Enhancing Federated Learning Performance Fairness via Collaboration Graph-based Reinforcement Learning

Yuexuan Xia[1⋆], Benteng Ma[1⋆], Qi Dou[2], and Yong Xia[1,3,4✉]

[1] National Engineering Laboratory for Integrated Aero-Space-Ground-Ocean Big Data Application Technology, School of Computer Science and Engineering, Northwestern Polytechnical University, Xi'an 710072, China
[2] Department of Computer Science and Engineering, The Chinese University of Hong Kong, Sha Tin, Hong Kong
[3] Research & Development Institute of Northwestern Polytechnical University in Shenzhen, Shenzhen 518057, China
[4] Ningbo Institute of Northwestern Polytechnical University, Ningbo 315048, China
yxia@nwpu.edu.cn

Abstract. Federated learning has recently developed into a pivotal distributed learning paradigm, wherein a server aggregates numerous client-trained models into a global model without accessing any client data directly. It is acknowledged that the impact of statistical heterogeneity in client local data on the pace of global model convergence, but it is often underestimated that this heterogeneity also engenders a biased global model with notable variance in accuracy across clients. Contextually, the prevalent solutions entail modifying the optimization objective. However, these solutions often overlook implicit relationships, such as the pairwise distances of site data distributions, which makes pairwise exclusive or synergistic optimization among client models. Such optimization conflicts compromise the efficacy of earlier methods, leading to performance imbalance or even negative transfer. To tackle this issue, we propose a novel aggregation strategy called Collaboration Graph-based Reinforcement Learning (FedGraphRL). By deploying a reinforcement learning (RL) agent equipped with a multi-layer adaptive graph convolutional network (AGCN) on the server-side, we can learn a collaboration graph from client state vectors, revealing the collaborative relationships among clients during optimization. Guided by an introduced reward that balances fairness and performance, the agent allocates aggregation weights, thereby promoting automated decision-making and improvements in fairness. The experimental results on two real-world multi-center medical datasets suggest the effectiveness and superiority of the proposed FedGraphRL.

Keywords: Federated learning · Performance fairness · Collaboration graph · Reinforcement learning.

⋆ Yuexuan Xia and Benteng Ma contributed equally to this work.

1 Introduction

The advent of deep learning has profoundly transformed the domain of medical image analysis, primarily due to the proliferation of large-scale clinical data sets [17,26,23,19]. However, privacy concerns from patients and institutions prevent centralized training from accessing the data across multiple centers [9,14]. Federated learning (FL) [20,4] thereby emerges as an effective solution that preserves privacy by distributing the model to data sources to train a global model without sharing their data directly. Nevertheless, conventional FL methods target minimizing an aggregated loss, potentially leading to biased model performance where accuracy loss may be borne unequally among clients [21,15] due to the statistical heterogeneity.

Fair federated learning (FFL) has become an important area of research to address the problem of client bias mentioned above. In FFL, the collection of gradients is a key way to balance the principles of utilitarianism and egalitarianism. Current FFL methods can be categorized into three groups. ***Gradient-based*** methods [27,7] identify an improved global gradient by minimizing local gradient discrepancy among clients. Although effective, these methods cannot guarantee the prevention of local performance decline when more than two clients are involved. ***Objective function-based*** methods reweight and minimize the aggregate loss according to the client's loss magnitude. [21] utilizes the min-max approach to safeguard the least performing clients, at the expense of compromising the adaptability of the optimization goal and consequently reducing overall performance. [15,13,28,6] propose a more flexible optimization objective that can be adjusted according to the desired fairness level. Unfortunately, these strategies neglect the implicit relationships among diverse clients, such as pairwise distances of site data distributions [2], which makes the optimization of client models either mutually exclusive or favorable. Even when loss magnitudes are nearly balanced (objective function-based methods fail), optimization conflicts can still result in performance unfairness or even negative transfer, thereby compromising optimization effectiveness. ***Client contribution estimation-based*** methods [8] assess the relative contributions of individual clients to determine aggregation weights. Nevertheless, these one-to-many constraint relations do not fully capture the pairwise implicit relationships among clients.

