



Inverse Distance Aggregation for Federated Learning with Non-IID Data

Yousef Yeganeh¹(✉), Azade Farshad¹, Nassir Navab^{1,2}, and Shadi Albarqouni^{1,3}

¹ Computer Aided Medical Procedures, Technical University of Munich,
Munich, Germany
y.yeganeh@tum.de

² Whiting School of Engineering, Johns Hopkins University, Baltimore, USA

³ Department of Computing, Imperial College London, London, UK

Abstract. Federated learning (*FL*) has been a promising approach in the field of medical imaging in recent years. A critical problem in *FL*, specifically in medical scenarios is to have a more accurate shared model which is robust to noisy and out-of distribution clients. In this work, we tackle the problem of statistical heterogeneity in data for *FL* which is highly plausible in medical data where for example the data comes from different sites with different scanner settings. We propose **IDA** (Inverse Distance Aggregation), a novel adaptive weighting approach for clients based on meta-information which handles unbalanced and non-iid data. We extensively analyze and evaluate our method against the well-known *FL* approach, Federated Averaging as a baseline.

Keywords: Deep learning · Federated learning · Distributed learning · Privacy-preserving. · Heterogeneous data · Robustness

1 Introduction

Federated learning (*FL*) was proposed as a decentralized learning scheme where the data in each client is private and not exposed to other participants, yet they contribute to generation of a shared (global) model in a server that represents the clients' data [12]. An aggregation strategy in the server is essential in *FL* for combining the models of all clients. Federated Averaging (*FedAvg*) [21] is one of the most well-known *FL* methods which uses the normalized number of samples in each client to aggregate the models in the server. Another aggregation approach using temporal weighting along with a synchronous learning strategy was proposed in [3]. Many recent approaches have been proposed in order to improve the generalization or personalization of the global model using the ideas of knowledge transfer, knowledge distillation, multi-task learning and meta-learning [1, 2, 4, 8, 9, 15, 29].

Project page: <https://ida-fl.github.io/>.

© Springer Nature Switzerland AG 2020
S. Albarqouni et al. (Eds.): DART 2020/DCL 2020, LNCS 12444, pp. 150–159, 2020.
https://doi.org/10.1007/978-3-030-60548-3_15

Even though *FL* has emerged into a promising and popular method to engage with privacy preserving distributed learning, it has faced some challenges: **a)** Expensive communication, **b)** privacy, **c)** systems heterogeneity and **d)** statistical heterogeneity [16]. Although a large number of recent works on *FL* such as [20, 25] are focused on communication efficiency due to its application on edge devices with unstable connections [16], commonly using approaches such as compressed networks or compact features, its most determining aspects in the medical field are data privacy and heterogeneity [11, 23]. Data heterogeneity assumption includes: **a)** Massively distributed: The data points are distributed among a very large number of clients. **b)** Non-iid (Not independent and identically distributed): Data in each node comes from a distinct distribution. The local data points are not representative of the whole data distribution (combination of all clients' data). **c)** Unbalancedness: The number of samples across clients has a high variance. Such heterogeneity is foreseeable in medical data due to many reasons, for example, class imbalance in pathology, intra-/inter-scanner variability (domain shift), intra-/inter-observer variability (noisy annotations), multi-modal data, and different tasks for clients.

There has been numerous works to handle each of these data assumptions [10]. Training a global model with *FL* in non-iid data is a challenging task. Model training in deep neural network suffers quality loss and may even diverge given non-iid data [5]. There has been multiple works dealing with this problem. Sattler et al. [24] propose clustering loss terms and using cosine similarity to overcome the divergence problem when clients have different data distributions. Zhao et al. [33] overcome the non-iid problem by creating a subset of data which is shared globally with the clients. In order to maintain system heterogeneity (affected by their main idea of nonuniform local updates), FedProx [17] proposes a proximal term to minimize the distance between the local and global models. Close to our approach, geometric median is used in [22] to decrease the effect of corrupted gradients on the federated model.

