



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.

FedMLP: Federated Multi-Label Medical Image Classification under Task Heterogeneity

Zhaobin Sun^{*1}, Nannan Wu^{*1}, Junjie Shi¹, Li Yu¹, Kwang-Ting Cheng², and Zengqiang Yan¹(✉)

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

{zbsun, wnn2000, shijunjie, hustlyu, z_yan}@hust.edu.cn

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

Abstract. Cross-silo federated learning (FL) enables decentralized organizations to collaboratively train models while preserving data privacy and has made significant progress in medical image classification. One common assumption is task homogeneity where each client has access to all classes during training. However, in clinical practice, given a multi-label classification task, constrained by the level of medical knowledge and the prevalence of diseases, each institution may diagnose only partial categories, resulting in task heterogeneity. How to pursue effective multi-label medical image classification under task heterogeneity is under-explored. In this paper, we first formulate such a realistic label missing setting in the multi-label FL domain and propose a two-stage method FedMLP to combat class missing from two aspects: pseudo label tagging and global knowledge learning. The former utilizes a warmed-up model to generate class prototypes and select samples with high confidence to supplement missing labels, while the latter uses a global model as a teacher for consistency regularization to prevent forgetting missing class knowledge. Experiments on two publicly-available medical datasets validate the superiority of FedMLP against the state-of-the-art both federated semi-supervised and noisy label learning approaches under task heterogeneity. Code is available at <https://github.com/szbonaldo/FedMLP>.

Keywords: Federated learning · Partial label · Multi-label classification and Task heterogeneity.

1 Introduction

Attributed to heightened privacy concerns, merging multiple medical image datasets into a unified one is often prohibited, presenting additional challenges in developing deep neural models for automated disease classification. Recently, federated learning (FL) [16], a technique that enables collaborative model training across decentralized sources without compromising privacy, has demonstrated

* Equal contribution.

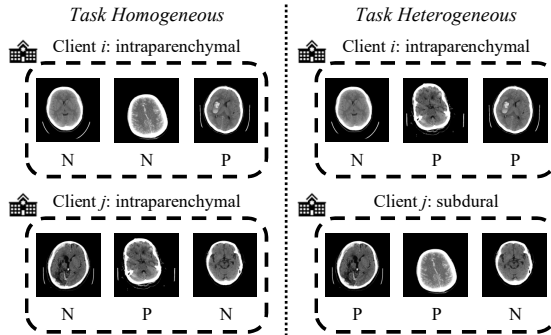


Fig. 1. Illustration of two decentralized data settings.

significant potential in overcoming this obstacle [26,11,15,10,8,24,23]. Nevertheless, current research predominantly operates under a critical assumption: cross-client data¹ is *task homogeneous* [7], as shown in Fig. 1. However, in practical data acquisition, a prevalent encountered scenario involves diverse institutions amassing data that is *task heterogeneous*, depending on their particular areas of interest. As depicted in Fig. 1, one hospital may prioritize intraparenchymal hemorrhage and thereby assemble a dataset. Conversely, another institution might gather data specifically for subdural hemorrhage reflecting its distinct focal points. This diversity in class interests and tasks among various clients introduces extra complexity to implementing FL, as the knowledge encapsulated by available labels is not universally shared across participants.

In this study, we conceptualize each client’s task as a multi-label classification problem characterized by partial annotations, reflecting that the areas of interest for certain clients represent a subset of the collective focus areas across all clients. Against this backdrop, our research is directed at tackling the critical challenge: *how to achieve data-efficient federated multi-label learning (FMLL) against heterogeneity in partially labeled classes?* Given diverse imperfect data caused by this heterogeneity, approaches in federated noisy label learning (FNLL) [25,27,12,22] and federated semi-supervised learning (FSSL) [13,14,15,3] fall short of effectively addressing it. Their designs, predominantly aimed at handling imperfect data in a homogeneous manner and focusing on multi-class classification, are sub-optimal when applied to FMLL scenarios. Thus, navigating this challenge remains a relatively untapped area of exploration.