Recognizing that client relationships are highly interdependent, we suggest modeling the pairwise constraint relationships between clients using a collaboration graph. Utilizing this graph-based design rationale, we propose a novel aggregation strategy for FFL - the collaboration graph-based reinforcement learning (**FedGraphRL**). This strategy seeks to guide managing model optimization conflicts among clients in FFL through optimizing the collaboration graph. A reinforcement learning (RL) agent, reinforced with a multi-layer adaptive graph convolutional neural network (AGCN), is deployed server-side to learn a collaboration graph from client state vectors. Nodes in this graph represent clients, while edge weights indicate the intensity of collaboration between two clients. A reward function is also introduced to strike a balance between model performance and fairness. Using this function, the proposed FedGraphRL agent ex-

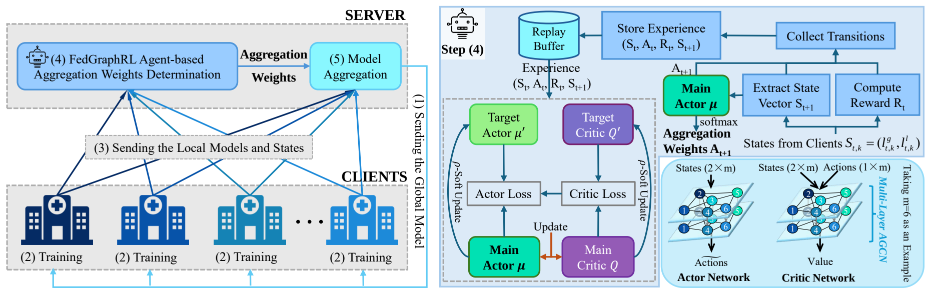

(a) Workflow of FedGraphRL

(b) Details of the FedGraphRL Agent-based Aggregation Weights Determination

Fig. 1. The framework of (a) the federated learning system which has five stages, and (b) FedGraphRL agent with multi-layer AGCN on the server side. Actions are actions predicted by the actor before softmax normalization.

tracts pertinent features from the collaboration graph to automatically allocate aggregation weights. To verify our approach, we conducted comparative experiments on two real-world multi-center medical image datasets. The experiment results demonstrate the superiority of the proposed model and the effectiveness of each component.

Our contributions are three-fold. (1) We introduce a collaboration graph, which represents the pairwise relationships among clients, to map out each client's collaborative relationships and enhance optimization outcomes. (2) We employ a learning-based optimization algorithm to allocate aggregation weights and balance performance and fairness. (3) Our experimental results indicate that the proposed method, which avoids compromising the performance of certain clients, could be a valuable asset to encourage more medical centers to engage in FL research.

2 Method

2.1 Problem Formalization and Method Overview

Given an FL system that involves m hospitals/clients, the i-th client has a local dataset $D_i = \{x_i^j, y_i^j\}_{j=1}^{n_i}$ sampled from a distribution P_i, where x_i^j and y_i^j represent the input and label, respectively. Due to the data heterogeneity across different medical centers, the data distributions $\{P_i\}_{i=1}^{m}$ are non-independent identically distribution (non-iid). Let F_k denote the task loss function of the client k, and $\varphi(\cdot)$ is a function aggregating local objectives. A fair FL system aims to minimize the objective function $\min_w \varphi(F_1(w), ...F_m(w))$ while ensuring that the global model w provides a more equitable solution by achieving uniform performance across all m clients. To tackle this optimization problem, our goal is to clarify the optimization conflict relationships between clients through a learnable collaboration graph, thereby devising a more equitable and precise aggregation strategy $\varphi(\cdot)$. The framework of FedGraphRL contains two components: a learning-based adaptive aggregation mechanism, and an AGCN component that embeds a collaboration graph into the optimization loop (see Fig. 1). We now delve into the details of each component.

2.2 Learning-based Adaptive Aggregation Mechanism

RL agents have been extensively utilized in FL for tasks such as client selection [29] and resource allocation [25], aiming for automated decision-making across various FL environments. We are dedicated to proposing a training method for agents that balances performance and fairness in FL. An off-policy actor-critic RL approach DDPG [16] is adopted, where the actor seeks the optimal aggregation strategy, and the critic evaluates the quality of the aggregation strategy.

