

Enable the Right to be Forgotten with Federated Client Unlearning in Medical Imaging

Zhipeng Deng¹, Luyang Luo¹, and Hao Chen^{1,2,3}✉

¹ Department of Computer Science and Engineering,

² Department of Chemical and Biological Engineering,

³ Division of Life Science,

The Hong Kong University of Science and Technology, Hong Kong, China

zdegaj@connect.ust.hk, cseluyang@ust.hk, jhc@cse.ust.hk

Abstract. The *right to be forgotten*, as stated in most data regulations, poses an underexplored challenge in federated learning (FL), leading to the development of federated unlearning (FU). However, current FU approaches often face trade-offs between efficiency, model performance, forgetting efficacy, and privacy preservation. In this paper, we delve into the paradigm of Federated Client Unlearning (FCU) to guarantee a client the right to erase the contribution or the influence, introducing the first FU framework in medical imaging. In the unlearning process of a client, the proposed model-contrastive unlearning marks a pioneering step towards feature-level unlearning, and frequency-guided memory preservation ensures smooth forgetting of local knowledge while maintaining the generalizability of the trained global model, thus avoiding performance compromises and guaranteeing rapid post-training. We evaluated our FCU framework on two public medical image datasets, including Intracranial hemorrhage diagnosis and skin lesion diagnosis, demonstrating that our framework outperformed other state-of-the-art FU frameworks, with an expected speed-up of 10-15 times compared with retraining from scratch. The code and the organized datasets can be found at: <https://github.com/dzp2095/FCU>.

Keywords: Federated Unlearning · Federated Learning

1 Introduction

To address the strict requirements on the collection, storage, and processing of personal data proposed in regulations like the General Data Protection Regulation (GDPR) [30] and the California Consumer Privacy Act (CCPA) [12], federated learning (FL) [23, 7, 21, 19] is regarded a promising privacy-preserving approach in medical imaging, which enables multiple parties to train a model collaboratively without sharing patient data. However, despite its decentralized nature, current FL research in medical imaging has not fully addressed the right to remove the influence of data from a trained global FL model, a right explicitly stated in GDPR as the *right to be forgotten* and in CCPA as the *right to*

delete. In centralized learning, the right to have data removed can be realized by Machine Unlearning (MU) [32, 2]. Regardless, the existing MU techniques are developed for the centralized scenarios, posing significant challenges to their direct application in distributed FL settings [11], highlighting the need for dedicated federated unlearning (FU) approaches.

Retraining from scratch without the target forgotten data is regarded as a naive way to achieve unlearning [1]. Nevertheless, this approach demands a large cost in communication and computation, especially in FL [25]. Hence, recent studies have proposed various federated unlearning (FU) methods including re-calibration of historical updates [20], gradient quantization [3], gradient modification [11] or knowledge distillation [31, 34]. However, these methods often compromise performance or privacy. For instance, FedEraser[20] accelerates retraining progress and removes the contribution of a target client iteratively by utilizing historical parameter updates of clients stored on the server side, yet this approach requires additional storage and poses a risk of data reconstruction by a malicious server [28]. FFMU [3] applies randomized gradient smoothing and quantization to execute unlearning operations on the target forgotten data, but may fail to retain the performance when a client decides to remove all their data. UPGA [11] formulates the unlearning process as a constrained maximization problem by limiting the unbounded loss to an ℓ_2 -norm sphere by a designated reference model that may be difficult to obtain. FUKD [31] removes the contribution of a target client by subtracting the historical parameter updates and recovering the model performance through knowledge distillation [14], necessitating unlabeled data on the server. Additionally, MoDe [34] adjusts pre-trained model parameters through two phases—knowledge erasure and memory guidance—to reduce discriminability for target forgotten data and restore performance, dependent on the initial state of the degraded model. These limitations underscore the necessity for a more effective FU framework that addresses these challenges without compromising efficiency or privacy.