To tackle this problem, the crux is to transfer knowledge from diverse tasks to each local client. We propose a prototype-based, two-stage method, named **Federated Multi-Label learning with Partial annotation (FedMLP)**. The initial stage involves a model warm-up to establish high-quality class prototypes. It is trained by a multi-label Weighted-Partial-Class (WPC) loss, combined with logit adjustment [17], to minimize the misclassification of false negatives in negative classes and mitigate the effects of class imbalance. Following the creation

¹ Data includes images and corresponding labels.

of initial class prototypes, the subsequent stage employs a dynamic Self-adaptive Threshold (ST) to select pseudo-labeling classes of samples whose features are more similar to class-wise prototypes, thus progressively facilitating task-specific knowledge transfer to each client, guided by further training under pseudo labels. To address the challenge of classes lacking pseudo labels, Consistency Regularization (CR) between global and local models is employed to diminish the potential negative effects of inaccurate knowledge within these classes.

The main contributions are summarized as follows. (1) An rarely-explored FMLL setting considering task heterogeneity. (2) A novel solution FedMLP to address heterogeneity in partially labeled classes. (3) Superior performance on two real-world medical datasets against SOTA FSSL and FNLL methods.

2 Methodology

2.1 Preliminaries and Overview

Given K participants, any k -th participant holds a dataset $D_k = \{(x_i, y_i)\}_{i=1}^{N_k}$ with N_k samples. Here, (x_i, y_i) is an image-label pair, with $x_i \in \mathcal{X} \subseteq \mathbb{R}^d$ and $y_i \in \mathcal{Y} = \{0, 1\}^C$, where C is the total number of classes. In scenarios with missing labels, each client k might not have labels for certain classes. To describe this, we define an active class set AC_k and a negative class set NC_k , with AC_k containing indices of classes with labels and NC_k comprising those without, ensuring $AC_k \cap NC_k = \emptyset$ and $AC_k \cup NC_k = [C]$. In the following context, active classes and negative classes are equivalent to the categories of labeled and missing classes respectively, and can be used interchangeably. Without data privacy leakage, we can get information of client-wise class-level annotation distribution $\mathbb{S} = \{\mathbb{S}_i\}_{i=1}^C$ at the server side, where \mathbb{S}_i denotes the set of client indexes recognizing class i , $\mathbb{S}_i \neq \emptyset$. In our setting, due to diversity in class interest, $|\mathbb{S}_i|$ is not identical for each class. A larger $|\mathbb{S}_i|$ implies a more common disease i , which is easier to learn than other classes, denoted as "hot classes", otherwise as "cool classes". The objective is to develop a global model capable of effectively detecting both hot and cool classes. The overview of FedMLP is illustrated in Fig. 2.

2.2 Warm-up with Missing Labels

Different from the typical FSSL scenario where at least one client or server has instances with complete labels, in FMLL with partial labels, from the perspective of each client, the only way to learn knowledge of local-unknown classes is from other clients. Previous works regard missing classes as negative [1,2,19,20] and use a binary cross entropy(BCE) loss on the multi-label setting formulated as

$$L_{BCE}(x_i, y_i) = -\frac{1}{C} \sum_{j=1}^C (y_i^j \cdot \log(\hat{y}_i^j) + (1 - y_i^j) \cdot \log(1 - \hat{y}_i^j)), \quad (1)$$

where \hat{y}_i^j denotes the predicted probability of class j by the Sigmoid function of sample i . Trained through this, clients tend to predict negatively if a label

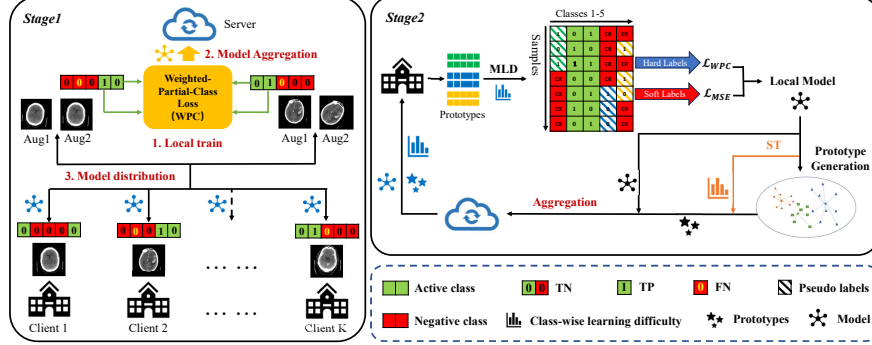


Fig. 2. Overview of FedMLP. Colors red and green represent missing and labeled classes respectively.

is missing, caused by the biased classifier and dense features, making the global model struggle to recognize cool classes completely. Though hot-class knowledge may be compensated by model aggregation, there is a risk of catastrophic forgetting when data distribution is quite heterogeneous.