State Space. In the topological graph G, which includes m clients, the system's state is a tuple of $2 \times m$ parameters. Specifically, the state of client k at round t is denoted as $S_{t,k} = (l^g_{t,k}, l^l_{t,k})$, where the ensemble of states at round t, S_t, consists of $\{S_{t,1}, S_{t,2}, \ldots, S_{t,m}\}$. $l^g_{t,k}$ denotes the k^{th} client's validation set loss of the global model distributed by the server in round t, reflecting the inter-client data correlation and the global model's capacity to adjust to each client's data. $l^l_{t,k}$ signifies the validation set loss of the k^{th} client after local training in round t, which reflects each client's quality of the data and the model performance.

Action Space. We utilize a continuous action space to determine the magnitude of aggregation weights. The action A_t is a tuple of m parameters, $A_t = \{\alpha_{t,1}, ..., \alpha_{t,m}\}$. Given the actor network output $a_{t,k}$ and a Gaussian noise term $\mathcal{N}$ for balancing exploration and exploitation [16], the agent then determines the aggregation weight $\alpha_{t,k}$ for client k as

$$\alpha_{t,k} = Softmax(a_{t,k} + \mathcal{N}). \tag{1}$$

Reward. The ultimate goal of the agent is to maximize the total reward $\mathscr{G} = \sum_{t=1}^{T} \gamma^{t-1} R_t$, where γ is the discount factor, and R_t is the immediate reward given by the FL system after the aggregation in round t. Following the principle of simplicity in reward design, we define R_t as

$$R_t = \frac{L_{t-1} - L_t}{L_{base}}. \tag{2}$$

where R_t is calculated as the aggregate loss difference between rounds $t-1$ and t, normalized by the aggregate loss of the initial round L_{base}. The aggregate loss L_t can evaluate both the system's performance and fairness simultaneously, formulated as

$$L_t = \sum_{k=1}^{m} \frac{p_k}{q+1} l^{q+1}_{t,k}(w), \tag{3}$$

where p_k indicates the data quantity proportion of client k within the entire FL system, $l_{t,k}(\cdot)$ corresponds to the state component $l^g_{t,k}$. Following [15], we use q to adjust the fairness level. Our approach is flexible and can be integrated with the objective functions of mainstream FFL to configure the reward function.

Algorithm 1 Overall Flow of FedGraphRL.

Input: A server with an w_0-parameterized global model, learning rate η, local step κ, m clients with separated datasets, maximum communication round T, θ-parameterized main critic network Q and ϕ-parameterized main actor network μ, target networks Q', μ' as copies of Q and μ, replay buffer P, warm-up rounds W, soft main-target update factor ρ;

for round $t = 0$ to $T - 1$ **do**
  Server broadcasts global model w_t
  Local:
  for client $k = 1$ to m in parallel **do**
    $w_{t,k} \leftarrow w_t$
    $l^g_{t,k} \leftarrow$Inference loss with $w_{t,k}$
    Local training in κ epochs:
    $w_{t,k} \leftarrow w_{t,k} - \eta \nabla F_i(w_{t,k})$
    $l^l_{t,k} \leftarrow$ Inference loss with $w_{t,k}$
    Send $S_{t,k} = (l^g_{t,k}, l^l_{t,k})$ and $w_{t,k}$ to the server
  end for
  Server:
  Concatenates $S_{t,k}$ to form S_t
  Computes reward R_{t-1}
  Stores $(S_{t-1}, A_{t-1}, R_{t-1}, S_t)$ into P
  $A_t \leftarrow$ ***ServerActionChoose***(S_t)
  Performs weighted aggregation
  $w_{t+1} = \sum_{k=1}^{m} \alpha_{t,k} \cdot w_{t,k}$
end for

Function: ***ServerActionChoose***(S_t)
  if round $< W$ **then**
    Randomly sample an action $\tilde{A}_t$
  else
    ServerAgentTrain(P)
    Select action $\tilde{A}_t = \mu(S_t) + \mathcal{N}$
  end if
  Computes aggregation weights
  $A_t \leftarrow softmax(\tilde{A}_t)$
  return A_t

