

Benchmarking FedAvg and FedCurv for Image Classification Tasks

Bruno Casella¹, Roberto Esposito¹, Carlo Cavazzoni² and Marco Aldinucci¹

¹*Department of Computer Science, University of Torino, Torino, Italy*

²*Leonardo S.p.A., Italy*

Abstract

Classic Machine Learning (ML) techniques require training on data available in a single data lake (either centralized or distributed). However, aggregating data from different owners is not always convenient for different reasons, including security, privacy and secrecy. Data carry a value that might vanish when shared with others; the ability to avoid sharing the data enables industrial applications where security and privacy are of paramount importance, making it possible to train global models by implementing only local policies which can be run independently and even on air-gapped data centres. Federated Learning (FL) is a distributed machine learning approach which has emerged as an effective way to address privacy concerns by only sharing local AI models while keeping the data decentralized. Two critical challenges of Federated Learning are managing the heterogeneous systems in the same federated network and dealing with real data, which are often not independently and identically distributed (non-IID) among the clients. In this paper, we focus on the second problem, i.e., the problem of statistical heterogeneity of the data in the same federated network. In this setting, local models might be strayed far from the local optimum of the complete dataset, thus possibly hindering the convergence of the federated model. Several Federated Learning algorithms, such as FedAvg, FedProx and Federated Curvature (FedCurv), aiming at tackling the non-IID setting, have already been proposed. This work provides an empirical assessment of the behaviour of FedAvg and FedCurv in common non-IID scenarios. Results show that the number of epochs per round is an important hyper-parameter that, when tuned appropriately, can lead to significant performance gains while reducing the communication cost. As a side product of this work, we release the non-IID version of the datasets we used so to facilitate further comparisons from the FL community.

Keywords

Federated Learning, Federated Curvature, FedCurv, non-IID

1. Introduction

The increasing availability of big data and computational resources supports the growth of accurate and reliable Machine Learning (ML) and Deep Learning (DL) models. More and more sectors, from the public to companies, follow data-driven approaches in their everyday decisions.

ITADATA 2022: The 1st Italian Conference on Big Data and Data Science, September 20–21, 2022, Milan, Italy

✉ bruno.casella@unito.it (B. Casella); roberto.esposito@unito.it (R. Esposito); carlo.cavazzoni@leonardo.com (C. Cavazzoni); marco.aldinucci@unito.it (M. Aldinucci)

🌐 <https://alpha.di.unito.it/bruno-casella/> (B. Casella); <https://www.unito.it/persona/esposito> (R. Esposito); <https://alpha.di.unito.it/marco-aldinucci/> (M. Aldinucci)

🆔 0000-0002-9513-6087 (B. Casella); 0000-0001-5366-292X (R. Esposito); 0000-0002-9589-4785 (C. Cavazzoni); 0000-0001-8788-0829 (M. Aldinucci)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

One of the difficulties in building systems that support these decisions is the necessity to obtain large datasets for training the models. In fact, in many situations, data is scattered between many parties, and it would be beneficial to merge it into a single place in order to gather the information needed to build high-quality models. However, due to increasing privacy as well as security concerns merging the data might not be a feasible approach. For instance, the European GDPR [1] regulation adds strong constraints on the possibility of sharing data between parties when sensible pieces of information are at stake. Moreover, industrial data are often not shared even when allowed by the law because they represent an essential advantage over competitors or because of the security concerns deriving from the need to expose part of the data to other parties. Consequently, many players resort to train their models on local datasets, and the resulting AI models often suffer from low reliability, leading to generalizability and overfitting issues [2].

Federated Learning (FL) [3] is a machine learning setting that emerged as an effective way to train on distributed datasets without requiring any exchange of data-related information between the parties. This not only solves the privacy concerns, but also allows the parties to implement local policies that can be deployed more securely even in extreme settings such those required by air gapped systems.

In a typical FL scenario, at each round, multiple clients take one or more steps of gradient descent on a shared model using their local data, and then a central node acts as an aggregator performing a weighted average (Federated Averaging [3]) of the resulting models. This approach generally performs well when the training data distributions are independent and identically distributed (IID). However, in a real-world scenario, due to the natural differences in the collection of data, FL faces difficulties due to data being non-independent and identically distributed (non-IID) among the involved parties. In these cases, the data of a single client does not represent the overall distribution of data among the federation, and this poses a key challenge for FL [4, 5]. For example, in image recognition, models trained on sunny days may not be accurate on cloudy days due to an unrecognizable feature distribution.