In the last few years, there has been a growing interest in applying *FL* in healthcare, in particular, to medical imaging. Sheller et al. [27] were among the first works who applied *FL* to multi-institutional data for Brain Tumor Segmentation task. To date, there has been numerous works on *FL* in Healthcare [7, 18, 19, 26, 32]. However, little attention has been paid to the aggregation mechanism given the data and system heterogeneity; for example, when the data is non-iid, or the participation rate of the clients is pretty low.

In this work, we try to overcome the challenges of statistical heterogeneity in data and propose a robust aggregation method at the server side (*cf.* Fig. 1). Our weighting coefficients are based on the meta-information extracted from the statistical properties of the model parameters. Our goal is to train a low variance global model given high variance local models which is robust to non-iid and unbalanced data. Our contributions are twofolds; **a)** A novel adaptive weighting scheme for federated learning which is compatible with other aggregation approaches, **b)** Extensive evaluation of different scenarios on non-iid data on multiple datasets.

Next, a brief overview of the federated learning concept is introduced in the methodology section before diving into the main contribution of the paper, the Inverse Distance Aggregation (IDA). Experiments and results on both machine learning datasets (Proof-of-Concept), and clinical use-cases are demonstrated and discussed.




Fig. 1. Federated learning with non-iid data - the data has different distributions among clients.

2 Method

Given a set of K clients with their own data distribution $p_k(x)$ and a shared neural network with parameters ω , the objective is to train a global model minimizing the following objective function;

$$\arg \min_{\omega_g^t} f(x; \omega_g^t), \quad \text{where} \quad f(x; \omega_g^t) = \sum_{k=1}^K f(x; \omega_k^t), \quad (1)$$

where ω_g^t, ω_k^t are the global and local parameters, respectively.

2.1 Client

Each randomly sampled client, from the total number of K clients (based on the participation rate pr), receives the global model parameter ω_g^t at communication round t , and trains the shared model, initialized by ω_g^t , on its own training data $p_k(x)$ for E iterations to minimize its local objective function $f_k(x) = \mathbf{E}_{x \sim p_k(x)} [f(x; \omega_k^t)]$ where ω_k^t is the weight parameters of the client k . The training data in each client is a subset of the whole training data, which can be sampled from different classes of data. The number of classes of data assigned to each client is denoted by n_{cc} .

2.2 Server

Each round t , the updated local parameters ω_k^t are sent back to the server and aggregated to form the updated global parameter ω_g^t ,

$$\omega_g^t = \sum_{k=1}^K \alpha_k \cdot \omega_k^{t-1}. \quad (2)$$

where α_k is the weighting coefficient. This procedure continues for the given total communication rounds T .

2.3 Inverse Distance Aggregation (IDA)

In order to reduce the inconsistency among the updated local parameters due to the non-iid problem, we propose a novel robust aggregation method, denoted as Inverse Distance Aggregation (**IDA**). The core of our method is the way the coefficients α_k are computed, which is based on the inverse distance of each client parameters to the average model of all clients. This allows us to reject or weigh less the models who are poisoning, *i.e.* out-of-distribution models.

To realize this, the ℓ_1 -norm is utilized as a metric to measure the distance of clients ω_k to the average one ω_{Avg} as

$$\alpha_k = \frac{1}{Z} \|\omega_{Avg}^{t-1} - \omega_k^{t-1}\|^{-1}, \quad (3)$$

where $Z = \sum_{k \in K} \|\omega_{Avg}^{t-1} - \omega_k^{t-1}\|^{-1}$ is a normalization factor. In practise, we add ϵ to both numerator and denominator to avoid any numerical instability. Note that $\alpha_k = 1$ when clients' parameters is equivalent to the average one, and $\alpha_k = n_k$ is equivalent to the *FedAvg* [21].

We also propose to use the training accuracy of clients in the final weighting which we denote by INTRAC (INverse TRaining ACcuracy) to penalize over-fitted models and encourage under-trained models in the aggregated model. To calculate the coefficients for INTRAC, We assign $\alpha'_k = \frac{Z'}{\max(\frac{1}{K}, acc_k)}$. The *max* function is used to assure all of the values are above chance level. Here acc_k is the training accuracy of client k , α'_k is the INTRAC coefficient and $Z' = \sum_{k \in K} \max(\frac{1}{K}, acc_k)$ is the normalization factor. We normalize the calculated coefficients α'_k once again to bring them to the range of $(0, 1]$. To combine different coefficient values (*i.e.* INTRAC, IDA, FedAvg), we multiply the acquired coefficients and normalize them in the range of $(0, 1]$.