Function: ***ServerAgentTrain***(P)
  Sample a batch $B = (\hat{S}, \hat{A}, \hat{R}, \hat{S}') \sim P$
  Compute targets:
  $y(\hat{R}, \hat{S}') \leftarrow \hat{R} + \gamma Q'(\hat{S}', \mu'(\hat{S}'))$;
  Update Q with gradient descent:
  $\nabla_\theta \frac{1}{|B|} \sum_{(S,A,R,S')\in B} (Q(S,A) - y(R,S'))^2$;
  Update μ with gradient ascent:
  $\nabla_\phi \frac{1}{|B|} \sum_{S\in B} Q(S, \mu(S))$;
  Update target networks with:
  $\phi' \leftarrow \rho\phi' + (1-\rho)\phi,\ \theta \leftarrow \rho\theta' + (1-\rho)\theta$;
  return μ

2.3 Enhancing RL Agent with Adaptive Graph Convolutional Neural Network

Collaboration Graph Feature Extraction. To embed graph adjacency information into the optimization loop, we utilize AGCN to process the collaboration graph within the RL agent. As shown in Fig. 1 (b), two-layer AGCN computes each client's hidden representation by aggregating the feature vectors of its neighbors, allowing nodes/clients to receive information from both nearby and distant nodes. The AGCN layer can be represented as

$$H^{(l+1)} = \sigma(\tilde{A} H^{(l)} \Theta^{(l)} + b^{(l)}), \tag{4}$$

where $\tilde{A}$ is the Laplacian matrix for spectral convolution [12], and $\Theta^{(l)}$ and $b^{(l)}$ are the layer-specific weight and bias. $\sigma(\cdot)$ is the activation function $ReLU(\cdot)$. $H^{(l)} \in \mathbb{R}^{m\times d}$ represents the hidden features of the l^{th} layer (m: number of nodes, d: feature dimension). $H^{(0)}$ is the state vector input into the RL network.

Adaptive Adjacency Matrix. Leveraging a data-driven strategy, we dynamically uncover latent client associations, constructing a more comprehensive adjacency matrix that elucidates the collaborative relationships among clients, overcoming the limitations of approaches based solely on prior knowledge like data distribution distances. Specifically, the relationships between nodes are measured by the similarity of the learnable node embedding vector $E_G \in \mathbb{R}^{m \times d_e}$ (m: number of nodes, d_e: node embedding dimension), and then applying softmax normalization to the matrix. The Laplacian matrix, derived from the adjacency matrix A and degree matrix D, formulated as

$$\tilde{A} = I + D^{-\frac{1}{2}} A D^{-\frac{1}{2}} = softmax(E_G \cdot E_G^T) \tag{5}$$

The details of FedGraphRL's workflow are as described in Algorithm 1. To simplify notation, the transition vector $(S_{t-1}, A_{t-1}, R_{t-1}, S_t)$ is denoted as (S, A, R, S') when training the FedGraphRL agent. Details of the actor-critic network are given in the supplementary materials.

3 Experiments and Results

3.1 Dataset and Experimental Setup

Dataset. We conducted case studies on two real-world multi-center medical datasets: (1) HAM10k [26], encompassing skin lesion classification data from four distinct sources, and (2) Fed-DRG, comprising data collected from six public datasets (APTOS [10], DeepDR [18], FGADR [30], e-ophtha [5], IDRiD [22], Messidor [1]) for diabetic retinopathy grading. More details can be found in the supplementary materials. Each institution/source, acting as a single client, contributes data that reflects realistic heterogeneous distributions. The data is randomly split into training, validation, and test sets, with ratios of 0.7, 0.1, and 0.2 for each client, respectively. All images are resized to 128×128.
Implementation Details. In each client, we train Efficientnet-b0 [24] using Adam [11] optimizer with a learning rate of 5e-4, momentum of 0.9 and 0.99, and employing cross-entropy loss, for 200 and 50 rounds on the HAM10k and Fed-DRG to ensure steady model convergence. The local epoch is 1 and the batch size is 64. AGCN layer's embedding dimension d_e is 5. The agent is trained with a learning rate of 1e-3 and a batch size of 8, using a soft main-target update factor of 5e-3 and a discount factor γ of 0.99. Warm-up rounds W is designated as 16 and 8 for HAM10k and Fed-DRG, respectively.
Evaluation Metrics. According to [6], we use accuracy (Acc) as the performance metric of clients. Following the definition of fairness from [15,27], we utilize the standard deviation (Std.) of clients' test Acc to assess performance fairness and the average (Avg.) test Acc to evaluate the FL system's performance.