The first and most used FL algorithm is FedAvg, proposed by Google in 2017, which can actually work on non-IID data when strong assumptions hold on to the loss function. Specifically, a recent work [6] shows that if the optimization problem is strongly convex and smooth, then FedAvg would converge even non-IID data. FedAvg averages local gradients of the nodes of the federation. Recent years have seen the deployment of several FL algorithms for addressing the non-IID setting, like FedProx [7], FedNova [8] and SCAFFOLD [9]. All these algorithms have already been tested in [5], which provides valuable information about their behaviour in non-IID settings. In this paper, we compare the performance of FedAvg with FedCurv [10] on several common non-IID settings. FedCurv is an algorithm that addresses the problem of catastrophic forgetting [11] in FL by adding a penalty term to the loss function in order to compel the local models to a shared optimum. We tested both algorithms on three public image datasets, i.e. MNIST [12], CIFAR10 [13], and MedMNIST [14]. We manipulate the dataset to simulate common non-IID settings and show that performance is often related to the number of epochs performed locally in each round of the FL algorithm.

The main contributions of this work are:

- we provide a benchmark datasets for five different non-IID settings (see Section 4):

- quantity skew, three versions of prior shift and covariate shift;
- we provide results of extensive experiments on FedAvg and FedCurv on the considered non-IID settings as well as on the IID case. To the best of our knowledge, this is the first work providing empirical evidence on the behaviour of FedCurv in common non-IID cases;
 - we discuss the results under several points of view, thus providing a faceted understanding of the behaviour of the algorithms;
 - as a result of the experimentation, we show that the number of epochs per round can be pivotal in obtaining good performances while limiting the communication costs.

The rest of the paper is organized as follows. In Section 2, we present the recent related works. Section 3 reviews the FL algorithms tackling non-iidness. In Section 5 are given some of the most typical non-IID partition strategies, and experimental results are shown and discussed. Finally, in Section 6, conclusions are drawn.

2. Related Work

Dealing with non-IID data represents one of the fundamental challenges in FL. [4] surveys this issue providing several non-IID data cases.

To the best of our knowledge, there are only a few benchmarks for FL in the non-IID scenario. One work [15] tests only FedAvg on MNIST and CIFAR10 in the quantity skew and labels quantity skew settings. Another work [5] considers quantity skew, labels quantity skew and three types of feature distribution skew, i.e. noise-based, synthetic and real-world feature imbalance. It tests FedAvg, FedProx [7], FedNova [8] and SCAFFOLD [9] on 9 public image datasets, including MNIST and CIFAR10. This work points out how non-iidness represents a challenge for FL algorithms in terms of accuracy, showing how none of those algorithms outperforms others in all the tested cases.

A more recent work [16] consider proposes three FL algorithms based on distributed boosting strategies rather than gradient-descent based methods: *DistBoost.F* and *PreWeak.F*, that are the adaptation to the FL setting of the distributed boosting algorithms DistBoost [17] and PreWeak [18], and *AdaBoost.F*, that is a novel algorithm developed for FL. The authors performed experiments on ten tabular UCI datasets, comparing their proposed algorithms in the same non-iid settings of this paper. This work shows that in most cases, the aggregated model outperforms the models that could have been learned on local data, but that non-iidness can hurt the efficiency of the federation.

Apart from these works, there are some benchmarks for FL that do not focus on non-IID data issues. LEAF [19] is a modular benchmarking framework for FL, which includes a set of both image and tabular datasets, an evaluation framework and a set of reference implementations. Some open-source datasets they provided are *FEMNIST* (Federated Extended MNIST), which is built by partitioning the data in Extended MNIST [20] based on the writer of the digit/character, and *Shakespeare*, a dataset built from *The Complete Works of William Shakespeare* [21], where each speaking role in each play is considered a different device. The Open Application Repository for Federated Learning (OARF) [22] is a benchmark for FL mimicking realistic scenarios with publicly available datasets as different data silos in image, text and structured data. FedML [23] is

an open research library and benchmark developed for fair algorithm performance comparison. It provides algorithm development of FedAvg, FedNova and FedOpt [24].

3. FL Algorithms on non-IID data

In a non-IID setting, the local data of a single client can not represent the overall distribution of the federation. In such a situation, the local models drift apart because local optima may be far from the global optima. This results in a reduction of accuracy of both local models, which are driven towards the local optima, and of the aggregated model because local updates may be large, in particular when each round has a lot of local epochs. In such conditions, the aggregated model can have a worse accuracy than models learnt using only local data.