Considering the different prevalences of diseases, BCE loss may make the model biased towards updating the majority classes. Related works like Logit Adjustment (LA) [17] were proposed for multi-class classification, fine-tuning model output by calculating the positive rates of each category and mitigating the impact of class imbalance, formulated as

$$l'_c = l_c + \tau \cdot \log(\pi_c), \quad (2)$$

where l_c denotes the logit of class c , τ is a hyper-parameter set as 1 by default, and π_c is the sample proportion of class c in the training set. In this work, we further extend LA to the multi-label setting:

$$\hat{y}_i^{c'} = \frac{\hat{y}_i^c \cdot \pi_{c,1}}{\hat{y}_i^c \cdot \pi_{c,1} + (1 - \hat{y}_i^c) \cdot \pi_{c,0}}, \quad (3)$$

where $\pi_{c,1}$ represents the proportion of positive samples in class c , and the sum of $\pi_{c,1}$ and $\pi_{c,0}$ equals to 1. Considering partial annotation makes the local model lazy to predict missing classes as negative, thus we propose partial-class loss to exclude the disturbance of false-negative labels, forming a Weighted-Partial-Class loss

$$L_{WPC}^k(x_i, y_i) = -\frac{1}{C} \sum_{j \in AC_k} ((1 - y_i^j) \cdot \log(1 - \hat{y}_i^{j'}) + y_i^j \cdot \log(\hat{y}_i^{j'})). \quad (4)$$

At the first t_1 communication rounds, we take L_{WPC}^k as the objection of local optimization and continue to use FedAvg for model aggregation. To enhance sample diversity, we further utilize data augment techniques, inspired by semi-supervise learning [18]. To this end, we get a warmed-up model to generate the representation F_i^k of sample x_i in client k with acceptable quality.

2.3 Missing Label Detection with Global Active Class Prototypes

Toward the phenomenon where traditional training methods tend to ignore cool classes under class interest diversity (*see the supplementary material*), we use prototypes to detect and choose confident samples for missing classes (Stage 2), enabling each client to gradually learn the missing knowledge of local samples.

Prototype Generation. For each client k , we only consider local active classes AC_k for prototype calculation to improve feature quality. Each class c has dual prototypes: $P_0^{k,c}$ and $P_1^{k,c}$ denoting local prototypes of negative and positive for class c in client k , updated at the end of each round after warm-up, defined as

$$P_n^{k,c} = \frac{\sum_{i=1}^{N_k} \mathbb{I}(y_i^c = n) \cdot F_i^k}{|D_{k,y_i^c=n}|}, n \in \{0, 1\} \quad (5)$$

where $c \in AC_k$ and $|D_{k,y_i^c=n}|$ is the number of samples being positive in class c . Each client only delivers its local prototypes and updated model to the server, maintaining the same communication cost as FedAvg. Then, the server aggregates clients' models and prototypes and sends the global prototypes of all classes and the model to clients. Global active class prototypes are updated by

$$P_n^c = \frac{\sum_{i=1}^K \mathbb{I}(i \in \mathbb{S}_c) \cdot P_n^{i,c}}{|\mathbb{S}_c|}, n \in \{0, 1\}. \quad (6)$$

According to previous research[9], taking the average of multiple iterations of features can effectively reduce the privacy leakage.