To make better use of the information contained in the teacher network, since [26], most knowledge distillation shifted from output distillation to feature distillation [13, 33], showing superior performance on various tasks. However, existing knowledge distillation-based unlearning methods focus on merely output distillation [34, 5, 17]. MoDe [34] achieves unlearning by using a degraded model (i.e., a “Bad Teacher”) that has never been trained on the forgotten data, to generate pseudo labels for the student model on forgotten data. Chundawat et al. [5] employ a teacher-student objective that minimizes KL-Divergence between the output of the “Bad Teacher” and the student, encouraging the student model to align closely with the “Bad Teacher” on the forgotten set. Similarly, SCRUB [17] suggests maximizing KL-Divergence to encourage the student model to “move away” from the trained teacher model on forgotten data. To the best of our knowledge, no method has yet considered encouraging the student model to learn from the “Bad Teacher” at the feature level to guarantee a higher level forgetting, marking an opportunity for innovation in unlearning.




Fig. 1. Overview of Federated Client Unlearning (FCU): The target client conducts unlearning locally and updates the server with the unlearned model, which then serves as the initial model for post-training on the remaining clients.

In this paper, we delve into the paradigm of Federated Client Unlearning and present the first FU framework in medical imaging to ensure the right of a target client to remove the contribution of their data from a trained global model efficiently. We use Model-Contrastive Unlearning (MCU) to encourage the model to perform similarly with a downgraded model and differently with the trained global model on the forgotten data, which pioneers a step towards achieving unlearning at the feature level. Meanwhile, to preserve the generalized knowledge and only remove client-specific knowledge of the target client, we use Frequency-Guided Memory Preservation (FGMP) to preserve the low-frequency components of the trained model, ensuring a rapid post-training on the remaining clients. We validate our proposed method on two real-world tasks, including intracranial hemorrhage (ICH) diagnosis and skin lesion diagnosis. Extensive experiments demonstrate that our method outperforms a number of state-of-the-art FU methods, without compromising privacy, and with an expected speed-up of 10-15 times compared with retraining from scratch.

2 Methodology

2.1 Preliminaries and Overview

We adopt a typical Federated Unlearning (FU) scenario as described in prior works [11, 20], involving K clients $\mathbf{C} = \{C_1, C_2, \dots, C_K\}$ and a central server participating in FL. The dataset held by each client is denoted as $\mathbf{D} = \{D_1, D_2, \dots, D_K\}$.

Suppose that after t rounds of FL, each client possesses a trained global model $\mathbf{M}_{tr}$. We refer to the client C_u wants to opt out as the *target client*, where C_u requests to remove the contribution of their data D_u from $\mathbf{M}_{tr}$. Following [20, 3, 11], the goal of this study is to unlearn the D_u , effectively eliminating its influence from $\mathbf{M}_{tr}$ to produce an unlearned model $\mathbf{M}_{un}$.

Our proposed framework FCU is presented in Fig. 1. The target client C_u initiates the unlearning process by performing local unlearning to generate the initial unlearned model $\mathbf{M}_{un}$, where the proposed Model-Contrastive Unlearning (MCU) is to make $\mathbf{M}_{un}$ perform similarly with a model that has never trained on D_u and differently with the trained model $\mathbf{M}_{tr}$ at the feature level. Furthermore, to preserve the generalized knowledge and only remove client-specific knowledge of the target client, we use Frequency-Guided Memory Preservation (FGMP) to preserve the low-frequency components of the trained model, thereby achieving a rapid post-training to generate the final unlearned model $\mathbf{M}_{un}$.

2.2 Model-Contrastive Unlearning

The intuition of our Model-Contrastive Unlearning (MCU) is to achieve unlearning by encouraging the unlearned model $\mathbf{M}_{un}$ perform similarly at feature level to a model that has never been trained on D_u , which we refer to as a downgraded model $\mathbf{M}_{down}$. Model-Contrastive Learning was first proposed in [18], where they aim to decrease the representation drift between the local model and the global model in FL. In contrast, we propose to encourage the unlearned model to output similar features as a downgraded model $\mathbf{M}_{down}$ (pull) and dissimilar features as the trained global model $\mathbf{M}_{tr}$ (push), where we refer to this process as Model-Contrastive Unlearning (MCU).

