Experiments

3.1 Datasets and Implementation Details

Datasets. Two public medical image segmentation datasets are included:

1. Two real-world multiple sclerosis datasets, focusing on the segmentation of white matter lesions (WML) in 3D magnetic resonance (MR) brain images, denoted as **MS**, including MSSEG-1 [1] and PubMRI [9]. In the task, we only use the FLAIR modality, in which the lesions are relatively clear.
2. The widely-used COVID dataset, aiming at segmentation and quantification of lung lesions caused by SARS-CoV-2 infection from computed tomography (CT) images, denoted as **LUNG**. [16]

Each dataset is divided into training and test sets by a ratio of 8:2, whose training set is then randomly split into four clients. For computational efficiency, all 3D samples are converted into 2D slices and resized to 256×256 pixels.

To verify the robustness to different degrees of incompleteness of our method, several settings are used for evaluation. Specifically, for **MS**, the annotation completeness rate of the k -th client is set as $20\% \times k - 10\% \times m + 40\%$. And we conduct four sets of experiments, *i.e.*, $m = 0, 1, 2, 3$. For **LUNG**, three different settings are used, and the completeness rates are formulated as $10\% \times k - 30\% \times m + 70\%$, where $m = 0, 1, 2$.

Incomplete Annotation Generation When doctors or other professionals label multi-lesion data, they tend to label one lesion at the 3D volume level before annotating another. Therefore, to simulate real noise and generate incomplete annotation \tilde{v}_j , lesions are randomly removed at the 3D level. This process allows us to mimic real-world conditions more accurately. Specifically, we first set the annotation completeness of each client as a_k , which is unknown during training. Then, we calculate the number of lesion-connected components c_j^n in each 3D sample v_j and randomly choose c_j^n lesion regions, where $c_j^n = c_j^g \cdot a_k$. Only the chosen lesion regions are kept as well-annotated while others are set as background (*i.e.*, incomplete/missing annotation).

Implementation Details. In this work, U-Net [15] is adopted as the foundational model architecture for FL. The FL training process is designed to include a total of 300 communication rounds, with each local training phase consisting of a single local epoch. During local training, the model undergoes optimization via the Adam optimizer with momentum terms set to (0.9, 0.99), a batch size of 4, and an initial learning rate of 1e-4. To accommodate the early learning strategy, the initial learning round, T , is set as 10 for **MS** and 40 for **LUNG**.

3.2 Comparison with State-of-the-art Methods

In our analysis, we benchmark FedIA against recent leading methods tailored to address label noise in both classification and segmentation tasks, including ELR

Table 1. Comparison results with state-of-the-art methods under the **MS** and **LUNG** settings. $c_0/c_1/c_2/c_3$ means the annotation completeness rates a_k of clients are $c_k \times 10\%$, corresponding to the setting of m . The average results (%) from the last ten rounds are reported. The best results are marked in bold.

Methods	From	MS				LUNG		
		$m = 0$	$m = 1$	$m = 2$	$m = 3$	$m = 0$	$m = 1$	$m = 2$
		4/6/8/10	3/5/7/9	2/4/6/8	1/3/5/7	7/8/9/10	4/5/6/7	1/2/3/4
FedAvg	AISTATS'17	60.88	55.77	34.65	13.98	54.40	46.74	28.47
ELR	NeurIPS'20	63.01	57.97	35.91	9.50	49.29	35.01	11.22
NR-Dice	TMI'20	69.19	64.42	60.16	23.60	58.36	54.32	33.75
ADELE	CVPR'22	61.34	58.63	25.01	0	54.33	44.14	17.89
FedCorr	CVPR'22	62.03	57.12	30.35	0	55.24	50.56	23.78
RMD	TMI'23	62.77	59.15	41.60	17.15	48.60	32.79	9.37
FedNoRo	IJCAI'23	67.09	60.78	39.74	31.53	49.86	37.25	19.16
FedIA	Ours	74.73	74.03	69.22	56.53	59.42	55.34	44.72




Fig. 3. Qualitative comparison on **MS** where *others* represents FedAvg, ELR, ADELE, and FedCorr failing to segment any lesion. Red, blue and green color show the prediction of true-positive, false-negative and false-positive regions, respectively.