3 Experiments and Results

We evaluated our method on commonly used databases to show a Proof-of-Concept (PoC) before we present some results on a clinical use-case. We compare the results of our method **IDA** against the baseline method *FedAvg* [21]. In the first set of PoC experiments, we investigate the following: 1) Non-iid vs. iid: Comparison of *FedAvg* and **IDA** in iid and non-iid with different datasets and architectures. 2) Ablation study: Investigation of effectiveness of IDA compared to FedAvg 3) Sensitivity analysis: Performance comparison in extreme situations.

Datasets We show the results of our evaluation on cifar-10 [13], fashion-mnist (f-mnist) [31] and HAM10K(multi-source dermatoscopic images of pigmented lesions)[30] datasets. f-mnist is a well-known variation of mnist with 50k images of 28×28 black and white clothing pieces. cifar-10 is another dataset with 60k 32×32 images of vehicles and animals, commonly used in computer vision. For the clinical study, we evaluate our method on HAM10k dataset which includes a total number of 10015 images of different pigmented skin lesions in 7 classes. The different classes and their number of samples in HAM10k are as follows: Melanocytic nevi: 6705, Melanoma: 1113, Benign keratosis: 1099, Basal cell carcinoma: 514, Actinic Keratoses: 327, Vascular: 142, Dermatofibroma: 115. We chose this dataset due to its heavy unbalancedness.

Implementation Details. The training settings for each dataset are: LeNet [14] for f-mnist with 10 classes, batchsize=128, learning rate (lr) = 0.05 and local iteration of 1 ($E = 1$), VGG11 [28] without batch normalization and dropout layers for cifar-10 with 10 classes and batchsize = 128, lr = 0.05 and $E = 1$. For HAM10K, we used Densenet-121 [6] with 7 classes, batchsize = 32, lr = 0.016 and $E = 1$. In all of the experiments 90% of the images are randomly sampled for training and the rest are employed for evaluation. All of the models are trained for a total number of 5000 rounds. The mentioned values are the default for all experiments unless otherwise specified.

Evaluation Metrics. In all of the experiments, we separate a part of each client’s dataset as its test set, and we report the accuracy of the global (aggregated) model on the union of the test sets of clients and the local accuracy of each client on its own local test data. This gives us an indication of how well the global model is representative of the aggregated dataset. We report the classification accuracy in all of the experiments.

3.1 Proof-of-Concept

Non-iid vs. Iid. In this section we evaluate and compare **IDA** with *FedAvg* on f-mnist and cifar-10 datasets given different scenarios of data distribution in clients. Table 1 demonstrates the results of balanced data distribution where all clients have the same or similar number of samples for $n_{cc} \in \{3, 5, 10(iid)\}$ and $pr \in \{30\%, 50\%, 100\%\}$. Our results show that **IDA** has slightly better or on-par performance to *FedAvg* in all scenarios of balanced data distribution.

Ablation Study. In this section, we investigate the effect of different components of the weighting coefficients. We evaluate all of the proposed components on cifar-10 and f-mnist and compare them with two baseline methods, namely *FedAvg*, and another baseline where $\alpha_k = 1$, denoted by *Mean* shown in Table 2. We also evaluate the combination of our weighting method with number of samples per client (**IDA + FedAvg**) and adding the training accuracy of each client to the weighting scheme (**IDA + INTRAC**). The results indicate that combining different weighting schemes can lead to a better performing global model in

Table 1. Comparison between our method and the baseline on cifar10 and f-mnist with different number of classes per client in non-iid and iid scenarios

Dataset	n_{cc}	3c			5c			iid
	pr	30%	50%	100%	30%	50%	100%	100%
Method								
cifar-10	FedAvg	63.20	65.11	69.81	19.68	83.11	80.94	87.77
	IDA	64.36	67.70	70.80	76.06	83.55	83.82	89.46
f-mnist	FedAvg	86.23	87.09	87.45	87.60	87.81	87.16	86.95
	IDA	87.64	87.61	87.44	87.93	87.89	87.46	87.10

FL. This supports our hypothesis, that if some of the clients have lower quality or poisonous models, *FedAvg* would be vulnerable, but our methods can lower the contribution of bad models (overfitted, low quality or poisonous models) so the final model performs better on the federated dataset.