Missing Label Detection (MLD). Typical FNLL clustered sample losses to recognize noisy clients and samples [25,27] and used pseudo labeling for correction. To some extent, our scenario can be considered a special case with an explicit noise transition matrix, which completely turns missing classes to negative. Utilizing global active class prototypes, we already have the confidence of any instance i for a local negative class c in client k

$$Z_i^{k,c} = \mathbf{cos}(P_0^c, F_i^k) - \mathbf{cos}(P_1^c, F_i^k), \quad (7)$$

where \mathbf{cos} is cosine similarity. $Z_i^{k,c} < 0$ indicates the sample is more likely to be positive in class c , otherwise more likely to be negative. $|Z_i^{k,c}|$ approaches 0 for uncertain samples. Following this, we propose a novel criterion to select high-confidence samples for pseudo labeling, according to

$$\mathbb{P}_{c,0}^k = \{i \mid Z_i^{k,c} > \mathbf{Rank}(\mathbb{Z}^{k,c}, \tau_0), Z_j^{k,c} \in \mathbb{Z}^{k,c} \geq 0\}, \quad (8)$$

$$\mathbb{P}_{c,1}^k = \{i \mid -Z_i^{k,c} > \mathbf{Rank}(-\mathbb{Z}^{k,c}, \tau_1), Z_j^{k,c} \in \mathbb{Z}^{k,c} < 0\}, \quad (9)$$

where $\mathbf{Rank}(\mathbb{Z}^{k,c}, \tau)$ denotes the set of top τ percent elements $\mathbb{Z}^{k,c}$. If a sample is in $\mathbb{P}_{c,0}^k$ or $\mathbb{P}_{c,1}^k$, it will be permanently tagged with the label 0 or 1. In later

rounds, $\mathbb{Z}^{k,c}$ only includes residual samples until all samples have pseudo labels. τ_0 and τ_1 denote the selection ratios of samples with pseudo labels to alleviate erroneous information. Given more difficult classes to learn, sample selection is more discreet indicating a smaller τ . To automate this, we propose a self-adaptive ratio mechanism as discussed in the following.

Self-adaptive Threshold. The way to determine the class learning difficulty is crucial to adjusting the sample selection ratio automatically. In the local view, as the active classes are deemed to have correct labels, it may reduce the difficulty gap between classes. Thus, we evaluate negative class difficulty in each client using its local data and average them on the server side at the end of each communication round, via uncertainty estimation:

$$d_c^k = \frac{\sum_{i=1}^{N_k} \mathbb{I}(\hat{y}_i^c < L \cup \hat{y}_i^c > R)}{|D_k|}, \text{ and } d_c^G = \sum_{k \in \mathbb{S}_c} \frac{|D_k|}{\sum_{i \in \mathbb{S}_c} |D_i|} \cdot d_c^k, \quad (10)$$

where d_c^k and d_c^G denote the learning degree of class c in client k and the global learning degree respectively, L, R are hyper-parameters. When a disease is easy to recognize, the numerator of d_c^k increases, implying more samples can be selected to train with pseudo labels. Then, τ_0 and τ_1 are formulated as $\tau_0 = d_c^G \cdot T_0$ and $\tau_1 = d_c^G \cdot T_1$ where hyper-parameters T_0 and T_1 controls the selected ratios of negative and positive samples.

2.4 Consistency Regularization of Uncertain Classes

The training process will become unstable if only partial classes are updated, *e.g.*, knowledge quickly fades away when gradients are not computed for a minority category. Thus, we further utilize the global model as a teacher to avoid clients forgetting the knowledge about negative classes without pseudo labels. In this way, each sample will be trained by at most three labels including the correct active class label, the pseudo label, and the soft label prompted by the global model. We use the weighted-partial-class loss mentioned in Sec. 2.2 for the first two labels and the mean-squared (MSE) loss for the soft labels.

3 Experiments

3.1 Experimental Setup

Dataset. Two multi-label medical image datasets are adopted. *ICH*: The RSNA Intracranial Hemorrhage dataset [5], with over 800,000 CT images, is condensed to a subset of 180,000 instances by mixing all lesion samples and partial healthy samples. From this, 25,000 samples are further used for training, with the rest designated as the test set. *ChestXray14*: The ChestXray14 dataset [21] contains chest X-ray images from over 30,000 patients, of which PA-view samples are selected and split into training and testing sets in a 7:3 ratio. Both datasets are pre-processed following protocols in [15,25], with an equal distribution of training data among clients.