We use a model with the same structure as the $\mathbf{M}_{tr}$ but only pretrained on ImageNet[6] while not being trained on D_u as the downgraded model $\mathbf{M}_{down}$, based on two intuitions: 1) $\mathbf{M}_{down}$ also serves as pretrained model for $\mathbf{M}_{tr}$ in FL training before FU, 2) $\mathbf{M}_{down}$ possess the ability to extract low-level features of images [24]. In contrast, MoDe[34] chose a randomly initialized model as a degraded model to generate pseudo labels for the unlearned model, where this randomly initialized model may be strongly dependent on the initial weights. Similar to [18], our Model-Contrastive Unlearning loss can be formulated as:

$$\mathcal{L}_{mcu} = -\log \left(\frac{\exp(\text{sim}(z, z_{down})/\tau)}{\exp(\text{sim}(z, z_{down})/\tau) + \exp(\text{sim}(z, z_{tr})/\tau)} \right), \quad (1)$$

where $\text{sim}(\cdot, \cdot)$ is a cosine similarity function, τ , known as the temperature, z represents the feature vector extracted by the model for a given input x , z_{down} is the feature vector extracted by the downgraded model $\mathbf{M}_{down}$, and z_{tr} is the feature vector extracted by the trained global model $\mathbf{M}_{tr}$.

2.3 Frequency-Guided Memory Preservation

The goal of the unlearning process is to remove the specific knowledge of the target client C_u without losing the generalized knowledge already learned in

a trained global model $\mathbf{M}_{tr}$. Inspired by the findings of [22, 4], which indicates that the low-frequency components of parameters may reflect the basis for global features across all clients while high-frequency components may contain specific knowledge for an individual client, we introduce Frequency-Guided Memory Preservation (FGMP). FGMP aims to preserve the low-frequency components of $\mathbf{M}_{tr}$ and high-frequency components of the unlearned model $\mathbf{M}_{un}$ to acquire a newly unlearned model after MCU. Intuitively, the newly unlearned model $\mathbf{M}'_{un}$ maintains the generalized knowledge inherited from $\mathbf{M}_{tr}$ and has the specific knowledge from D_u removed in high-frequency components by MCU.

Specifically, we conduct FGMP every T_{FGMP} iterations (e.g. T_{FGMP} is set to 10 in our experiment) while MCU is continuously executed, resulting in an unlearned model $\mathbf{M}_{un}$. We use FFT to convert the parameters of $\mathbf{M}_{un}$ and $\mathbf{M}_{tr}$ into frequency domain, and preserve the low-frequency part of $\mathbf{M}_{tr}$ and high-frequency part of $\mathbf{M}_{un}$. Afterwards, the newly unlearned model $\mathbf{M}'_{un}$ is constructed by inverse FFT (IFFT). We conduct FFT and IFFT of model parameters similarly as [4]. To clarify, for the weights w in a convolutional layer that has N input channels, H output channels and a kernel with size $d_1 \times d_2$, we reshape $w \in \mathbb{R}^{N \times H \times d_1 \times d_2}$ into a 2-D matrix $w' \in \mathbb{R}^{d_1 N \times d_2 H}$ to ease the process of FFT and IFFT. Then, we can obtain the amplitude map $\mathcal{F}^A$ and phase map $\mathcal{F}^P$ through the Fourier transform $\mathcal{F} = \mathcal{F}^A e^{j\mathcal{F}^P}$:

$$\mathcal{F}(w)(m, n) = \sum_{x, y} w'(x, y) e^{-j2\pi \left(\frac{x}{d_1 N} m + \frac{y}{d_2 H} n \right)}, \quad j^2 = -1. \quad (2)$$