Table 2. Ablation study on different weighting combinations on f-mnist and cifar-10 datasets.

Method	Settings			
	f-mnist $n_{cc} = 3$ $pr = 30\%$		cifar-10 $n_{cc} = 3$ $pr = 30\%$	
	K = 10		K = 10	K = 20
Mean	87.47		65.82	84.80
FedAvg	86.23		63.20	22.84
IDA	87.64		64.36	83.98
IDA + FedAvg	86.67		67.29	82.14
IDA +INTRAC	88.33		64.93	85.23

Sensitivity Analysis. In real-life scenarios, stability of learning process in unfavorable conditions is critical. In *FL* it is not mandatory for the members to contribute in each round, so the participation rate can be different in each round of training, and we might have lower quality models in any round. It is very likely that some clients have very few data samples, and some other clients have a lot of data. In this section we investigate the global model’s performance given low participation rate and severe non-iidness.

Low Participation Rate in Non-iid Distribution. To investigate the effect of participation rate, we used 1000 clients on f-mnist dataset with (batchsize = 30, $lr = 0.016$ and $n_{cc} = 3$ and each client has up to 500 samples). In this experiment, we observe that despite the fact that this dataset is relatively easy to learn,

decreasing the participation rate of clients lowers the performance (*cf.* Fig. 2). When the participation rate is at 1%, the model trained using *FedAvg* collapses. However, when we increase the participation rate to 5% the model continues to learn. We observe a robust performance for both **IDA** and **IDA + FedAvg** in both scenarios.




Fig. 2. Left: participation rate (pr) of 0.01; Right: participation rate of 0.05. The pr affects the stability of federated learning, and it is shown that **IDA** has stable performance comparing to FedAvg.

Severity of Non-IID. To analyze the effect of non-iidness on the performance of our method, we design an experiment by increasing the data samples of the low performing clients. To achieve this, first we train our models in a normal fashion as mentioned in previous sections. Then we choose three clients with the lowest accuracy at the end of the initial training and double the amount of their samples in the training data distribution. We repeat the training using the newly generated data distribution. We propose this experiment to see the effect of *FedAvg* weighting in a scenario where low performing clients are given higher weight. It can be seen in Fig. 3 that before increasing the number of samples, **IDA** performs marginally better compared to other methods; however, after we increase the number of samples in those three clients, *FedAvg* collapses at the beginning of training. Considering the performance of *Mean* aggregation, we see that **IDA** is the main contributing factor to the learning process.

3.2 Clinical Use-Case

We evaluate our proposed method on HAM10k dataset and show our results in Table 3. Even though the global accuracy of the model using **IDA** is on par with *FedAvg*, it can be seen that the local accuracy (accuracy of clients on their own test set) using **IDA** is superior to *FedAvg* in all scenarios. This indicates that **IDA** has a better generalization and lower variance in local accuracy of clients.