Table 1. Quantitative comparison on *ICH* under four levels of partial annotation settings. The results (%) based on the last-epoch model on the testing set are reported. The best and second-best results are marked in bold and underlined.

Type	Method	Missing 1 class			Missing 2 classes			Missing 3 classes			Missing 4 classes			Average		
		BACC	AUC	mAP	BACC	AUC	mAP	BACC	AUC	mAP	BACC	AUC	mAP	BACC	AUC	mAP
Base	FedAvg	74.58	90.65	70.32	66.13	88.85	64.76	58.64	85.43	54.53	50.00	76.32	32.62	62.34	85.31	55.56
	FedAvg-PL	80.52	91.32	70.94	76.12	87.88	58.98	63.40	83.12	44.18	61.46	79.48	40.60	70.38	85.45	53.68
FSSL	RSCFed	80.26	90.83	70.79	79.09	89.24	67.27	67.55	85.40	61.43	67.16	71.16	36.89	73.52	84.16	59.10
	FedFixMatch	80.75	<u>91.52</u>	<u>71.03</u>	80.75	90.19	69.10	66.26	77.75	46.35	50.00	74.90	31.05	69.44	83.59	54.38
	FedIRM	79.77	90.93	70.63	72.78	87.12	63.17	66.10	83.92	49.68	69.39	77.19	36.73	72.01	84.79	55.05
	CBAFed	76.13	91.43	70.58	71.97	87.50	63.27	55.66	74.85	43.94	52.28	72.15	34.99	64.01	81.48	53.20
FNL	FedLSR	75.81	91.28	70.52	65.48	89.04	66.76	56.67	84.62	52.70	50.00	74.14	29.73	61.99	84.77	54.93
	FedNoRo	<u>82.47</u>	90.92	70.88	<u>81.90</u>	<u>91.12</u>	<u>69.53</u>	<u>80.98</u>	90.99	65.73	<u>74.64</u>	<u>85.09</u>	<u>59.28</u>	<u>80.00</u>	<u>89.53</u>	<u>66.36</u>
Ours	FedMLP	83.50	92.84	73.80	82.26	91.75	70.52	81.10	<u>90.87</u>	<u>65.53</u>	80.07	88.47	59.31	81.73	90.98	67.29

Partial Label Generation. We randomly remove an equal number of categories from each client, ensuring that no client possesses a complete set of labels. *ICH* comprises five subtypes, whereas *ChestXray14* includes 14 categories. Our focus narrows to the eight categories with the highest incidence of positive cases. To validate the performance in dealing with extreme label scarcity where each client labels only one category, we adjust the number of clients to match the number of categories, resulting in five clients for *ICH* and eight for *ChestXray14*.

Implementation Details. For both datasets and all comparison methods, we utilized ResNet-18 [6] pre-trained on ImageNet[4] as the backbone, the Adam optimizer with a learning rate of 3e-5, L_2 regularization on model parameters, and a batch size of 32. In the federated setting, we configured each client to conduct one local training round before each communication round. The warm-up round t_1 was set at 50, and the total number of rounds was 500. The parameters L and R were set to 0.3 and 0.7, respectively, with T_0 and T_1 fixed at 0.5% and 1%. All clients participate in aggregation during each communication round. More details of comparative methods can be found in appendix.

3.2 Comparison with State-of-the-art Methods

FedAvg [16] and its variant coupled by a partial loss to include only active classes (denoted as FedAVG-PL) are selected as baseline approaches. By regarding missing sample labels as absent class labels, FSSL is extendable to such a setting. Thus, SOTA FSSL methods are introduced for comparison, including RSCFed (CVPR’22) [14] using a sub-consensus model and local distillation for unknown classes, FedFixMatch by extending FixMatch (NeuIPS’20) [18] with consistency regularization through weak and strong augmentation to FedAvg, FedIRM (MICCAI’21) [15] exploring relational class knowledge and aligning clients, and CBAFed (CVPR’23) [13] implementing pseudo-labeling with dynamic thresholds. Our scenario can also be regarded as a setting where all positive samples with unlabeled classes are mislabeled as negative due to noise, addressed by FNNL. SOTA FNNL methods are introduced, including FedLSR

Table 2. Quantitative comparison on *ChestXray14* under four levels of partial annotation settings. The results (%) based on the last-epoch model on the testing set are reported. The best and second-best results are marked in bold and underlined.