To extract the low-frequency components, we define a mask matrix M of the same dimensions as w' , $M \in \{0, 1\}^{d_1 N \times d_2 H}$. $M_{ij} = 1$ for the central region, and $M_{ij} = 0$ elsewhere. The central region is a rectangle centered around the middle of M , with dimensions $\lfloor rd_1 N \rfloor \times \lfloor rd_2 H \rfloor$, where r is the ratio of the low-frequency part to be preserved, and $\lfloor \cdot \rfloor$ denotes the floor function to ensure integer dimensions. Hence, the newly unlearned model in the frequency domain that preserves the low frequency part of $\mathbf{M}_{tr}$ and high frequency part of $\mathbf{M}_{un}$ can be formulated as: $\hat{\mathcal{F}}^A(w'_{adj}) = M \odot \mathcal{F}^A(w_{pre}) + (1 - M) \odot \mathcal{F}^A(w_{adj})$, where $\odot$ means element-wise multiplication. Finally, we apply IFFT $\mathcal{F}^{-1}$ to convert the amplitude and phase maps back to the parameter as $w'_{adj} = \mathcal{F}^{-1}(\hat{\mathcal{F}}^A(w'), \mathcal{F}^P(w'))$.

After FGMP, we obtain a newly unlearned model $\mathbf{M}'_{un}$. Intuitively, the high-frequency components of $\mathbf{M}_{un}$ have been made to forget the specific knowledge of the target client due to the application of MCU, while the low-frequency components retain the generalized knowledge. This selective retention and forgetting forms a solid foundation for post-training, ensuring that the model preserves its ability to generalize well while removing client-specific knowledge.

2.4 Overall of FCU framework

After local unlearning, the server sends the unlearned model $\mathbf{M}_{un}$ to the remaining clients request them to conduct post-training using FedAvg [23], where

$\mathbf{M}_{un}$ serves as the initial global model and the global objective can be formulated as: $\min_W \mathbf{L}(W) = \sum_{k \in \{1, \dots, K\} \setminus \{u\}} \mathcal{L}_k(W)$, where $\mathcal{L}_k$ is the local objective of client C_k , and W is the parameters of the global model. The global model parameter W is iteratively updated with the aggregation of local models on the remaining clients, which is defined as $W^{t+1} = \sum_{k \in \{1, \dots, K\} \setminus \{u\}} \frac{n_k}{n - n_u} W_k^t$, where: W^{t+1} represents the updated global model parameters, n_k and n_u denote the sample sizes of the k -th client and client u , respectively, n is the total sample size from all clients, and W_k^t are the parameters from the k -th client's model. Due to the memory preservation achieved by FGMP, our post-training can efficiently restore model performance on remaining datasets with a few rounds.

3 Experiments

3.1 Experiment Setup

Datasets. We evaluated our method on two public real-world medical image classification tasks: 1) Intracranial hemorrhage (ICH) diagnosis. We use the RSNA-ICH dataset [8] and follow [16] to perform the binary diseased-or-healthy classification, and randomly sample 25,000 slices. 2) ISIC2018 skin lesion diagnosis. We conducted skin lesion diagnosis with HAM10000 [29], which contains 10,015 dermoscopy images. Training, validation and testing sets for both datasets were divided into 7:1:2. For both tasks, to simulate heterogeneous multi-source data, following [27], Dirichlet distribution, i.e. $Dir(\alpha = 1.0)$, is used to divide the training set to 5 clients.

Implementation Details. For both tasks, We used DenseNet121[15] as the backbone. The network was optimized by Adam optimizer where the momentum terms were set to 0.9 and 0.99, with learning rate set to $1e^{-5}$ at target client and $1e^{-4}$ for other clients. The total batch size was 64 in both local training and local unlearning, the local unlearning iterations were 100. The temperature τ in MCU loss was 0.5 by default like [18]. The interval T_{FGMP} to execute FGMP in MCU was set to 10. During post-training, the local training iterations were 20, and the total communication rounds were 10. The images in both tasks were resized to 224×224 . For Task 1, data augmentation included a combination of random flip, rotation, translation, scaling, and gaussian blur. For Task 2, we employed the random flip, rotation, and translation.