(NeurIPS'20) [12] and ADELE (CVPR'22) [11], which leverage the early learning phenomenon to prevent model overfitting to noisy labels; RMD (TMI'23) [3], which mitigates annotation noise in medical imaging through mutual distillation; NR-Dice (TMI'20) [17], introducing a noise-robust Dice loss to combat noisy labels; FedNoRo (IJCAI'23) [21], designed to manage class-imbalanced noisy data; and FedCorr (CVPR'22) [23], employing the LID score to identify noisy clients. Additionally, we incorporate the universally recognized FL framework, FedAvg [13], as a baseline for comparison. *Detailed implementations of these methods are available in the supplementary material.*

Quantitative comparison results on **MS** and **LUNG** measured by Dice coefficient (%) are summarized in Table 1. Notably, our FedIA exhibits consistent performance even as annotation completeness diminishes, in contrast to some methods whose effectiveness wanes with decreased annotation completeness. This demonstrates that FedIA surpasses other sophisticated methods across various datasets and configurations, underscoring the robustness and efficacy of our strategy in tackling the challenge.

Exemplar qualitative comparison on **MS** in annotation completeness setting of 10%, 30%, 50%, 70% is illustrated in Fig. 3. FedIA effectively recalls all lesions with fewer false positives, leading to the best segmentation performance. Com-

Table 2. Component-wise study.**Table 3.** Impact of the Round T

FedAvg	ACAG	CAC	MS				T	LUNG		
			0	1	2	3		0	1	2
✓			60.88	55.77	34.65	13.98	10	56.53	52.25	46.88
✓	✓		71.91	70.34	66.95	43.40	20	59.36	54.78	45.37
✓		✓	74.23	73.67	67.43	37.74	30	59.52	55.07	44.74
✓	✓	✓	74.73	74.03	69.22	56.53	40	59.42	55.34	44.72

paratively, FedNoRo, RMD, and NR-Dice suffer from extensive false negatives, resulting in noisy segmentation. What’s worse, FedAvg and other comparison methods completely fail to segment any lesion, indicating the difficulty in addressing heterogeneous annotation completeness and the necessity and value of FedIA. *More qualitative results are available in the supplementary material.*

3.3 Ablation Study

Component-wise Study. We conduct an ablation study by separately removing the Annotation Completeness-Aware Aggregation (ACAG) and the Client-wise Adaptive Correction (CAC) components from FedIA as summarized in Table 2. We observe that FedAvg can benefit from both components, particularly under the lowest annotation completeness settings. This phenomenon demonstrates the effectiveness of our designs against annotation noise. The best performance is typically achieved when both components are incorporated.

Impact of the Early Learning Round T . It is worth noting that the early learning phase, denoted by T , is set differently for the two tasks as **LUNG** presenting a more complex learning challenge compared to **MS**, essentially requiring a longer early learning period. This variation prompts a relevant question regarding the optimal number of training rounds necessary for effective early training. To address this, we perform ablation studies on **LUNG** under various T settings: 10, 20, 30, and 40 as summarized in Table 3. The results indicate that our method exhibits considerable robustness to changes in T . Notably, FedIA consistently outperforms the baseline FedAvg across all tested T selections, demonstrating its robustness and superior performance irrespective of the early learning duration. *More ablation studies are available in the supplementary material for reference.*

4 Conclusion

In this study, we tackle a significant yet overlooked challenge in federated medical image segmentation: how to enhance FL against heterogeneity in annotation completeness. We approach incomplete annotations as akin to noisy data, employing strategies to mitigate their negative impacts denoted as FedIA. FedIA

involves initially assessing the level of annotation completeness at the client level through designed indicators. Then, it prioritizes clients with greater annotation completeness and undertakes corrective measures for those with incomplete ones, aiming to ensure that the training process is mainly based on accurate knowledge. After rigorously evaluated through a line of experiments on two extensively utilized medical image segmentation datasets, experimental results affirm the effectiveness of FedIA, showcasing its advantage over current leading approaches. We believe that the issue formulated and the proposed solution will pave the way for more practical FL applications in complex medical scenarios.

Acknowledgement. This work was supported in part by the National Natural Science Foundation of China under Grants 62202179 and 62271220, and in part by the Natural Science Foundation of Hubei Province of China under Grant 2022CFB585. The computation is supported by the HPC Platform of HUST.