Several approaches have already been proposed to tackle the problem of non-iidness. The most popular algorithm for FL in the non-IID settings are FedAvg [3], FedCurv [10], FedProx [7], FedNova [8] and SCAFFOLD [9]. The main features of these algorithms are described below.

FedAvg: Federated Averaging [3] has been the first algorithm to be proposed for the FL setting. Basically, the shared model parameters are initialized by the aggregator at the beginning of the training. Afterwards, each client trains a local copy of the model on its local data and sends the result to the server, which sets the shared model to be the (weighted) average of the received local models. Then the aggregator sends the shared model again to each client, and the process repeats. Typically in cross-device settings, the server sends the global model to a random subset of clients in order to cope with the high number of parties in the federation, which can be in the order of thousands or even more. Two parameters that can be controlled and might have a large impact on the results are the number of local training epochs per round (E) and the local batch size (B). If $B = \infty$ and $E = 1$ the algorithm is called FedSGD [3]. With $E > 1$, the number of communication rounds can decrease; however, local models can be driven towards the local optima, leading to bad accuracy. In [6] it is shown that if the loss is strongly convex and smooth, the rate of convergence of FedAvg is $O(1/T)$ (where T is the number of rounds) and that weight decay is a necessary condition for optimal convergence.

FedCurv: Federated Curvature (FedCurv) [10] builds on the idea of Lifelong Learning [25] to prevent catastrophic forgetting [11] in FL. It is an adaptation for FL of Elastic Weight Consolidation (EWC) [26], an algorithm for sequentially training a model on new tasks without forgetting old ones. The basic assumption of EWC is that neural networks are sufficiently over-parameterized that there are good chances of finding an optimal solution B^* to task B in the neighbourhood of a solution A^* learned on a previous task A. EWC uses the diagonal of the Fisher information matrix to choose the most important weights of the previous task. EWC adds a penalty term to the optimization function in order to force the model parameters with the higher Fisher information for task A to maintain their actual value while learning task B. FedCurv adds this EWC penalty term for minimizing the model disparity across the clients of a federation. During each round, FedCurv works just like FedAvg, but each client sends its local model together with the diagonal of Fisher's information matrix. Since this method allows for less frequent communication fewer steps are required in order to reach the desired accuracy. However, in each round, the number of parameters to be transmitted is about three times that

in FedAvg.

FedProx: FedProx [7] is a re-parametrization of FedAvg aiming to tackle two key challenges in FL: systems heterogeneity (variability of the devices) and statistical heterogeneity (non-iidness). FedProx alleviates systems heterogeneity starting from the idea of allowing a variable amount of work across devices based on their resource constraints. In the case of statistical heterogeneity, local models may drift apart. FedProx restricts the size of local updates by adding an L2 regularization term in the cost function compelling the local model to stay close to the global model. L2 regularization is controlled by an additional term μ that needs to be tuned carefully. FedProx does not introduce communication overhead, but it increases the computation overhead of the devices.

FedNova: Federated Normalized Averaging (FedNova) [8] proposes a slightly modified version of FedAvg to overcome the problem of objective inconsistency while preserving fast error convergence. The paradigm is the same as FedAvg, but FedNova normalizes and scales local updates of the parties based on their number of local steps (mini-batches of local training).

SCAFFOLD: Stochastic Controlled Averaging for federated learning (SCAFFOLD) [9] proposes a solution for client-drift. SCAFFOLD requires significantly fewer communication rounds, however, it doubles the communication size per round due to the additional control variates.

4. Non-IID data

An existing study [4] identifies five different non-IID scenarios: 1) prior shift (label distribution skew), 2) covariate shift (feature distribution skew), 3) same label but different features, 4) same features but different labels, and 5) quantity skew. Many of these cases (1, 2 and 5) have already been tested in [7]. In our work, we focus on the iid case (uniform data distribution) and on five types of non-iidness: quantity skew, three versions of prior shift and covariate shift (feature distribution skew). These scenarios have already been tested in [16] with tabular datasets and non-gradient-descent methods. We briefly describe the adopted settings.

