

# FedA<sup>3</sup>I: Annotation Quality-Aware Aggregation for Federated Medical Image Segmentation against Heterogeneous Annotation Noise

Nannan Wu, Zhaobin Sun, Zengqiang Yan\*, Li Yu

School of Electronic Information and Communications, Huazhong University of Science and Technology, Wuhan, China  
wnn2000@hust.edu.cn, zbsun@hust.edu.cn, z\_yan@hust.edu.cn, hustlyu@hust.edu.cn

## Abstract

Federated learning (FL) has emerged as a promising paradigm for training segmentation models on decentralized medical data, owing to its privacy-preserving property. However, existing research overlooks the prevalent annotation noise encountered in real-world medical datasets, which limits the performance ceilings of FL. In this paper, we, for the first time, identify and tackle this problem. For problem formulation, we propose a contour evolution for modeling non-independent and identically distributed (Non-IID) noise across pixels within each client and then extend it to the case of multi-source data to form a heterogeneous noise model (*i.e.*, Non-IID annotation noise across clients). For robust learning from annotations with such two-level Non-IID noise, we emphasize the importance of data quality in model aggregation, allowing high-quality clients to have a greater impact on FL. To achieve this, we propose **Federated learning with Annotation quAlity-aware AggregatIon**, named **FedA<sup>3</sup>I**, by introducing a quality factor based on client-wise noise estimation. Specifically, noise estimation at each client is accomplished through the Gaussian mixture model and then incorporated into model aggregation in a layer-wise manner to up-weight high-quality clients. Extensive experiments on two real-world medical image segmentation datasets demonstrate the superior performance of FedA<sup>3</sup>I against the state-of-the-art approaches in dealing with cross-client annotation noise. The code is available at <https://github.com/wnn2000/FedAAAI>.

## 1 Introduction

Accurate medical image segmentation is crucial as it is typically the first step in analytical workflows. Recent developments in segmentation suggest collecting massive data to train deep learning (DL) models (Ronneberger, Fischer, and Brox 2015; Gu et al. 2019; Valanarasu and Patel 2022). However, growing privacy concerns in medical scenarios prohibit data transfer among sources (Rieke et al. 2020; Dou et al. 2021; Wu et al. 2023b; Yan et al. 2020; Li et al. 2021), making it infeasible to construct datasets at scale. Fortunately, federated medical image segmentation (FMIS) (Liu et al. 2021) based on federated learning (FL) (McMahan et al. 2017; Huang et al. 2023a,b) offers a solution




Figure 1: Annotation noise in multi-source datasets. The blue and red curves represent the contours of clean (*i.e.*, ideal) and noisy annotations, respectively. Given a certain sample, annotations between two curves are noisy, indicating that noise is distributed near contours instead of being IID among pixels. Given samples from different clients, annotation noise is heterogeneous where noise on client  $i$  causes under-segmentation and noise on client  $j$  causes over-segmentation.

by enabling multiple medical institutions to train a unified segmentation model cooperatively without data sharing. As FMIS enriches training data in a privacy-preserving manner, it has emerged as a promising paradigm in practice.

Existing research on FMIS has mainly focused on distribution shift (Liu et al. 2021; Xu et al. 2022a; Wang, Jin, and Wang 2022; Jiang et al. 2023b; Wang et al. 2023), client-side unfairness (Jiang et al. 2023a), and region-of-interest (ROI) inconsistency (Xu et al. 2023) caused by decentralized data. However, prevalent noisy annotations, referring to annotations containing errors, in medical datasets (Ji et al. 2021; Zhang et al. 2020a) have often been overlooked, which are inevitable in real-world medical datasets. On the one hand, annotating is subjective that can be biased by annotators' expertness (*i.e.*, **preference**). For instance, front-line clinicians tend to mark ROIs larger to avoid missing positive parts (*i.e.* tumors/lesions) in ambiguous regions (Calisto et al. 2021; Salehi, Erdogmus, and Gholipour 2017), while non-professionals may only recognize well-characterized core areas, leading to smaller annotated regions. On the other hand, the level of annotator engagement can also affect the

\*Corresponding author

quality of annotations, resulting in random mislabeling (*i.e.*, **randomness**) when annotators are distracted or fully occupied. It should be noted that noisy annotations can seriously degrade the performance of segmentation models under supervised learning paradigms (Liu et al. 2022). Thus, addressing cross-client annotation noise in FMIS is an urgent task demanding insightful investigation.

Despite extensive studies on addressing related issues like

1. Noisy labels for classification (Xu et al. 2022b; Wu et al. 2023a) and weak annotations for segmentation (Wicaksana et al. 2022; Zhu, Chen, and Yuan 2023) in FL,
2. Noisy annotations for segmentation in centralized learning (CL) (Liu et al. 2022; Yao et al. 2023),

noisy annotations in FMIS may not be well addressed based on existing methods. It is because those methods fail to address the two-level Non-IID noise in annotations across both pixels and clients in FMIS, as shown in Fig. 1. Specifically, at the pixel level, annotation noise tends to be distributed in indistinct regions near contours, caused by inaccurate contour delineation, making it difficult to analyze the noise as the first type of studies mainly focuses on addressing relatively uniform noise within each classification instance. As for the client level, as annotation noise is highly annotator-dependent, noise is heterogeneous across clients due to non-overlapping annotators under the context of multi-source data in FL (Wu et al. 2023a). Such cross-client scenario makes it almost impossible to model annotation noise from a universal perspective, posing additional challenges when applying strategies from the second type of studies that primarily address homogeneous annotation noise. Consequently, how to effectively address noisy annotations in FMIS remains unexplored.

In this paper, we, for the first time, identify and formulate this new FMIS problem, *i.e.*, how to train a federated segmentation model with noisy pixel-wise annotations. Different from weak annotations like bounding boxes studied in previous FMIS works (Wicaksana et al. 2022; Zhu, Chen, and Yuan 2023), noisy annotations studied in this paper are more common and realistic in clinical practice and thus more worthy of being addressed. To formulate annotation noise, inspired by the real-world annotating process of delineating ROI contours, we first propose a generic model to generate noisy annotations named contour evolution model (CEM). In CEM, noise is generated by making a clean contour evolve into a corrupted one, determined by two main parameters corresponding to annotators' preference and randomness. Following a CEM with certain parameters, annotation noise within a client/annotator is produced where noise among pixels is Non-IID. To generate heterogeneous annotation noise in multi-source data, different clients are assigned with CEMs with different parameters, resulting in Non-IID annotation noise across clients. To address such two-level Non-IID annotation noise in FMIS, our insight resides in that model aggregation in FL should consider not only data quantity as FedAvg (McMahan et al. 2017) but also data quality (*i.e.*, annotation quality in this paper). Therefore, we incorporate **Federated learning with Annotation quality-aware Aggregation**, named