3.2 Comparison with SOTA Methods

We compared our method against the baseline approach FedAvg [20] and several current SOTA methods for FFL. FedFV [27] and FedMGDA+ [7] aim to find a more optimal global gradient by reducing local gradient conflicts. AFL [21] employs a min-max approach to protect the lowest-performing client.

Table 1. Performance of our FedGraphRL and nine competing methods. For each task, **best** and <u>second best</u> ranks are marked.

Task	Skin Lesion Classification						Diabetic Retinopathy Grading							
Client	1	2	3	4	Avg.	Std.	1	2	3	4	5	6	Avg.	Std.
FedAvg	96.33	74.26	57.08	67.05	73.68	14.43	97.17	79.75	85.87	70.45	68.93	70.49	78.78	10.2
AFL	95.06	69.64	66.15	85.22	79.02	11.72	96.76	73.25	91.85	82.95	71.84	70.77	80.86	10.15
q-FedAvg	92.92	70.83	66.81	84.09	78.66	10.42	95.14	73.75	92.12	72.73	76.7	72.21	80.44	9.47
TERM	96.21	70.39	54.20	76.14	74.23	15.02	97.17	80.00	86.68	73.86	70.87	69.91	79.75	9.68
FedFV	93.93	67.86	57.96	80.68	75.11	13.53	96.90	70.50	87.50	78.41	77.67	69.34	80.05	9.61
FedMGDA+	94.44	77.98	62.39	85.23	80.01	11.73	96.49	70.25	85.87	80.68	76.70	69.63	79.94	9.31
PropFair	95.95	72.17	58.63	76.14	75.72	13.36	97.30	77.75	86.96	67.05	77.67	71.06	79.63	10.04
Prop-FFL	95.32	76.49	66.37	81.82	80.00	10.44	96.22	76.00	92.12	73.91	75.73	71.92	80.98	9.49
FedCE	96.21	80.21	67.70	72.73	79.21	10.78	96.09	78.25	91.57	76.14	74.76	70.20	81.17	9.36
FedGraphRL	93.05	81.55	65.27	86.36	<u>81.56</u>	<u>10.25</u>	95.68	82.75	91.85	73.86	77.67	70.77	<u>82.10</u>	<u>9.09</u>
FedGraphRL*	94.31	81.55	68.14	86.36	**82.59**	**9.51**	95.41	80.50	92.39	82.95	77.67	73.07	**83.67**	**7.88**

Table 2. Performance of FedGraphRL, its four variants, and one baseline method.

Method		HAM10k		Fed-DRG	
		Avg.	Std.	Avg.	Std.
Baseline	FedAvg	73.68	14.43	78.78	10.20
Analysis of AGCN	RL	78.29	**9.39**	81.51	9.90
	RL + GCN w/ Pre-Defined Relation Graph	80.92	10.58	81.75	8.48
	RL +AGCN (FedGraphRL*)	**82.59**	9.51	**83.67**	**7.88**
Analysis of RL	AGCN	78.08	10.63	80.78	9.44
	AGCN + RL (FedGraphRL)	81.56	10.25	82.10	9.09
	AGCN + RL (FedGraphRL*)	**82.59**	**9.51**	**83.67**	**7.88**

q-Fedavg [15], TERM [13], Prop-Fair [28] and Prop-FFL [6] devise weighted aggregation schemes for model/gradient aggregation that flexibly balance accuracy and fairness. FedCE [8] uses the local client's Shapley Value estimates as weights for model aggregation. We treat one execution of Algorithm 1 as an episode. In Table 1, FedGraphRL is compared against other methods in a single episode to ensure comparison fairness under the same number of communication rounds. We also explore FedGraphRL* over multiple episodes, catering to scenarios of (1) high accuracy and fairness demands with sufficient communication resources, and (2) transferring the server-side agent to different FL settings. FedGraphRL* and FedGraphRL achieved top positions in average accuracy and the lowest standard deviations, demonstrating superior performance and fairness.

3.3 Further Analysis

Ablation Study on Different Parts of Our Method. The detailed ablation studies shown in Table 2 further validate the effectiveness of each component in FedGraphRL. In the analysis of the AGCN, "GCN w/ Pre-Defined Relation Graph" uses pairwise distances from [2] as its adjacency matrix. Lack of or singular client collaboration inputs, like distribution distance, impairs federated system optimization. Conversely, AGCN's adaptive adjacency matrix cap-

Table 3. The effect of q on the average and standard deviation of clients' test accuracy.