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

References

1. Commowick, O., Istace, A., Kain, M., Laurent, B., Leray, F., Simon, M., Pop, S.C., Girard, P., Ameli, R., Ferré, J.C., et al.: Objective evaluation of multiple sclerosis lesion segmentation using a data management and processing infrastructure. *Scientific reports* **8**(1), 13650 (2018) [6](#)
2. Dou, Q., So, T.Y., Jiang, M., Liu, Q., Vardhanabhuti, V., Kaissis, G., Li, Z., Si, W., Lee, H.H., Yu, K., et al.: Federated deep learning for detecting COVID-19 lung abnormalities in CT: a privacy-preserving multinational validation study. *npj Digit. Medicine* **4**(1), 60 (2021) [1](#)
3. Fang, C., Wang, Q., Cheng, L., Gao, Z., Pan, C., Cao, Z., Zheng, Z., Zhang, D.: Reliable mutual distillation for medical image segmentation under imperfect annotations. *IEEE Trans. Med. Imaging* (2023) [7](#)
4. Huang, W., Ye, M., Du, B.: Learn from others and be yourself in heterogeneous federated learning. In: *CVPR* (2022) [2](#)
5. Huang, W., Ye, M., Shi, Z., Du, B.: Generalizable heterogeneous federated cross-correlation and instance similarity learning. *TPAMI* (2023) [2](#)
6. Huang, W., Ye, M., Shi, Z., Li, H., Du, B.: Rethinking federated learning with domain shift: A prototype view. In: *CVPR* (2023) [1](#)
7. Huang, W., Ye, M., Shi, Z., Wan, G., Li, H., Du, B., Yang, Q.: A federated learning for generalization, robustness, fairness: A survey and benchmark. *arXiv* (2023) [2](#)
8. Jiang, M., Roth, H.R., Li, W., Yang, D., Zhao, C., Nath, V., Xu, D., Dou, Q., Xu, Z.: Fair federated medical image segmentation via client contribution estimation. In: *CVPR*. pp. 16302–16311 (2023) [1](#)
9. Lesjak, Ž., Galimzianova, A., Koren, A., Lukin, M., Pernuš, F., Likar, B., Špičlin, Ž.: A novel public mr image dataset of multiple sclerosis patients with lesion segmentations based on multi-rater consensus. *Neuroinformatics* **16**, 51–63 (2018) [6](#)
10. Li, J., Socher, R., Hoi, S.C.: DivideMix: Learning with noisy labels as semi-supervised learning. In: *ICLR* (2020) [3](#)

11. Liu, S., Liu, K., Zhu, W., Shen, Y., Fernandez-Granda, C.: Adaptive early-learning correction for segmentation from noisy annotations. In: CVPR. pp. 2606–2616 (2022) [2](#), [3](#), [7](#)
12. Liu, S., Niles-Weed, J., Razavian, N., Fernandez-Granda, C.: Early-learning regularization prevents memorization of noisy labels. In: NeurIPS. vol. 33, pp. 20331–20342 (2020) [2](#), [3](#), [7](#)
13. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: AISTATS. pp. 1273–1282 (2017) [1](#), [4](#), [7](#)
14. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3DV. pp. 565–571 (2016) [4](#)
15. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241. Springer (2015) [6](#)
16. Roth, H.R., Xu, Z., Tor-Díez, C., Jacob, R.S., Zember, J., Molto, J., Li, W., Xu, S., Turkbey, B., Turkbey, E., et al.: Rapid artificial intelligence solutions in a pandemic—the covid-19-20 lung ct lesion segmentation challenge. *MIA* **82**, 102605 (2022) [6](#)
17. Wang, G., Liu, X., Li, C., Xu, Z., Ruan, J., Zhu, H., Meng, T., Li, K., Huang, N., Zhang, S.: A noise-robust framework for automatic segmentation of COVID-19 pneumonia lesions from CT images. *IEEE Trans. Med. Imaging* **39**(8), 2653–2663 (2020) [7](#)
18. Wang, J., Jin, Y., Wang, L.: Personalizing federated medical image segmentation via local calibration. In: ECCV. pp. 456–472 (2022) [1](#)
19. Wu, N., Kuang, Z., Yan, Z., Yu, L.: From optimization to generalization: Fair federated learning against quality shift via inter-client sharpness matching. In: IJCAI (2024) [1](#)
20. Wu, N., Sun, Z., Yan, Z., Yu, L.: FedA3I: Annotation quality-aware aggregation for federated medical image segmentation against heterogeneous annotation noise. In: AAAI (2024) [1](#), [2](#), [5](#)
21. Wu, N., Yu, L., Jiang, X., Cheng, K.T., Yan, Z.: FedNoRo: Towards noise-robust federated learning by addressing class imbalance and label noise heterogeneity. In: IJCAI (2023) [2](#), [7](#)
22. Wu, N., Yu, L., Yang, X., Cheng, K.T., Yan, Z.: FedIIC: Towards robust federated learning for class-imbalanced medical image classification. In: MICCAI (2023) [1](#)
23. Xu, J., Chen, Z., Quek, T.Q., Chong, K.F.E.: FedCorr: Multi-stage federated learning for label noise correction. In: CVPR. pp. 10184–10193 (2022) [2](#), [7](#)