**FedA<sup>3</sup>I**, for noise-robust FMIS. In FedA<sup>3</sup>I, the annotation noise of each client is estimated by the Gaussian mixture model (GMM) with the learning difficulty of regions around contours, which is used to compute the quality-based aggregation weight. This step is conducted on the region-level rather than the pixel-level to mitigate the impact of Non-IID noise across pixels on noise analysis. After obtaining clients' quality- and quantity-based (*i.e.*, FedAvg) weights, instead of directly adjusting the corresponding client-wise weights for model-wise aggregation (Liang et al. 2022; Wicaksana et al. 2022; Wu et al. 2023a), we dynamically combine the two weights in a layer-wise manner (Ma et al. 2022). To be specific, quantity-based weights contribute more to updating shallower layers while quality-based weights dominate the updates of deeper layers. It is motivated by the consensus that noisy labels have a more severe impact on deep layers (Bai et al. 2021). The main contributions are as follows:

- A new FMIS problem where two-level Non-IID noise is identified to be addressed for the first time.
- A generic annotation noise model, named CEM, to formulate an annotator's preference and randomness, and a heterogeneous annotation noise model to formulate cross-annotator/-client annotation noise based on CEMs with different parameters.
- A novel FL framework FedA<sup>3</sup>I to address two-level Non-IID noise in FMIS, where both data quantity and quality are included for layer-wise model aggregation.
- Superior performance against the state-of-the-art (SOTA) approaches validated on two publicly-available real-world medical datasets.

## 2 Related Work

### 2.1 Federated Medical Image Segmentation

FMIS has shown great potential for decentralized medical datasets under privacy preservation. Existing research on FMIS mainly emphasizes addressing data heterogeneity and its derived problems, *e.g.*, distribution shift (Liu et al. 2021; Xu et al. 2022a; Wang, Jin, and Wang 2022; Jiang et al. 2023b; Wang et al. 2023), inconsistent ROIs (Xu et al. 2023) and inter-client unfairness (Jiang et al. 2023a). Recently, some studies (Wicaksana et al. 2022; Zhu, Chen, and Yuan 2023) address weak annotations in FMIS. FedMix (Wicaksana et al. 2022) explores how to learn from mixed-type annotations including complete and weak annotations, while FedDM (Zhu, Chen, and Yuan 2023) deals with typical bounding box annotations.

### 2.2 Learning with Noisy Labels/Annotations

It is important to design noise-robust training strategies given the inevitable mislabeling phenomenon. Existing studies on this are conducted under both CL (Han et al. 2018; Yu et al. 2019; Wang et al. 2019; Zhang and Sabuncu 2018; Wang et al. 2020; Liu et al. 2020, 2022; Yao et al. 2023; Fang et al. 2023; Li, Socher, and Hoi 2020) and FL paradigms (Xu et al. 2022b; Wu et al. 2023a; Jiang et al. 2022), mainly based on sample selection (Han et al. 2018; Yu et al. 2019), noise-robust loss functions (Wang et al. 2019; Zhang and

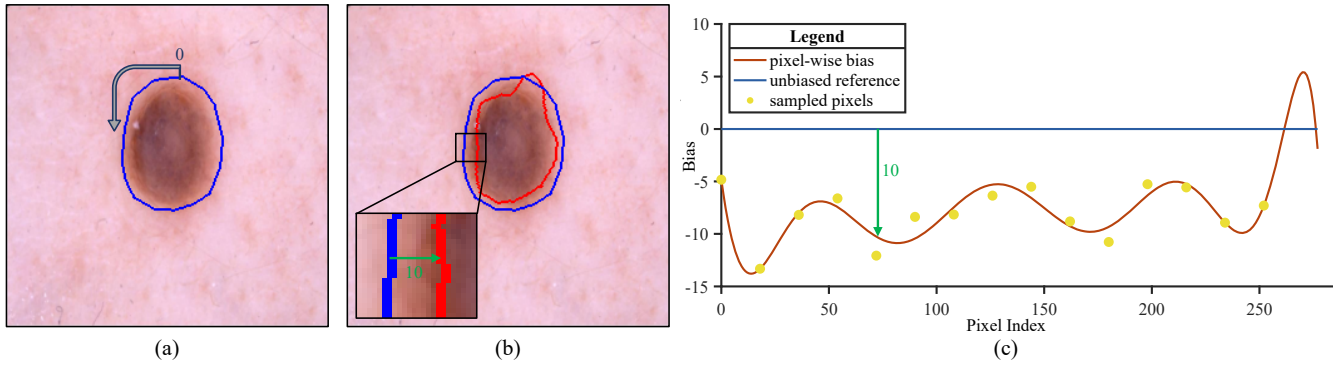


Figure 2: An example of using CEM (*i.e.*,  $C(-8, 2)$ ) to generate annotation noise. In (a), the arrow indicates the positive annotation direction starting from the origin (*i.e.*, 0) of the contour (*i.e.*, the blue curve). Then, the bias of each pixel is calculated based on the polynomial function fitted by sampled pixels as illustrated in (c), and used to control the movements of pixels of the contour as indicated by the green arrow, resulting in the noisy contour (*i.e.*, the red curve) in (b).

Sabuncu 2018; Wang et al. 2020), noise-robust regularization (Liu et al. 2020; Jiang et al. 2022), label correction (Liu et al. 2022; Xu et al. 2022b; Yao et al. 2023), label softening (Wu et al. 2023a; Fang et al. 2023), semi-supervised learning (Li, Socher, and Hoi 2020), *etc.*

### 3 Methodology

#### 3.1 Preliminaries

For simplification, this paper focuses on the typical binary segmentation problem<sup>1</sup> following previous studies (Wicansana et al. 2022; Zhu, Chen, and Yuan 2023). Given  $K$  participants in FL, each client  $k$  holds a 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  $n_k$  represents the size of  $D_k$ , and  $(x_i, \tilde{y}_i)$  denotes an available image-annotation pair with the height, width, and channel dimensions of  $H$ ,  $W$ , and  $C$ , respectively. Note that the annotation  $\tilde{y}_i$  is noisy due to the inevitable mislabeling phenomenon, described as  $\tilde{y}_i = y_i \oplus \xi_i$ , where  $y_i$  is the underlying true annotation,  $\xi_i \in \mathcal{Y}$  represents the pixel-wise noise, and  $\oplus$  is the pixel-wise binary XOR operation. In our work, we focus on two-level Non-IID noise, *i.e.*, the pixel level and the client level. Our goal is to develop a noise-robust FL algorithm to train a binary segmentation model  $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{H \times W}$  parameterized by  $\theta$ .