Quantity Skew: a collection of datasets exhibits a quantity skew if different clients (datasets) can hold vastly different amounts of data. In this case, the sampling distribution may still be consistent among the parties; it is, however, interesting to see the effect of the quantity imbalance on the convergence of the FL algorithm. In this case the proportion of example to be assigned to client $x \in \{0 \dots N - 1\}$ is determined by the Power distribution $P(x; \alpha) = \alpha x^{\alpha-1}$. Here $\alpha > 0$ is a parameter affecting the distribution as it follows: if $\alpha = 1$ then examples are uniformly distributed across clients; if $\alpha = 2$ examples are "linearly" distributed across users (our case); if $\alpha \geq 3$ the examples are power law distributed; in general, as $\alpha \rightarrow \infty$, it is increasingly true that a single user obtains most of the examples, and the remaining users only get very few examples.

Prior shift: let us consider the conjunct distribution of data and labels $P(x_i, y_i) = P(x_i|y_i)P(y_i)$, a collection of datasets exhibit a prior shift if the labels' prior $P(y_i)$ varies across clients. This is a common scenario in many real-world FL applications. For instance, labels' prior can vary when devices are spread across different world regions, and certain elements are present only in some countries. The types of prior shift implemented are the following:

Labels quantity skew: this is the simplest version of prior shift. Labels are partitioned

between clients and each client receives samples that belongs to only a fixed number of classes. In our experiments, we fixed the number of classes per client to 2.

Dirichlet labels skew: in this case the number of examples of a given class that each client receives is distributed according to samples extracted from the Dirichlet distribution. More specifically, for each class k , we extract an N dimensional vector p_k from $\text{Dir}_N([\beta, \dots, \beta])$ and assign a proportion of $p_{k,j}$ samples of class k to client j . The N -dimensional vector of β s is the *concentration parameter*. The lower β , the greater the imbalance. In our experiments we fixed $\beta = 0.5$ as in [16, 7].

Pathological labels skew: this skewness was designed in [3]. First of all, data are sorted by label and divided into shards of size $N \cdot \text{shards_per_client}$. Each client owns shards_per_client shards. We fixed $\text{shards_per_client} = 2$. Basically, in this case, most parties will only have a limited number of labels.

Covariate shift: in covariate shift, also known as distribution skew, the marginal distributions $P(x_i)$ may vary across clients, even if $P(y|x)$ is shared. It is common to encounter covariate shift in several fields of machine learning. For example, in speech recognition, a model trained in a specific language can have trouble when it encounters particular accents or dialects of that language; in image recognition, a model trained on sunny days may not be accurate on cloudy days due to an unrecognizable feature distribution. To simulate covariate shift on our datasets, we followed the procedure outlined in [16], where examples are assigned to clients according to a distribution based on the results of a Principal Component Analysis (PCA) performed on the data.

5. Experiments

To conduct our experiments, we adopted Open Federated Learning (OpenFL) [27], the new framework for FL developed by the Intel Internet of Things Group (IOTG) and Intel Labs. OpenFL is a Python 3 library for FL that is Deep Learning framework-agnostic. Training of statistical models may be done with any deep learning framework, such as TensorFlow or PyTorch, via a plugin mechanism. OpenFL employs an Aggregator-Collaborator workflow: a Collaborator is a client in the federation that trains a global model on a local dataset, while the Aggregator aggregates the model updates received from Collaborators. All the experiments were computed in a distributed environment encompassing $N = 10$ collaborators. Each collaborator is run on an Intel Xeon CPU (8 cores per CPU).

Dataset: We compared FedAvg and FedCurv on MNIST [12], CIFAR10 [13] and MedMNIST [14]. CIFAR10 and MNIST are default benchmarks in NN literature; MedMNIST has been selected due to the increasing interest of DL and FL in medical domains. In particular, we used OrganAMNIST, a 2D dataset on Abdominal CT contained in MedMNIST. The details of the datasets are summarized in Table I.

Preprocessing: 2D datasets MNIST and MedMNIST were kept at 28x28, while CIFAR10 was rescaled to 64x64. As for data augmentation, we performed random horizontal flips and angle rotation of 10° with a probability of 80%. All the datasets were normalized according to their mean and standard deviation.

Table 1
Statistics of the datasets.

Dataset	Train samples	Test samples	# labels
MNIST	60.000	10.000	10
CIFAR10	50.000	10.000	10
MedMNIST	34.581	17.778	11

Table 2
Classification accuracy in the non-federated setting.