Evaluation Metrics. We adopt four widely recognized metrics as in recent representative FU study FFMU [3] and other commonly recognized metrics to assess machine unlearning performance across three dimensions: **1) Fidelity** that assesses whether the unlearning methods preserve the original model's performance. This includes measuring F1-score, Accuracy, errors on the retained data $Error^r$ (evaluated on D_r , the data held by the remaining clients $\mathbf{C} \setminus C_u$), and errors on the test dataset $Error^t$. **2) Efficacy** that evaluates the success of an FU method in eradicating the influence of the target client's data, D_u . This is gauged by $Error^f$, the classification errors on the forgotten dataset D_f (i.e. the dataset on the target client C_u). A model's $Error^f$ close to that of a retrained

Table 1. Comparison with state-of-the-art methods on ICH diagnosis. * denotes the retrained model is regarded as the gold-standard for $Error^f$.

Methods	Fidelity				Efficacy	Efficiency
	Accuracy	F1	$Error^t$	$Error^r$	$Error^f$	Runtime (s)
Origin	86.84	86.84	13.15	8.44	6.59(−5.7)	3612
Retrain	85.52	85.40	14.48	10.60	12.28*(0.0)	2873
Finetune	86.11	86.01	13.88	7.86	7.25(−5.0)	301
FFMU [3]	85.09	84.98	14.90	10.28	9.79(−2.5)	249
MoDe [34]	80.88	80.62	19.12	14.59	18.34(+6.1)	509
UPGA [11]	78.96	78.93	21.04	20.43	20.25(+8.0)	1031
FedEraser [20]	84.32	84.30	15.67	11.72	15.29(+3.0)	906
FUKD [31]	83.66	83.54	16.34	12.49	15.63(+3.4)	1347
w/o FGMP	85.42	85.36	14.58	10.21	11.80(−0.5)	1831
w/o post-training	82.57	82.53	17.42	12.06	16.80(+4.5)	14
Ours	86.40	86.32	13.60	8.11	11.37(−0.9)	177

model (which has never encountered D_f) is considered favorable [3, 10, 9]. **3) Efficiency** that measures the reduction in communication and computational overheads by quantifying runtime, with each method trained to convergence for a fair comparison. All the results were averaged over 3 runs.


3.2 Comparison with state-of-the-arts

We compared our method with the trained global model denoted as **origin**, the model **finetuned** on the trained global model as a baseline, and the model **retrained** from scratch as gold-standard for efficacy [3, 10, 9]. Besides, we compared with five recent state-of-the-art (SOTA) methods, including **FFMU** [3] applying random gradient quantization, **MoDe** [34] utilizing a degraded model to unlearn, **UPGA** [11] employing projected gradient ascent to maximize the empirical loss on the target client, **FedEraser** [20] eliminating the influence of the target client by historical parameter updates iteratively, and **FUKD** [31] that erases the contributions of clients by subtracting the historical parameter updates and restore the performance by knowledge distillation.

The quantitative results for two tasks are presented in Table 1 and Table 2. Our method leads in fidelity, achieving improvements of approximately 1.5% in accuracy and 1.6% in F1 score for Task 1, and 1.9% in accuracy and 5.0% in F1 score for Task 2, compared to the second-best FU method, showing robustness in model performance maintenance. In efficacy, it nearly matches the retrained model considered the standard in forgetting [3, 10, 9], ensuring effective unlearning. Besides, we proved that the finetune method failed to achieve unlearning. For Efficiency, we see an impressive runtime reduction, achieving roughly a 10 to 15 times speed-up compared with retraining from scratch.

Table 2. Comparison with state-of-the-art methods on Skin Lesion diagnosis. * denotes the retrained model is regarded as the gold-standard for $Error^f$.