#### 3.2 Noise Model

To formulate the problem, we first present a noise model to simulate pixel-wise annotations and then generalize it to the case of multi-source datasets.

**Contour Evolution Model** Inspired by the clinical practice where annotating is typically performed by delineating ROI contours and noise is essentially introduced by biased delineation, we propose a contour evolution model (CEM) to describe annotation noise as a biased contour. In this way, contour bias can be quantified as a pixel bias sequence  $\langle b_k \rangle_{k=1}^l = \langle b_1, b_2, \dots, b_l \rangle$ , where  $l$  is the number of contour

pixels. Each pixel moves outward (with a positive bias) or inward (with a negative bias) along the normal direction according to the absolute value of its bias, forming a corrupted contour. To further analyze the bias, we decompose it into two orthogonal components as  $b_k = c + r_k$ , where  $c$  is a fixed component and  $r_k$  is the  $k$ -th element of a random component  $\langle r_k \rangle_{k=1}^l$  with  $\sum_{k=1}^l r_k = 0$ . These two components reflect two real-world factors contributing to the bias, namely preference and randomness. In clinical practice, annotators may draw contours of different sizes based on expertise and preference (Zhang et al. 2020a), resulting in annotated objects being larger (*i.e.*  $c > 0$ ) or smaller (*i.e.*  $c < 0$ ) than ground truth. As for randomness, annotators may make mistakes, resulting in random distortion being introduced along contours (*i.e.*,  $r_k$ ).

Based on the analysis above, we can generate annotator-specific annotation noise by controlling the two components. Specifically, given  $l$  contour pixels, a subset of  $l_{sub}$  pixels are selected with equal intervals, and  $l_{sub}$  values are sampled from a normal distribution  $\mathcal{N}(\mu, \sigma^2)$ . Then,  $l_{sub}$  values, along with their  $l_{sub}$  indices, are employed to fit a  $p$ -th degree polynomial function  $\mathcal{P}$  that is used to recompute the bias sequence with indices, *i.e.*,  $\langle b_k \rangle_{k=1}^l = \langle \mathcal{P}(1), \mathcal{P}(2), \dots, \mathcal{P}(l) \rangle$ . With this design, a pixel-wise bias with a fixed component approximately equal to  $\mu$  and a random component controlled by  $\sigma$  is generated. In addition, the bias sequence  $\langle b_k \rangle_{k=1}^l$  is continuous, ensuring contour continuity and smoothness. Each contour pixel is then moved along the normal direction to a new position based on the corresponding element in  $\langle b_k \rangle_{k=1}^l$ , resulting in annotator-specific noise controlled by  $(\mu, \sigma)$ . For convenience, we denote such a contour evolution model as  $C(\mu, \sigma)$  as illustrated in Fig. 2.

Compared to existing noise models like class-conditional noise (CCN) (Han et al. 2018; Yu et al. 2019; Chen et al. 2021), morphological operation noise (MON) (Zhang et al. 2020a; Liu et al. 2022), and Markov label noise (MLN) (Yao et al. 2023), CEM generates more realistic annotation noise as it directly simulates the real-world annotating process.

<sup>1</sup>This can be easily generalized to multi-class segmentation by studying each class separately.

Additionally, CEM considers the Non-IID nature of noise at the pixel level more comprehensively. To illustrate this, we revisit different noise models from the pixel-dependent noise (PDN) perspective by introducing the signed distance function (Yao et al. 2023).

**Definition 1 (PDN).** For an underlying clean annotation, denote the set of all pixels as  $\mathcal{A}$  and the signed distance function as  $\mathcal{D} : \mathcal{A} \rightarrow \mathbb{R}$ . Assume the imposed noise  $\xi$  is drawn from a random matrix  $\Xi = (\Xi_{i,j})$ .  $\xi$  is pixel-dependent if

1.  $\exists \epsilon > 0$ , s.t.  $\sum_{(i,j) \in \mathcal{I}} \frac{\mathbb{E}[\Xi_{i,j}]}{|\mathcal{I}|} > \sum_{(i,j) \in \mathcal{O}} \frac{\mathbb{E}[\Xi_{i,j}]}{|\mathcal{O}|}$ , where  $\mathcal{I} = \{(i,j) \in \mathcal{A} \mid |\mathcal{D}((i,j))| < \epsilon\}$ , and  $\mathcal{O} = \mathcal{A} \setminus \mathcal{I}$ .
2.  $\exists \delta \in \mathbb{R}$  and two different pixels  $(u_1, v_1), (u_2, v_2) \in \mathcal{A}$ , satisfying  $\mathcal{D}((u_1, v_1)) = \mathcal{D}((u_2, v_2)) = \delta$ , s.t.  $\Xi_{u_1, v_1}$  and  $\Xi_{u_2, v_2}$  are non-identically distributed.

The first condition ensures that noise is localized near contours, and the second condition ensures that noise is heterogeneous along contours. PDN indicates that noise in pixel-wise annotations is Non-IID across both the normal and tangential directions of contours, which fits real-world scenarios due to the following two reasons. First, when an annotator segments an object by marking its contours, noise is introduced to the regions between the marked and the underlying clean contours. That is, from the normal direction, noise mainly exists around contours. Second, the difficulty of marking different contour parts could vary since medical objects often lack isotropy. As a consequence, Non-IID noise along the tangential direction could arise. However, existing noise models (e.g., CCN, MON, MLN) fail to guarantee pixel-dependent due to the strong assumption of noise distribution. Comparatively, CEM well satisfies the above conditions by formulating annotation noise as a biased contour and altering the bias of different contour parts flexibly. More discussions and proof can be found in Appendix<sup>2</sup> A.

**Heterogeneous Noise Model with CEMs** In FMIS, annotators are non-overlapping across clients, resulting in heterogeneous noise due to its annotator-dependent nature. Based on CEM to generate annotator-specific noise, we present a heterogeneous noise model to emulate real-world noise in multi-source datasets by assigning CEMs with different parameters to clients. For each client  $i$ , we use a specific CEM  $C(\mu_i, \sigma_i)$  to generate its annotation noise by setting

$$\mu_i = \begin{cases} \mu \sim \mathcal{U}(0, \mu^{max}), & \text{with probability } p_d \\ \mu \sim \mathcal{U}(\mu^{min}, 0), & \text{with probability } 1 - p_d, \end{cases} \quad (1)$$

and  $\sigma_i = \sigma \sim \mathcal{U}(\frac{\sigma^{max}}{2}, \sigma^{max})$ , where  $\mathcal{U}$  denotes the uniform distribution. In this way, different clients can generate heterogeneous noise according to client-dependent CEMs. Such a multi-source noise model is denoted as  $M(\mu^{max}, \mu^{min}, \sigma^{max}, p_d)$  for simplification.