Type	Method	Missing 1 class			Missing 3 classes			Missing 5 classes			Missing 7 classes			Average		
		BACC	AUC	mAP	BACC	AUC	mAP	BACC	AUC	mAP	BACC	AUC	mAP	BACC	AUC	mAP
Base	FedAvg	63.27	76.46	26.89	57.21	75.45	23.31	51.98	70.21	18.65	50.00	66.64	12.16	55.62	72.19	20.25
	FedAvg-PL	63.88	76.02	27.37	60.14	73.28	21.38	56.87	69.33	19.35	53.30	64.80	11.49	58.55	70.86	19.90
FSSL	RSCFed	65.07	77.08	26.61	63.30	76.37	<u>26.41</u>	58.87	70.81	23.96	50.49	65.92	14.70	59.43	72.55	22.92
	FedFixMatch	62.90	77.94	27.70	58.07	75.48	25.68	58.02	69.53	19.65	54.52	65.59	11.42	58.38	72.14	21.11
	FedIRM	61.62	76.03	26.99	57.84	73.91	23.67	54.68	67.79	20.98	50.01	63.20	13.44	56.04	70.23	21.27
FNL	CBAFed	<u>68.77</u>	77.33	26.95	62.98	74.28	24.02	52.97	68.26	19.67	50.00	52.25	7.53	58.68	68.03	19.54
	FedLSR	62.59	76.95	<u>28.25</u>	55.16	74.88	25.25	52.94	70.52	19.65	50.00	50.00	6.79	55.17	68.09	19.99
	FedNoRo	69.65	78.99	27.96	69.27	78.73	26.26	<u>67.68</u>	<u>74.75</u>	<u>25.97</u>	<u>66.54</u>	<u>74.23</u>	<u>20.31</u>	<u>68.29</u>	<u>76.68</u>	<u>25.13</u>
Ours	FedMLP	68.55	<u>78.44</u>	28.87	<u>68.37</u>	<u>78.51</u>	28.01	68.62	76.02	26.63	70.08	75.76	22.48	68.91	77.18	26.50

Table 3. Component-wise ablation study.

FedAvg	MLD	WPC	CR	ST	<i>ICH</i>			<i>ChestXray14</i>		
					BACC	AUC	mAP	BACC	AUC	mAP
✓					50.00	76.32	32.62	50.00	66.64	12.16
✓	✓				73.41	82.32	46.74	62.82	69.80	13.75
✓	✓	✓			78.26	85.98	52.10	66.63	72.59	16.37
✓	✓	✓	✓		79.73	88.44	58.81	68.85	75.06	19.56
✓	✓	✓	✓	✓	80.07	88.47	59.31	70.08	75.76	22.48

(CIKM’22) [12] via self-regularization during local training and FedNoRo (IJ-CAI’23) [25] detecting noisy clients with Gaussian mixture model of which only the second stage is trained as all clients are noisy in our setting.

Quantitative comparison measured by balanced accuracy (BACC, for multi-class) and AUC along with mAP (for multi-label) are summarized in Tables 1 and 2. With increasing class absences, existing FL approaches struggle particularly with less prevalent “cool classes”. Specifically, FedAvg and FedLSR fail completely with only one class labeled, defaulting to classify all samples as healthy. Comparatively, FedMLP leverages class prototypes to extract knowledge from each client, consistently achieving top or near-top results on two medical image datasets and maintaining strong performance across varying levels of label absence.

3.3 Ablation Study

Component-wise ablation study of FedMLP on *ICH* and *ChestXray14* is summarized in Table 3. Each client was configured to focus on a single category, highlighting the contribution of each component under conditions of extreme label scarcity. In such scenarios, FedAvg struggles to detect positive cases, resulting in a decline in BACC down to 50%. MLD, utilizing class prototypes, empowers clients to access information on unlabeled classes, mitigating the knowledge loss typically incurred during model aggregation. Integrating the remaining three components further improves model efficacy across all evaluation metrics.