Methods	Fidelity				Efficacy	Efficiency
	Accuracy	F1	$Error^t$	$Error^r$	$Error^f$	Runtime (s)
Origin	79.98	57.90	20.01	9.66	29.95(−5.4)	2469
Retrain	81.52	54.76	18.47	5.40	35.37*(0.0)	2038
Finetune	80.52	55.43	19.47	8.43	30.49(−4.9)	289
FFMU [3]	79.67	44.64	20.33	10.32	30.99(−4.4)	176
MoDe [34]	73.04	33.63	26.96	16.55	50.84(+15.5)	431
UPGA [11]	75.08	38.43	24.91	17.51	48.28(+12.9)	899
FedEraser [20]	80.13	53.64	19.87	8.72	37.49(+2.1)	672
FUKD [31]	78.11	41.27	21.89	12.69	39.72(+4.4)	1132
w/o FGMP	80.91	53.92	19.09	6.69	33.69(−1.6)	1579
w/o post-training	76.83	40.79	23.16	13.85	42.58(+7.2)	13
Ours	81.67	56.32	18.33	8.22	34.45(−0.9)	156

**Fig. 2.** Ablation study of FGMP and MCU across local unlearning iterations, showing forgotten $Error^f$ and Accuracy for the two tasks.

Ablation study. We conducted ablation studies to assess the effectiveness of the primary components of our FCU framework. As shown in Table 1 and Table 2, the performance drops significantly without post-training, and the runtime increases a lot without FGMP. The effectiveness of our FGMP in facilitating local unlearning across various iterations is demonstrated in Fig. 2, which can maintain the performance at its maximum level. Without FGMP, there is a noticeable decline in test accuracy even before the model has fully executed unlearning across the entire dataset. Besides, we replace our MCU with KL-based knowledge distillation [5] and pseudo label-based knowledge distillation [34], which shows that our MCU can better preserve the performance on test set when achieve at similar forgotten error.

4 Conclusion

We present the first federated unlearning framework in medical imaging, which facilitates the right for a client to be forgotten. In the local unlearning phase, our FCU utilizes Model-Contrastive unlearning (MCU) to encourage the model to perform similarly to a model that has never seen the forgotten data at the feature level. To preserve the generalized knowledge, we use Frequency-Guided Memory Preservation (FGMP) to preserve the low-frequency components of the trained global model, ensuring a smooth forgetting process. Benefited from FGMP, our FCU framework quickly restores performance with minimal post-train rounds, achieving a 10-15 times speed-up over retraining from scratch, while demonstrating remarkable federated unlearning effectiveness for medical imaging.

Acknowledgments. This work was supported by the Hong Kong Innovation and Technology Fund (Project No. MHP/002/22) and and HKUST 30 for 30 Research Initiative Scheme (No. 3030_024).

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

References

1. Bourtole, L., Chandrasekaran, V., Choquette-Choo, C.A., Jia, H., Travers, A., Zhang, B., Lie, D., Papernot, N.: Machine unlearning. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 141–159. IEEE (2021)
2. Cao, Y., Yang, J.: Towards making systems forget with machine unlearning. In: 2015 IEEE symposium on security and privacy. pp. 463–480. IEEE (2015)
3. Che, T., Zhou, Y., Zhang, Z., Lyu, L., Liu, J., Yan, D., Dou, D., Huan, J.: Fast federated machine unlearning with nonlinear functional theory. In: Fortieth International Conference on Machine Learning (2023)
4. Chen, Z., Zhu, M., Yang, C., Yuan, Y.: Personalized retrogress-resilient framework for real-world medical federated learning. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24. pp. 347–356. Springer (2021)
5. Chundawat, V.S., Tarun, A.K., Mandal, M., Kankanhalli, M.: Can bad teaching induce forgetting? unlearning in deep networks using an incompetent teacher. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 7210–7217 (2023)
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
7. Deng, Z., Luo, L., Chen, H.: Scale federated learning for label set mismatch in medical image classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 118–127. Springer (2023)
8. 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)