### 3.3 Annotation Quality-Aware Aggregation

**Motivation and Overview** When adopting FedAvg (McMahan et al. 2017) for FMIS, the optimization objective

<sup>2</sup>The appendix is included in the arXiv version, which can be found in <https://github.com/wnn2000/FedAAAI>

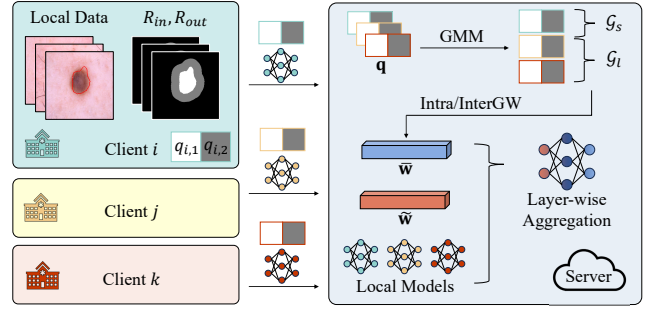


Figure 3: Overview of FedA<sup>3</sup>I. In each client  $i$ , we compute the learning difficulty of the inner and outer regions (the white and gray regions around contours respectively), denoted as  $q_{i,1}$  and  $q_{i,2}$ , and upload them to the server. These indicators are used to fit a Gaussian mixture model (GMM), which divides all clients into two subsets. Based on this, we compute the quality-based weights using two components, namely IntraGW and InterGW. Finally, both quality-based and quantity-based weights are utilized for model aggregation in a layer-wise manner.

can be formulated as

$$\min_{\theta} [F(\theta) := \frac{\sum_{i=1}^K n_i F_i(\theta, D_i)}{\sum_{i=1}^K n_i}], \quad (2)$$

where  $F_i(\theta, D_i)$  represents the local objective function of client  $i$ . The contribution of each client to the global model is determined solely by its data amount  $n_i$  in FedAvg. Consequently, the global model can be negatively diverted by clients with large amounts of noisy annotations. To address this, we propose FedA<sup>3</sup>I to make global model aggregation quality-dependent, thereby making high-quality clients dominate the training process as illustrated in Fig. 3. Details are presented as follows.

**Noise Estimation** We analyze annotation quality by examining the noise in each client. As annotation noise is heterogeneous across clients and annotated objects can be either larger or smaller than ground truth, the global model typically learns an intermediate pattern that is relatively clean for a faster drop in loss minimization during the early stage of training. This phenomenon, known as early learning (Liu et al. 2022; Li, Socher, and Hoi 2020; Liu et al. 2020), allows us to obtain a model that is relatively free from noise and has preliminary segmentation ability for noise estimation. Specifically, assuming an early model  $f_{\theta_e}$  warmed up by FedAvg with  $T_1$  rounds, we first calculate the client-wise learning difficulty  $\mathbf{q} = (q_{i,j}) \in \mathbb{R}^{K \times 2}$  with

$$q_{i,1} = \frac{1}{n_i} \sum_{(x,\tilde{y}) \in D_i} \sum_{(u,v) \in R_{in}(\tilde{y})} \frac{\ell_{ce}(f_{\theta_e}(x), \tilde{y})_{u,v}}{|R_{in}(\tilde{y})|} \quad (3)$$

and

$$q_{i,2} = \frac{1}{n_i} \sum_{(x,\tilde{y}) \in D_i} \sum_{(u,v) \in R_{out}(\tilde{y})} \frac{\ell_{ce}(f_{\theta_e}(x), \tilde{y})_{u,v}}{|R_{out}(\tilde{y})|}, \quad (4)$$

where  $\ell_{ce}$  is the pixel-wise cross entropy loss.  $R_{in}(\tilde{y})$  and  $R_{out}(\tilde{y})$  denote the inner and outer regions, starting from the contour of  $\tilde{y}$  to a distance  $d$ . Here,  $d$  is set as the maximal distance determined by gradually increasing  $d$  till either  $|R_{in}(\tilde{y})|$  or  $|R_{out}(\tilde{y})|$  no longer changes. In other words, we localize the largest regions that are most likely to contain annotation noise and calculate the corresponding average losses via Eqs. 3 and 4. The two loss values reflect the degree of divergence between  $f_{\theta_e}$ 's predictions and annotations around contours. It should be noted that using the largest possible regions is to mitigate inaccurate loss values due to wrongly-predicted pixels made by the early model  $f_{\theta_e}$ , while focusing on symmetric regions (*i.e.*,  $|R_{in}(\tilde{y})|$  and  $|R_{out}(\tilde{y})|$ ) can alleviate the impact of imbalanced sizes between inner and outer regions on divergence estimation (*e.g.*, a larger region will be less divergent as regions distant from contours are more likely to be predicted correctly). Thanks to the roughly accurate segmentation results from  $f_{\theta_e}$ , a larger value of  $q_{i,1}$  for client  $i$  suggests that it may have larger human-annotated objects, while a larger value of  $q_{i,2}$  implies the opposite. Therefore, noise in each client can be estimated by comparing its degree of divergence with that of the whole in globe. To this end, we compute  $q_{i,1}$  and  $q_{i,2}$  in each client and upload them to the server for obtaining  $\mathbf{q}$ . Given  $\mathbf{q}$ , a two-component Gaussian mixture model (GMM) is used to fit  $\mathbf{q}$  and classify all clients' indices into two groups:  $\mathcal{G}_l$  and  $\mathcal{G}_s$ , where  $\mathcal{G}_l$  and  $\mathcal{G}_s$  represent the indices of clients corresponding to Gaussian distributions with mean vectors  $(\mu_{l,1}, \mu_{l,2})^T$  and  $(\mu_{s,1}, \mu_{s,2})^T$  respectively under the constraint of  $\mu_{l,1} - \mu_{l,2} > \mu_{s,1} - \mu_{s,2}$ . In this way, clients in two groups own different patterns of noise. Client  $i$  ( $i \in \mathcal{G}_l$ ) tends to annotate objects larger due to greater inconsistency between annotations and model predictions in inner regions, while client  $j$  ( $j \in \mathcal{G}_s$ ) tends to annotate objects smaller. Note that this process is privacy-preserving as only average loss values are shared without leaking any information about local data distributions. After client grouping, noise strength can be estimated in each group by

$$s_i = \begin{cases} q_{i,1} - q_{i,2}, & \text{for } i \text{ in } \mathcal{G}_l, \\ q_{i,2} - q_{i,1}, & \text{for } i \text{ in } \mathcal{G}_s, \end{cases} \quad (5)$$

*i.e.*, the divergence difference between inner and outer regions. In each group, a greater  $s_i$  indicates stronger noise in client  $i$ .