4 Conclusion

In this paper, we propose a novel federated multi-label partial label setting and identify the catastrophic forgetting problem of "cool classes" with the imbalance of class attention. Our method FedMLP, which combines class prototype generation and global knowledge regularization, enables the majority of clients to leverage their local unlabeled data to enhance limited labeled information. Experimental results demonstrate its superior performance and robustness to high missing rates on two real-world medical datasets. We believe this study is helpful in building real-world FL systems under complicated data/task heterogeneity.

Acknowledgment 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. Bucak, S.S., Jin, R., Jain, A.K.: Multi-label learning with incomplete class assignments. In: CVPR. pp. 2801–2808. IEEE (2011)
2. Chen, M., Zheng, A., Weinberger, K.: Fast image tagging. In: ICML. pp. 1274–1282. PMLR (2013)
3. Cho, Y.J., Joshi, G., Dimitriadis, D.: Local or global: Selective knowledge assimilation for federated learning with limited labels. In: CVPR. pp. 17087–17096 (2023)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255. Ieee (2009)
5. Flanders, A.E., Prevedello, L.M., Shih, G., Halabi, S.S., Kalpathy-Cramer, J., Ball, R., Mongan, J.T., Stein, A., Kitamura, F.C., Lungren, M.P., et al.: Construction of a machine learning dataset through collaboration: the rsna 2019 brain ct hemorrhage challenge. *Radiology: Artificial Intelligence* **2**(3), e190211 (2020)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
7. Huang, W., Ye, M., Du, B.: Learn from others and be yourself in heterogeneous federated learning. In: CVPR (2022)
8. Huang, W., Ye, M., Shi, Z., Du, B.: Generalizable heterogeneous federated cross-correlation and instance similarity learning. *IEEE TPAMI* (2023)
9. Huang, W., Ye, M., Shi, Z., Li, H., Du, B.: Rethinking federated learning with domain shift: A prototype view. In: CVPR (2023)
10. 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)
11. Jiang, M., Yang, H., Li, X., Liu, Q., Heng, P.A., Dou, Q.: Dynamic bank learning for semi-supervised federated image diagnosis with class imbalance. In: MICCAI. pp. 196–206 (2022)

12. Jiang, X., Sun, S., Wang, Y., Liu, M.: Towards federated learning against noisy labels via local self-regularization. In: CIKM. pp. 862–873 (2022)
13. Li, M., Li, Q., Wang, Y.: Class balanced adaptive pseudo labeling for federated semi-supervised learning. In: CVPR. pp. 16292–16301 (2023)
14. Liang, X., Lin, Y., Fu, H., Zhu, L., Li, X.: Rscfed: Random sampling consensus federated semi-supervised learning. In: CVPR. pp. 10154–10163 (2022)
15. Liu, Q., Yang, H., Dou, Q., Heng, P.A.: Federated semi-supervised medical image classification via inter-client relation matching. In: MICCAI. pp. 325–335 (2021)
16. 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)
17. Menon, A.K., Jayasumana, S., Rawat, A.S., Jain, H., Veit, A., Kumar, S.: Long-tail learning via logit adjustment. In: ICLR (2021)
18. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: NeurIPS. vol. 33, pp. 596–608 (2020)
19. Sun, Y.Y., Zhang, Y., Zhou, Z.H.: Multi-label learning with weak label. In: AAAI. vol. 24, pp. 593–598 (2010)
20. Wang, Q., Shen, B., Wang, S., Li, L., Si, L.: Binary codes embedding for fast image tagging with incomplete labels. In: ECCV. pp. 425–439. Springer (2014)
21. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: CVPR. pp. 2097–2106 (2017)
22. Wu, C., Li, Z., Wang, F., Wu, C.: Learning cautiously in federated learning with noisy and heterogeneous clients. arXiv preprint arXiv:2304.02892 (2023)
23. 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)
24. Wu, N., Sun, Z., Yan, Z., Yu, L.: Fed3i: Annotation quality-aware aggregation for federated medical image segmentation against heterogeneous annotation noise. In: AAAI (2024)
25. 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. IJCAI (2023)
26. 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)
27. 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)