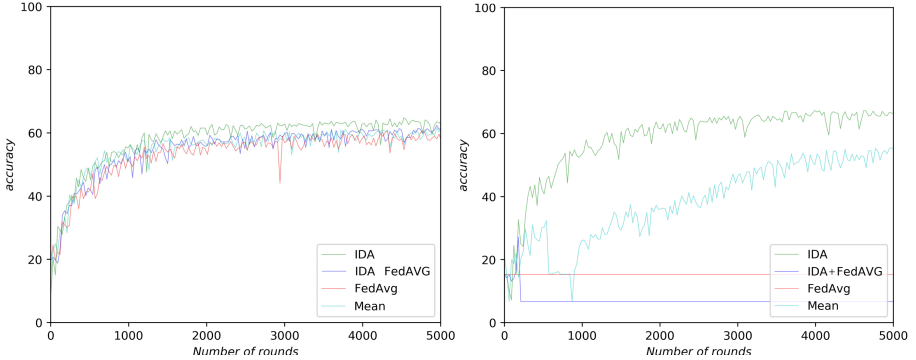


Fig. 3. Accuracy of global model of clients with non-iid data distribution on cifar-10: in the right we have the same clients, and the same learning hyperparameters of the left, but the number of samples in three of the clients with poor performances increased. The local distribution of data points in those three clients remained the same. This experiment is performed on cifar-10 dataset with $K = 10$ clients, $n_{cc} = 3$, $E = 2$, $lr = 0.01$ and random number of samples per class per client up to 1000 samples.

Table 3. Investigation on an unbalanced data distribution among the clients in federated setting, with five random classes per client, and random number of samples per client for HAM10k.

Method	n_{cc}	Global accuracy	Local accuracy
FedAvg	1	69.72	60.52 ± 9.20
IDA	1	69.16	61.21 ± 8.79
FedAvg	2	62.23	57.14 ± 10.84
IDA	2	61.21	60.21 ± 5.48
FedAvg	10 (iid)	63.5	52.88 ± 15.73
IDA	10 (iid)	63.72	57.38 ± 10.56

4 Discussion and Conclusion

In this work, we proposed a novel weighting scheme for aggregation of client models in a federated learning setting for non-iid and unbalanced data distribution. Our weighting is calculated based on the statistical meta-information which gives higher weights in aggregation to the clients that their data has a lower distance to the global average. We also propose another weighting approach called INTRAC that normalizes models to lower the contribution of overfitted models to the shared model. Our extensive experiments show that our proposed method outperforms FedAvg in terms of classification accuracy in non-iid scenario. Our proposed method is also resilient to low quality or poisonous data in the clients. For instance, if the majority of clients are rather aligned, then they can rule out the out-of-distribution models. This is not the case with FedAvg, however, which is based on the presumption that the clients with more data, have a better

distribution compared to other models, and they should have more voting power in the global model. Future research directions concerning the out-of-distribution models detection and robust aggregation schemes should be further considered.

Acknowledgements. S.A. is supported by the PRIME programme of the German Academic Exchange Service (DAAD) with funds from the German Federal Ministry of Education and Research (BMBF). A.F. is supported by Munich Center for Machine Learning (MCML) with funding from the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036B. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan V GPU used for this research.

References

1. Beel, J.: Federated meta-learning: Democratizing algorithm selection across disciplines and software libraries. *Science (AICS)* **210**, 219 (2018)
2. Chen, F., Dong, Z., Li, Z., He, X.: Federated meta-learning for recommendation. arXiv preprint [arXiv:1802.07876](https://arxiv.org/abs/1802.07876) (2018)
3. Chen, Y., Sun, X., Jin, Y.: Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. *IEEE Trans. Neural Netw. Learn. Syst.* (2019)
4. Corinzia, L., Buhmann, J.M.: Variational federated multi-task learning. arXiv preprint [arXiv:1906.06268](https://arxiv.org/abs/1906.06268) (2019)
5. Hsieh, K., Phanishayee, A., Mutlu, O., Gibbons, P.B.: The Non-IID data quagmire of decentralized machine learning. arXiv preprint [arXiv:1910.00189](https://arxiv.org/abs/1910.00189) (2019)
6. 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)
7. Huang, L., et al.: Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records. *J. Biomed. Inform.* **99**, 103291 (2019)
8. Jeong, E., Oh, S., Kim, H., Park, J., Bennis, M., Kim, S.L.: Communication-efficient on-device machine learning: Federated distillation and augmentation under Non-IID private data. arXiv preprint [arXiv:1811.11479](https://arxiv.org/abs/1811.11479) (2018)
9. Jiang, Y., Konečný, J., Rush, K., Kannan, S.: Improving federated learning personalization via model agnostic meta learning. arXiv preprint [arXiv:1909.12488](https://arxiv.org/abs/1909.12488) (2019)
10. Kairouz, P., et al.: Advances and open problems in federated learning. arXiv preprint [arXiv:1912.04977](https://arxiv.org/abs/1912.04977) (2019)
11. Kaissis, G.A., Makowski, M.R., Rückert, D., Braren, R.F.: Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Mach. Intell.*, 1–7 (2020)
12. Konečný, J., McMahan, B., Ramage, D.: Federated optimization: distributed optimization beyond the datacenter. arXiv preprint [arXiv:1511.03575](https://arxiv.org/abs/1511.03575) (2015)
13. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
14. LeCun, Y., et al.: Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1**(4), 541–551 (1989)