**Aggregation Weight Calculation** As discussed above, quantity-based aggregation is vulnerable to noise. To address this, clients with high-quality annotations should be given more attention, which is achievable with the help of noise estimation above. Therefore, we compute a quality-based aggregation weight  $\bar{w}_i$  for each client  $i$ , defined as:

$$\bar{w}_i = \begin{cases} r \cdot \frac{\max_{j \in \mathcal{G}_l} s_j - s_i}{|\mathcal{G}_l| \max_{j \in \mathcal{G}_l} s_j - \sum_{t \in \mathcal{G}_l} s_t}, & \text{for } i \text{ in } \mathcal{G}_l, \\ (1-r) \cdot \frac{\max_{j \in \mathcal{G}_s} s_j - s_i}{|\mathcal{G}_s| \max_{j \in \mathcal{G}_s} s_j - \sum_{t \in \mathcal{G}_s} s_t}, & \text{for } i \text{ in } \mathcal{G}_s, \end{cases} \quad (6)$$

where  $r \in [0, 1]$  is a balancing coefficient set as 0.5 by default. In this way, a client with higher noise strength is

Noise Type	Dataset	
	<i>SKIN</i>	<i>BREAST</i>
$N_S$	$M(20, -20, 10, 0.2)$	$M(25, -15, 5, 0.2)$
$N_E$	$M(25, -25, 10, 0.8)$	$M(25, -15, 5, 0.8)$

Table 1: Details of two annotation noise settings.

given a smaller weight in each group, denoted as intra-group weighting (IntraGW). However, in real-world scenarios,  $\mathcal{G}_l$  and  $\mathcal{G}_s$  may differ significantly in size (*i.e.*,  $|\mathcal{G}_l|$  and  $|\mathcal{G}_s|$ ), and simply weighting clients within each group can still make the model misled by the majority noise. To address this, we introduce another parameter  $r$  to balance the weights between the two groups, denoted as inter-group weighting (InterGW). Our objective is to enable the global model to learn fairly from both annotation noise types, allowing it to learn accurate knowledge that goes beyond noise rather than over-fitting to the majority noise.

Given quality-based weights  $\bar{\mathbf{w}} = (\bar{w}_1, \dots, \bar{w}_K)$ , together with quantity-based weights  $\tilde{\mathbf{w}} = (\tilde{w}_1, \dots, \tilde{w}_K) = \frac{1}{\sum_{i=1}^K n_i} (n_1, \dots, n_K)$ , one straight-forward way for noise-robust aggregation is to use  $\bar{\mathbf{w}}$  to adjust  $\tilde{\mathbf{w}}$  for client-wise re-weighting (Liang et al. 2022; Wicaksana et al. 2022; Wu et al. 2023a). However, this method may be sub-optimal as it greatly underestimates the value of low-quality data for model training, especially in shallow layers. As discussed in (Bai et al. 2021), noisy labels would affect deep layers more compared to shallow layers. It is because the shallow layers of a model mainly learn general low-level features while the deep layers integrate semantic information for decision-making. In other words, low-quality data can still contribute to the training of shallow layers to maximize the utilization of training data, while high-quality data should dominate the training of deep layers to ensure that the decision-making process is not biased by noise. Therefore, we calculate a noise-robust aggregation weight  $w_{i,j}$  for any client  $i$  when updating the  $j$ -th layer of the global model, defined as

$$w_{i,j} = \mathcal{L}(j)\bar{w}_i + (1 - \mathcal{L}(j))\tilde{w}_i, j \in [L], \quad (7)$$

where  $L$  denotes the total number of layers in the trained model.  $\mathcal{L}(j) := \frac{j-1}{L-1}$  is a layer-wise trade-off coefficient which increases linearly as  $j$  increases to let  $\tilde{w}_i$  and  $\bar{w}_i$  dominate the updates of shallow and deep layers respectively. The layer-wise aggregation weight is employed in the subsequent process for noise-robust FL.

## 4 Experiments

### 4.1 Datasets and Evaluation Metric

Following prior studies addressing weak/noisy annotations (Wicaksana et al. 2022; Yao et al. 2023), two public medical image segmentation datasets are adopted, including

1. The widely-used ISIC 2017 skin lesion segmentation dataset (Codella et al. 2018), denoted as *SKIN*, contains 2000 images for training and 600 images for testing.
2. Three real-world breast ultrasound datasets for breast tumor segmentation, denoted as *BREAST*, including BUS

Loss	From	Method	Dataset			
			<i>SKIN</i>		<i>BREAST</i>	
			Noise Type			
			$N_S$	$N_E$	$N_S$	$N_E$
<i>CE</i>	NeurIPS'18	GCE	66.43 ± 2.21	73.47 ± 0.37	65.67 ± 1.51	67.12 ± 0.95
	CVPR'19	SCE	62.49 ± 4.31	74.69 ± 0.32	67.71 ± 1.37	67.73 ± 0.79
	NeurIPS'20	ELR	65.46 ± 1.43	74.31 ± 0.41	65.65 ± 1.29	67.93 ± 0.81
	CVPR'22	ADELE	62.99 ± 3.63	70.62 ± 0.95	61.79 ± 1.29	70.15 ± 2.41
	TMI'23	RMD	62.75 ± 2.60	75.41 ± 0.37	64.12 ± 2.64	68.09 ± 1.06
	MLSys'20	FedProx	64.35 ± 1.50	75.76 ± 0.59	68.10 ± 1.09	66.55 ± 1.01
	TMI'23	FedDM	42.44 ± 0.77	67.92 ± 2.32	40.76 ± 2.91	64.16 ± 3.58
	CVPR'22	FedCorr	61.88 ± 2.19	74.60 ± 1.17	67.15 ± 1.39	68.43 ± 1.10
	IJCAI'23	FedNoRo	66.35 ± 4.62	74.22 ± 0.30	59.51 ± 1.28	66.29 ± 5.03
	Ours	FedA <sup>3</sup> I	<b>72.69 ± 1.28</b>	<b>78.81 ± 0.75</b>	<b>69.03 ± 1.95</b>	<b>74.35 ± 1.90</b>
<i>DC</i>	TMI'20	NR-Dice	73.33 ± 1.38	73.75 ± 0.68	74.02 ± 0.64	65.46 ± 0.58
	TMI'22	FedMix	67.99 ± 1.78	73.00 ± 1.22	72.99 ± 0.15	69.04 ± 1.92
	Ours	FedA <sup>3</sup> I	<b>75.92 ± 0.99</b>	<b>77.05 ± 0.56</b>	<b>75.28 ± 0.53</b>	<b>76.68 ± 0.49</b>

Table 2: Comparison evaluation measured in the mean (%) ± standard deviation (%) of Dice coefficients against SOTA methods under different annotation noise settings. The best results are marked in bold.