9. Fu, S., He, F., Tao, D.: Knowledge removal in sampling-based bayesian inference. In: International Conference on Learning Representations (2021)
10. Golatkar, A., Achille, A., Soatto, S.: Eternal sunshine of the spotless net: Selective forgetting in deep networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9304–9312 (2020)
11. Halimi, A., Kadhe, S., Rawat, A., Baracaldo, N.: Federated unlearning: How to efficiently erase a client in fl? arXiv preprint arXiv:2207.05521 (2022)
12. Harding, E.L., Vanto, J.J., Clark, R., Hannah Ji, L., Ainsworth, S.C.: Understanding the scope and impact of the california consumer privacy act of 2018. *Journal of Data Protection & Privacy* **2**(3), 234–253 (2019)
13. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1921–1930 (2019)
14. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
15. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
16. Jiang, M., Zhong, Y., Le, A., Li, X., Dou, Q.: Client-level differential privacy via adaptive intermediary in federated medical imaging. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 500–510. Springer (2023)
17. Kurmanji, M., Triantafillou, P., Hayes, J., Triantafillou, E.: Towards unbounded machine unlearning. *Advances in Neural Information Processing Systems* **36** (2024)
18. Li, Q., He, B., Song, D.: Model-contrastive federated learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10713–10722 (2021)
19. Li, X., Jiang, M., Zhang, X., Kamp, M., Dou, Q.: Fedbn: Federated learning on non-iid features via local batch normalization. arXiv preprint arXiv:2102.07623 (2021)
20. Liu, G., Ma, X., Yang, Y., Wang, C., Liu, J.: Federaser: Enabling efficient client-level data removal from federated learning models. In: 2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS). pp. 1–10. IEEE (2021)
21. Liu, Q., Yang, H., Dou, Q., Heng, P.A.: Federated semi-supervised medical image classification via inter-client relation matching. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24. pp. 325–335. Springer (2021)
22. Liu, Z., Xu, J., Peng, X., Xiong, R.: Frequency-domain dynamic pruning for convolutional neural networks. *Advances in neural information processing systems* **31** (2018)
23. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Artificial intelligence and statistics. pp. 1273–1282. PMLR (2017)
24. Neyshabur, B., Sedghi, H., Zhang, C.: What is being transferred in transfer learning? *Advances in neural information processing systems* **33**, 512–523 (2020)
25. Romandini, N., Mora, A., Mazzocca, C., Montanari, R., Bellavista, P.: Federated unlearning: A survey on methods, design guidelines, and evaluation metrics. arXiv preprint arXiv:2401.05146 (2024)
26. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. arXiv preprint arXiv:1412.6550 (2014)

27. Shang, X., Lu, Y., Huang, G., Wang, H.: Federated learning on heterogeneous and long-tailed data via classifier re-training with federated features. *arXiv preprint arXiv:2204.13399* (2022)
28. Singhal, K., Sidahmed, H., Garrett, Z., Wu, S., Rush, J., Prakash, S.: Federated reconstruction: Partially local federated learning. *Advances in Neural Information Processing Systems* **34**, 11220–11232 (2021)
29. 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)
30. Voigt, P., Von dem Bussche, A.: The eu general data protection regulation (gdpr). A Practical Guide, 1st Ed., Cham: Springer International Publishing **10**(3152676), 10–5555 (2017)
31. Wu, C., Zhu, S., Mitra, P.: Federated unlearning with knowledge distillation. *arXiv preprint arXiv:2201.09441* (2022)
32. Xu, H., Zhu, T., Zhang, L., Zhou, W., Yu, P.S.: Machine unlearning: A survey. *ACM Computing Surveys* **56**(1), 1–36 (2023)
33. Zhao, B., Cui, Q., Song, R., Qiu, Y., Liang, J.: Decoupled knowledge distillation. In: *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*. pp. 11953–11962 (2022)
34. Zhao, Y., Wang, P., Qi, H., Huang, J., Wei, Z., Zhang, Q.: Federated unlearning with momentum degradation. *IEEE Internet of Things Journal* (2023)