15. Li, D., Wang, J.: FedMD: heterogenous federated learning via model distillation. arXiv preprint [arXiv:1910.03581](https://arxiv.org/abs/1910.03581) (2019)
16. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated learning: challenges, methods, and future directions. arXiv preprint [arXiv:1908.07873](https://arxiv.org/abs/1908.07873) (2019)
17. Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., Smith, V.: Federated optimization in heterogeneous networks. arXiv preprint [arXiv:1812.06127](https://arxiv.org/abs/1812.06127) (2018)
18. Li, W., et al.: Privacy-preserving federated brain tumour segmentation. In: Suk, H.-I., Liu, M., Yan, P., Lian, C. (eds.) MLMI 2019. LNCS, vol. 11861, pp. 133–141. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32692-0_16
19. Li, X., Gu, Y., Dvornek, N., Staib, L., Ventola, P., Duncan, J.S.: Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: abide results. arXiv preprint [arXiv:2001.05647](https://arxiv.org/abs/2001.05647) (2020)
20. Liang, P.P., Liu, T., Ziyin, L., Salakhutdinov, R., Morency, L.P.: Think locally, act globally: federated learning with local and global representations. arXiv preprint [arXiv:2001.01523](https://arxiv.org/abs/2001.01523) (2020)
21. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., et al.: Communication-efficient learning of deep networks from decentralized data. arXiv preprint [arXiv:1602.05629](https://arxiv.org/abs/1602.05629) (2016)
22. Pillutla, K., Kakade, S.M., Harchaoui, Z.: Robust aggregation for federated learning. arXiv preprint (2019)
23. Rieke, N., et al.: The future of digital health with federated learning. arXiv preprint [arXiv:2003.08119](https://arxiv.org/abs/2003.08119) (2020)
24. Sattler, F., Müller, K.R., Samek, W.: Clustered federated learning: model-agnostic distributed multi-task optimization under privacy constraints. arXiv preprint [arXiv:1910.01991](https://arxiv.org/abs/1910.01991) (2019)
25. Sattler, F., Wiedemann, S., Müller, K.R., Samek, W.: Robust and communication-efficient federated learning from Non-IID data. *IEEE Trans. Neural Netw. Learn. Syst.* (2019)
26. Sheller, M.J., et al.: Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci. Rep.* **10**(1), 1–12 (2020)
27. Sheller, M.J., Reina, G.A., Edwards, B., Martin, J., Bakas, S.: Multi-institutional deep learning modeling without sharing patient data: a feasibility study on brain tumor segmentation. In: Crimi, A., Bakas, S., Kuijff, H., Keyvan, F., Reyes, M., van Walsum, T. (eds.) *BrainLes 2018*. LNCS, vol. 11383, pp. 92–104. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11723-8_9
28. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
29. Smith, V., Chiang, C.K., Sanjabi, M., Talwalkar, A.S.: Federated multi-task learning. In: *Advances in Neural Information Processing Systems*, pp. 4424–4434 (2017)
30. Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Sci. Data* **5**, 180161 (2018)
31. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint [arXiv:1708.07747](https://arxiv.org/abs/1708.07747) (2017)
32. Xu, J., Wang, F.: Federated learning for healthcare informatics. arXiv preprint [arXiv:1911.06270](https://arxiv.org/abs/1911.06270) (2019)
33. Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., Chandra, V.: Federated learning with Non-IID data. arXiv preprint [arXiv:1806.00582](https://arxiv.org/abs/1806.00582) (2018)