(Al-Dhabyani et al. 2020), BUSIS (Zhang et al. 2022), and UDIAT (Yap et al. 2017). We perform stratified sampling to split each of the three datasets into a training set and a test set with an 8:2 ratio.

Following standard practice in dealing with noisy annotations (Yao et al. 2023; Zhu, Shi, and Wu 2019; Zhang et al. 2020b; Li, Gao, and He 2021), all images in *SKIN* and *BREAST* are resized to  $256 \times 256$  pixels in resolution. The Dice coefficient, as the standard metric for medical image segmentation, is employed for performance evaluation.

## 4.2 Experimental Setup

**Data Partition and Noise Generation** The training set of each dataset is partitioned into several subsets to simulate FL scenarios. Specifically, we randomly split both *SKIN* and *BREAST* in 50 subsets/clients. For *BREAST*, a client can only contain data from one of BUS, BUSIS, or UDIAT, and thus data across clients is Non-IID due to various sources.

For detailed evaluation, we use the proposed heterogeneous noise model to generate two types of annotation noise, *i.e.*  $N_E$  with  $p_d > 0.5$  and  $N_S$  with  $p_d < 0.5$ .  $N_E$  indicates that the majority of clients tend to annotate objects larger, while  $N_S$  represents the opposite. Detailed annotation noise settings are stated in Tab. 1.

**Implementation Details** We select U-Net (Ronneberger, Fischer, and Brox 2015) as the segmentation model. For FL training, the local epoch and the maximal communication round are set as 5 and 100 respectively. During local training, the model is trained by the Adam optimizer with momentum terms of (0.9, 0.99), a batch size of 8, and a constant learning rate of  $5e-3$ . The warm-up round  $T_1$  is set as 10 and 30 by default for *SKIN* and *BREAST* respectively. In each experiment, the training process is independently repeated five times, and the averaged mean and standard deviation of Dice coefficients are reported to eliminate randomness. All experiments are conducted using PyTorch on a NVIDIA RTX-3090 GPU with 24 GB of memory.

## 4.3 Comparison with SOTA Methods

For a comprehensive evaluation of FedA<sup>3</sup>I on addressing heterogeneous annotation noise, a set of SOTA methods are selected for comparison, including GCE (Zhang and Sabuncu 2018), SCE (Wang et al. 2019), ELR (Liu et al. 2020), ADELE (Liu et al. 2022), RMD (Fang et al. 2023), NR-Dice (Wang et al. 2020), FedProx (Li et al. 2020), FedMix (Wicaksana et al. 2022), FedDM (Zhu, Chen, and Yuan 2023), FedCorr (Xu et al. 2022b), and FedNoRo (Wu et al. 2023a). The introduction and implementation details of these methods can be found in Appendix B.

The above methods can be divided into two groups according to the used loss functions: Cross entropy loss-based methods (denoted as *CE*) and Dice loss-involved methods (denoted as *DC*). For a fair comparison, FedA<sup>3</sup>I is implemented by two types of loss functions separately, namely 1) cross-entropy loss and 2) a combination of cross-entropy loss and Dice loss (Milletari, Navab, and Ahmadi 2016). Quantitative comparison results on both datasets under two noise settings are summarized in Tab. 2. Compared to other methods, FedA<sup>3</sup>I achieves the best performance across different datasets and settings, validating its effectiveness in handling heterogeneous annotation noise. It confirms the necessity of specific designs for the two-level Non-IID annotation noise in FMIS. In addition, different from those methods introducing complex correction (Xu et al. 2022b; Liu et al. 2022; Zhu, Chen, and Yuan 2023) or regularization (Liu et al. 2020; Wu et al. 2023a), no additional training overhead is introduced in FedA<sup>3</sup>I, making it more easy-to-deploy in practice.

## 4.4 Ablation Study

To analyze the effectiveness of FedA<sup>3</sup>I, a series of in-depth ablation studies are conducted.

**Component-wise Ablation Study** IntraGW and InterGW are added to the quantity-based aggregation strategy (FedAvg) sequentially to validate whether quality-based factors

Loss	Component			Dataset			
				SKIN		BREAST	
	FedAvg	IntraGW	InterGW	Noise Type			
			$N_S$	$N_E$	$N_S$	$N_E$	
CE	✓			$63.85 \pm 2.41$	$75.69 \pm 0.56$	$65.33 \pm 1.03$	$68.49 \pm 0.67$
	✓	✓		$67.45 \pm 1.55$	$77.33 \pm 0.18$	$67.87 \pm 0.87$	$69.26 \pm 1.80$
	✓	✓	✓	<b><math>72.69 \pm 1.28</math></b>	<b><math>78.81 \pm 0.75</math></b>	<b><math>69.03 \pm 1.95</math></b>	<b><math>74.35 \pm 1.90</math></b>
DC	✓			$73.50 \pm 1.24$	$75.06 \pm 0.41$	$72.31 \pm 1.11$	$67.27 \pm 0.70$
	✓	✓		$74.59 \pm 1.03$	$76.70 \pm 0.52$	$74.49 \pm 0.94$	$68.31 \pm 0.46$
	✓	✓	✓	<b><math>75.92 \pm 0.99</math></b>	<b><math>77.05 \pm 0.56</math></b>	<b><math>75.28 \pm 0.53</math></b>	<b><math>76.68 \pm 0.49</math></b>

Table 3: Component-wise evaluation measured in the mean (%)  $\pm$  standard deviation (%) of Dice coefficients under different annotation noise settings. The best results are marked in bold.