Task	Skin Lesion Classification						Diabetic Retinopathy Grading							
q	1	2	3	4	Avg.	Std.	1	2	3	4	5	6	Avg.	Std.
1	94.44	82.14	70.80	80.68	82.02	**8.40**	93.79	78.50	89.95	79.55	76.70	75.07	82.26	**7.03**
0.5	93.55	81.99	68.36	85.23	82.28	9.08	94.73	76.00	91.57	80.68	78.64	73.92	82.59	7.81
0.2	94.94	83.18	68.36	82.95	82.36	9.42	94.06	74.50	92.66	81.82	79.61	74.79	82.91	7.83
0.1	94.31	81.55	68.14	86.36	**82.59**	9.51	95.41	80.50	92.39	82.95	77.67	73.07	**83.67**	7.88
0	94.94	85.27	65.27	82.95	82.11	10.71	95.41	80.75	91.03	78.41	75.73	73.07	82.40	8.10

tures client relations more flexibly. For the RL component, since using only the AGCN component does not generate aggregation weights, we mimic [3] by using the weighted average of neighbor nodes as the aggregation weight. RL utilizes AGCN's collaboration graph for improved automated decision-making and fairness. As training episodes increase, replay buffer diversity improves, and FedGraphRL* achieves more precise decisions than FedGraphRL.

The Effect of the Fairness Impact Factor q. The aggregate loss in FedGraphRL's reward balances between performance and fairness using q. With

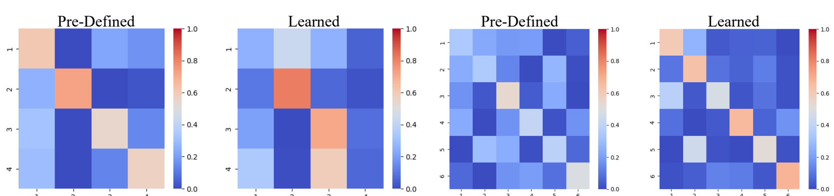

(a) Similarity of Collaboration Graph on HAM10k: 77.67% (b) Similarity of Collaboration Graph on Fed-DRG: 88.96%

Fig. 2. Visualization of pre-defined graph and learned graph on two federated datasets.

Visualization of Collaboration Graph G. Fig. 2 visualizes the normalized pre-defined graph used in the ablation study (left half) and the collaboration graph learned by AGCN (right half). On two federated datasets, the learned graph is more "comprehensive" than the pre-defined one, which solely depends on pairwise distribution distances. AGCN advances beyond learning from the pre-defined graph and reveals intricate hidden client relationships.

4 Conclusion

In this paper, we introduce a FedGraphRL-based model aggregation strategy for FFL. Equipped with a multi-layer AGCN, the FedGraphRL agent learns the collaboration graph among clients to reveal the collaborative relationships among clients in optimization, and captures the graph's features precisely. Through a learning-based mechanism automatically allocates aggregation weights to each client model, significantly improving fairness. Experiments on two multi-center medical image datasets demonstrate that FedGraphRL outperforms baseline models and the SOTA FFL method.

Acknowledgement. This work was supported in part by the National Natural Science Foundation of China under Grants 62171377, in part by Shenzhen Science and Technology Program under Grants JCYJ20220530161616036, in part by the Ningbo Clinical Research Center for Medical Imaging under Grant 2021L003 (Open Project 2022LYKFZD06), and in part by the Innovation Foundation for Master Dissertation of Northwestern Polytechnical University under Grant PF2024013.