Figure 4: Ablation studies on the balance coefficient  $r$  in InterGW. The solid line and transparent areas represent the mean and standard deviation respectively. *Second best* means the performance of the second-best method in Tab. 2.

can help aggregation with underlying noisy clients/models. Experiments are conducted following the settings in Sec. 4.3, and quantitative results are summarized in Tab. 3. FedAvg struggles with noisy annotations, as it measures the contribution of each client only based on its data amount, making the global model biased towards clients with large amounts of noisy data. To mitigate this, IntraGW is introduced to up-weight the clients with higher-quality annotations, leading to a performance boost. Further introducing InterGW to balance the two client groups  $\mathcal{G}_l$  and  $\mathcal{G}_s$  would help the model learn general knowledge instead of overfitting to the majority noise, achieving the best performance.

**Impact of the Balance Coefficient  $r$  in InterGW** In InterGW,  $r$ , set as 0.5 by default, is to balance the two client groups  $\mathcal{G}_l$  and  $\mathcal{G}_s$ . To evaluate the impact of  $r$ , ablation studies under various settings of  $r$  are conducted on SKIN with noise settings stated in Tab 1, and quantitative results are illustrated in Fig. 4. Comparison results indicate that  $r$  is a noise-related parameter, and InterGW can even fail without a proper setting. It is because  $r$  is to re-weight  $\mathcal{G}_l$  and  $\mathcal{G}_s$ , and the optimal weighting strategy can vary under different noise settings (e.g.  $N_S$  and  $N_E$ ). Fortunately, using the default value of 0.5 brings a stable boost for InterGW, leading to better performance than the second-best SOTA method.

**Impact of the Warm-up Round  $T_1$**  In FedA<sup>3</sup>I, we perform noise estimation with a warm-up global model trained by FedAvg. A natural question arises: how many training




Figure 5: Ablation studies on the warm-up round  $T_1$ . The solid line and transparent areas represent the mean and standard deviation respectively. *Second best* means the performance of the second-best method in Tab. 2.

rounds are needed for warm-up training? Therefore, ablation studies under different settings of  $T_1$  are conducted on SKIN as illustrated in Fig 5. FedA<sup>3</sup>I stably outperforms FedAvg and other SOTA methods with  $T_1$  in a certain range (i.e. 10-20), demonstrating that FedA<sup>3</sup>I is relatively robust to the selection of  $T_1$ . Hence, during the early stage of model learning (but not too early),  $T_1$  can be chosen relatively arbitrarily to reduce the workload of hyper-parameter tuning.

## 5 Conclusion

This paper presents a novel FMIS problem: training a segmentation model with decentralized data containing heterogeneous cross-client annotation noise. To formulate this problem, we first propose a general model called CEM to generate annotator-specific noise for pixel-wise annotations, and then extend it to construct a heterogeneous cross-client noise model for multi-source data. For problem-solving, we propose FedA<sup>3</sup>I for noise-robust FMIS. Specifically, a quality factor, calculated based on noise estimation, is incorporated into the model aggregation phase in a layer-wise manner to emphasize the contribution of high-quality clients. Experimental results on two publicly-available datasets demonstrate the superiority of FedA<sup>3</sup>I against the state-of-the-art approaches. We believe that the presented problem, the noise model, and the solution FedA<sup>3</sup>I will serve as valuable inspiration for future research in developing realistic medical image segmentation systems from decentralized data.

## Acknowledgments

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. We also thank Dr. Zhenyu Liao for his assistance in mathematical proof.

## References

- Al-Dhabyani, W.; Gomaa, M.; Khaled, H.; and Fahmy, A. 2020. Dataset of breast ultrasound images. *Data in brief*, 28: 104863.
- Bai, Y.; Yang, E.; Han, B.; Yang, Y.; Li, J.; Mao, Y.; Niu, G.; and Liu, T. 2021. Understanding and improving early stopping for learning with noisy labels. In *NeurIPS*, 24392–24403.
- Calisto, F. M.; Santiago, C.; Nunes, N.; and Nascimento, J. C. 2021. Introduction of human-centric AI assistant to aid radiologists for multimodal breast image classification. *Int. J. Hum. Comput. Stud.*, 150: 102607.
- Chen, P.; Ye, J.; Chen, G.; Zhao, J.; and Heng, P.-A. 2021. Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise. In *AAAI*, 11442–11450.
- Codella, N. C.; Gutman, D.; Celebi, M. E.; Helba, B.; Marchetti, M. A.; Dusza, S. W.; Kalloo, A.; Liopyris, K.; Mishra, N.; Kittler, H.; et al. 2018. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *ISBI*, 168–172. IEEE.
- Dou, Q.; So, T. Y.; Jiang, M.; Liu, Q.; Vardhanabhuti, V.; Kaissis, G.; Li, Z.; Si, W.; Lee, H. H.; Yu, K.; et al. 2021. Federated deep learning for detecting COVID-19 lung abnormalities in CT: a privacy-preserving multinational validation study. *npj Digit. Medicine*, 4(1): 60.
- Fang, C.; Wang, Q.; Cheng, L.; Gao, Z.; Pan, C.; Cao, Z.; Zheng, Z.; and Zhang, D. 2023. Reliable Mutual Distillation for Medical Image Segmentation under Imperfect Annotations. *IEEE Trans. Medical Imaging*.
- Gu, Z.; Cheng, J.; Fu, H.; Zhou, K.; Hao, H.; Zhao, Y.; Zhang, T.; Gao, S.; and Liu, J. 2019. CE-Net: Context encoder network for 2D medical image segmentation. *IEEE Trans. Medical Imaging*, 38(10): 2281–2292.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*.
- Huang, W.; Ye, M.; Shi, Z.; and Du, B. 2023a. Generalizable heterogeneous federated cross-correlation and instance similarity learning. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Huang, W.; Ye, M.; Shi, Z.; Wan, G.; Li, H.; Du, B.; and Yang, Q. 2023b. Federated Learning for Generalization, Robustness, Fairness: A Survey and Benchmark. *arXiv:2311.06750*.
- 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 *CVPR*, 12341–12351.
- Jiang, M.; Roth, H. R.; Li, W.; Yang, D.; Zhao, C.; Nath, V.; Xu, D.; Dou, Q.; and Xu, Z. 2023a. Fair Federated Medical Image Segmentation via Client Contribution Estimation. In *CVPR*, 16302–16311.
- Jiang, M.; Yang, H.; Cheng, C.; and Dou, Q. 2023b. IOP-FL: Inside-outside personalization for federated medical image segmentation. *IEEE Trans. Medical Imaging*.
- Jiang, X.; Sun, S.; Wang, Y.; and Liu, M. 2022. Towards federated learning against noisy labels via local self-regularization. In *CIKM*, 862–873.
- Li, J.; Socher, R.; and Hoi, S. C. 2020. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In *ICLR*.
- Li, S.; Gao, Z.; and He, X. 2021. Superpixel-guided iterative learning from noisy labels for medical image segmentation. In *MICCAI*, 525–535. Springer.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. In *MLSys*, 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. In *ICLR*.
- Liang, X.; Lin, Y.; Fu, H.; Zhu, L.; and Li, X. 2022. RSCFed: Random Sampling Consensus Federated Semi-supervised Learning. In *CVPR*, 10154–10163.
- Liu, Q.; Chen, C.; Qin, J.; Dou, Q.; and Heng, P.-A. 2021. FedDG: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *CVPR*, 1013–1023.
- Liu, S.; Liu, K.; Zhu, W.; Shen, Y.; and Fernandez-Granda, C. 2022. Adaptive early-learning correction for segmentation from noisy annotations. In *CVPR*, 2606–2616.
- Liu, S.; Niles-Weed, J.; Razavian, N.; and Fernandez-Granda, C. 2020. Early-learning regularization prevents memorization of noisy labels. In *NeurIPS*, volume 33, 20331–20342.
- Ma, X.; Zhang, J.; Guo, S.; and Xu, W. 2022. Layer-wised model aggregation for personalized federated learning. In *CVPR*, 10092–10101.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 1273–1282.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 565–571. Ieee.
- Rieke, N.; Hancox, J.; Li, W.; Milletari, F.; Roth, H. R.; Albarqouni, S.; Bakas, S.; Galtier, M. N.; Landman, B. A.; Maier-Hein, K.; et al. 2020. The future of digital health with federated learning. *npj Digit. Medicine*, 3(1): 119.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 234–241. Springer.

- Salehi, S. S. M.; Erdogmus, D.; and Gholipour, A. 2017. Tversky loss function for image segmentation using 3D fully convolutional deep networks. In *Machine Learning in Medical Imaging - 8th International Workshop*, 379–387. Springer.
- Valanarasu, J. M. J.; and Patel, V. M. 2022. UNeXt: MLP-based rapid medical image segmentation network. In *MIC-CAI*, 23–33. Springer.
- Wang, G.; Liu, X.; Li, C.; Xu, Z.; Ruan, J.; Zhu, H.; Meng, T.; Li, K.; Huang, N.; and Zhang, S. 2020. A noise-robust framework for automatic segmentation of COVID-19 pneumonia lesions from CT images. *IEEE Trans. Medical Imaging*, 39(8): 2653–2663.
- Wang, J.; Jin, Y.; Stoyanov, D.; and Wang, L. 2023. FedDP: Dual personalization in federated medical image segmentation. *IEEE Trans. Medical Imaging*.
- Wang, J.; Jin, Y.; and Wang, L. 2022. Personalizing federated medical image segmentation via local calibration. In *ECCV*, 456–472.
- Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric cross entropy for robust learning with noisy labels. In *CVPR*, 322–330.
- Wicaksana, J.; Yan, Z.; Zhang, D.; Huang, X.; Wu, H.; Yang, X.; and Cheng, K.-T. 2022. FedMix: Mixed Supervised Federated Learning for Medical Image Segmentation. *IEEE Trans. Medical Imaging*.
- Wu, N.; Yu, L.; Jiang, X.; Cheng, K.-T.; and Yan, Z. 2023a. FedNoRo: Towards Noise-Robust Federated Learning by Addressing Class Imbalance and Label Noise Heterogeneity. In *IJCAI*.
- Wu, N.; Yu, L.; Yang, X.; Cheng, K.-T.; and Yan, Z. 2023b. FedIIC: Towards Robust Federated Learning for Class-Imbalanced Medical Image Classification. In *MIC-CAI*.
- Xu, A.; Li, W.; Guo, P.; Yang, D.; Roth, H. R.; Hatamizadeh, A.; Zhao, C.; Xu, D.; Huang, H.; and Xu, Z. 2022a. Closing the generalization gap of cross-silo federated medical image segmentation. In *CVPR*, 20866–20875.
- Xu, J.; Chen, Z.; Quek, T. Q.; and Chong, K. F. E. 2022b. FedCorr: Multi-Stage Federated Learning for Label Noise Correction. In *CVPR*, 10184–10193.
- Xu, X.; Deng, H. H.; Gateno, J.; and Yan, P. 2023. Federated Multi-organ Segmentation with Inconsistent Labels. *IEEE Trans. Medical Imaging*.
- Yan, Z.; Wicaksana, J.; Wang, Z.; Yang, X.; and Cheng, K.-T. 2020. Variation-aware federated learning with multi-source decentralized medical image data. *IEEE J. Biomed. Health Informatics*, 25(7): 2615–2628.
- Yao, J.; Zhang, Y.; Zheng, S.; Goswami, M.; Prasanna, P.; and Chen, C. 2023. Learning to Segment from Noisy Annotations: A Spatial Correction Approach. In *ICLR*.
- Yap, M. H.; Pons, G.; Marti, J.; Ganau, S.; Sentis, M.; Zwiggelaar, R.; Davison, A. K.; and Marti, R. 2017. Automated breast ultrasound lesions detection using convolutional neural networks. *IEEE J. Biomed. Health Informatics*, 22(4): 1218–1226.
- Yu, X.; Han, B.; Yao, J.; Niu, G.; Tsang, I.; and Sugiyama, M. 2019. How does disagreement help generalization against label corruption? In *ICML*, 7164–7173. PMLR.
- Zhang, L.; Tanno, R.; Xu, M.-C.; Jin, C.; Jacob, J.; Ciccarelli, O.; Barkhof, F.; and Alexander, D. 2020a. Disentangling human error from ground truth in segmentation of medical images. In *NeurIPS*, volume 33, 15750–15762.
- Zhang, T.; Yu, L.; Hu, N.; Lv, S.; and Gu, S. 2020b. Robust medical image segmentation from non-expert annotations with tri-network. In *MICCAI*, 249–258. Springer.
- Zhang, Y.; Xian, M.; Cheng, H.-D.; Shareef, B.; Ding, J.; Xu, F.; Huang, K.; Zhang, B.; Ning, C.; and Wang, Y. 2022. BUSIS: a benchmark for breast ultrasound image segmentation. In *Healthcare*, volume 10, 729. MDPI.
- Zhang, Z.; and Sabuncu, M. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*.
- Zhu, H.; Shi, J.; and Wu, J. 2019. Pick-and-learn: Automatic quality evaluation for noisy-labeled image segmentation. In *MICCAI*, 576–584. Springer.
- Zhu, M.; Chen, Z.; and Yuan, Y. 2023. FedDM: Federated Weakly Supervised Segmentation via Annotation Calibration and Gradient De-conflicting. *IEEE Trans. Medical Imaging*.