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

References

1. Abràmoff, M.D., Folk, J.C., Han, D.P., Walker, J.D., Williams, D.F., Russell, S.R., Massin, P., Cochener, B., Gain, P., Tang, L., et al.: Automated analysis of retinal images for detection of referable diabetic retinopathy. JAMA ophthalmology **131**(3), 351–357 (2013)
2. Bao, W., Wang, H., Wu, J., He, J.: Optimizing the collaboration structure in cross-silo federated learning. In: ICML. vol. 202, pp. 1718–1736 (2023)
3. Chen, F., Long, G., Wu, Z., Zhou, T., Jiang, J.: Personalized federated learning with a graph. In: IJCAI. pp. 2575–2582 (2022)
4. Chen, J., Ma, B., Cui, H., Xia, Y.: Think twice before selection: Federated evidential active learning for medical image analysis with domain shifts. In: CVPR (2024)
5. Decenciere, E., Cazuguel, G., Zhang, X., Thibault, G., Klein, J.C., Meyer, F., Marcotegui, B., Quellec, G., Lamard, M., Danno, R., et al.: Teleophta: Machine learning and image processing methods for teleophthalmology. Irbm **34**(2), 196–203 (2013)
6. Hosseini, S.M., Sikaroudi, M., Babaie, M., Tizhoosh, H.: Proportionally fair hospital collaborations in federated learning of histopathology images. IEEE TMI (2023)
7. Hu, Z., Shaloudegi, K., Zhang, G., Yu, Y.: Federated learning meets multi-objective optimization. IEEE Transactions on Network Science and Engineering **9**(4), 2039–2051 (2022)
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)
9. Kairouz, P., McMahan, H.B., Avent, B., et al.: Advances and open problems in federated learning. Foundations and Trends in Machine Learning **14**(1–2), 1–210 (2021)

10. Karthick, M., Sohier, D.: Aptos 2019 blindness detection. Kaggle https://kaggle.com/competitions/aptos2019-blindness-detection Go to reference in chapter (2019)
11. Kinga, D., Adam, J.B., et al.: A method for stochastic optimization. In: ICLR. vol. 5, p. 6 (2015)
12. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017)
13. Li, T., Beirami, A., Sanjabi, M., Smith, V.: Tilted empirical risk minimization. In: ICLR (2021)
14. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated learning: challenges, methods, and future directions. IEEE Signal Processing Magazine **37**(3), 50–60 (2020)
15. Li, T., Sanjabi, M., Beirami, A., Smith, V.: Fair resource allocation in federated learning. In: ICLR (2020)
16. Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. In: ICLR (2016)
17. Liu, Q., Dou, Q., Yu, L., Heng, P.A.: Ms-net: multi-site network for improving prostate segmentation with heterogeneous mri data. IEEE TMI **39**(9), 2713–2724 (2020)
18. Liu, R., Wang, X., Wu, Q., Dai, L., Fang, X., Yan, T., Son, J., Tang, S., Li, J., Gao, Z., et al.: Deepdrid: Diabetic retinopathy—grading and image quality estimation challenge. Patterns **3**(6) (2022)
19. Ma, B., Zhang, J., Xia, Y., Tao, D.: Vnas: Variational neural architecture search. International Journal of Computer Vision pp. 1–25 (2024)
20. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: AISTATS (2017)
21. Mohri, M., Sivek, G., Suresh, A.T.: Agnostic federated learning. In: ICML (2019)
22. Porwal, P., Pachade, S., Kamble, R., Kokare, M., Deshmukh, G., Sahasrabuddhe, V., Meriaudeau, F.: Indian diabetic retinopathy image dataset (idrid): a database for diabetic retinopathy screening research. Data **3**(3), 25 (2018)
23. Roth, H.R., Chang, K., Singh, P., Neumark, N., Li, W., Gupta, V., Gupta, S., Qu, L., Ihsani, A., Bizzo, B.C., et al.: Federated learning for breast density classification: A real-world implementation. In: DART at MICCAI (2020)
24. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: ICML. pp. 6105–6114 (2019)
25. Tang, X., Yu, H.: Competitive-cooperative multi-agent reinforcement learning for auction-based federated learning. In: IJCAI (2023)
26. Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. Scientific data **5**(1), 1–9 (2018)
27. Wang, Z., Fan, X., Qi, J., Wen, C., Wang, C., Yu, R.: Federated learning with fair averaging. In: IJCAI. pp. 1615–1623 (2021)
28. Zhang, G., Malekmohammadi, S., Chen, X., Yu, Y.: Proportional fairness in federated learning. TMLR **2023** (2023)
29. Zhang, S.Q., Lin, J., Zhang, Q.: A multi-agent reinforcement learning approach for efficient client selection in federated learning. In: AAAI (2022)
30. Zhou, Y., Wang, B., Huang, L., Cui, S., Shao, L.: A benchmark for studying diabetic retinopathy: segmentation, grading, and transferability. IEEE TMI **40**(3), 818–828 (2020)