Datasets	10 epochs	100 epochs
MNIST	98.67%	99.12%
CIFAR10	64.03%	71.25%
MedMNIST	84.85%	89.13%

Model: We employed ResNet-18 [28] as classification model, trained by minimizing the cross-entropy loss with mini-batch gradient descent using the Adam optimizer with learning rate 10^{-4} . The local batch size was 64.

The algorithms have been compared using the standard classification metric of top-1 accuracy. Results are reported in tables 2 to 6, with table 2 reporting about a non-federated baseline, i.e., the typical AI scenario where the data are aggregated in a single device. The remaining tables reports the results about experiments in different non-iid settings of the two tested federated algorithms.

5.1. Discussion

The results can be analyzed from different points of view, showing interesting insights:

Epochs per round perspective: It can be noted that the accuracy increases as the number of epochs per round E increases. This means that local optima may be close to the global optima, and so, with the same amount of rounds, training for more epochs can be beneficial. This pattern is clearly shown for each setting and algorithm, except for the FedAvg on MedMNIST in the uniform setting (Table 3) and in the case of the Labels Quantity Skew (Table 5).

Distribution perspective: prior shift (see table 5) is the most challenging non-IID setting. In particular, among the different types of prior shift analyzed, the labels quantity skew is the most detrimental. Both FedAvg and FedCurv perform worse on the labels quantity skew. By design, the pathological labels skew is the most similar to the labels quantity skew because, in both cases, each client has examples that belong to only a small subset of the possible classes. Indeed, the pathological labels skew is the second hardest scenario after the labels quantity skew. The Dirichlet labels skew is the less challenging prior shift case. FedAvg and FedCurv perform well on the quantity skew setting (see table 4). This is reasonable because both algorithms adopt

Table 3

Comparison between FedAvg and FedCurv in the uniform setting.

Datasets	Epochs	10 rounds		100 rounds	
		FedAvg	FedCurv	FedAvg	FedCurv
MNIST	1	97.16%	97.17%	97.66%	97.53%
	10	99.05%	99.07%	99.35%	99.33%
	30	99.28%	99.33%	99.47%	99.55%
CIFAR10	1	56.34%	56.13%	54.89%	54.78%
	10	70.95%	70.40%	74.03%	75.57%
	30	74.05%	74.07%	78.89%	78.57%
MedMNIST	1	68.98%	79.70%	72.42%	83.50%
	10	68.19%	86.90%	72.96%	89.49%
	30	46.00%	88.77%	71.16%	90.23%
# best performance		2	7	4	5

Table 4

Comparison between FedAvg and FedCurv in the quantity skew setting.

Datasets	Epochs	10 rounds		100 rounds	
		FedAvg	FedCurv	FedAvg	FedCurv
MNIST	1	96.33%	96.72%	96.84%	97.24%
	10	99.07%	99.15%	99.37%	99.40%
	30	99.48%	99.38%	99.49%	99.42%
CIFAR10	1	56.24%	55.13%	56.10%	54.73%
	10	70.49%	70.35%	74.67%	73.46%
	30	74.67%	73.42%	76.91%	76.24%
MedMNIST	1	80.42%	79.44%	83.24%	84.28%
	10	87.91%	87.23%	90.24%	88.96%
	30	89.08%	89.19%	89.88%	90.31%
# best performance		6	3	5	4

a weighted averaging of the parameters, and the distribution of the examples (except for the quantity of the examples) is uniform among parties, which is the easiest setting and the one that is more similar to the non-federated scenario. The covariate shift (see table 6) presents only a low loss of accuracy.

Algorithm perspective: despite FedCurv was born for tackling non-IID data in FL, it obtains better results than FedAvg in the uniform (Table 3) and in the covariate shift settings (Table 6). FedAvg performs well on quantity skew and prior shift scenarios, confirming that it can work on non-IID data [6]. It is interesting to note that most of the time, FedCurv wins after 100 rounds, showing that it may require more training steps to convergence.

Table 5
Comparison between FedAvg and FedCurv in the prior shift setting.

Category	Datasets	Epochs	10 rounds		100 rounds	
			FedAvg	FedCurv	FedAvg	FedCurv
Labels Quantity Skew	CIFAR10	1	41.36%	39.93%	47.60%	37.06%
		10	45.76%	39.77%	57.65%	46.08%
		30	41.48%	26.91%	62.52%	39.65%
	MedMNIST	1	53.64%	48.75%	61.82%	54.01%
		10	46.03%	38.98%	60.03%	65.68%
		30	37.50%	53.90%	58.12%	56.65%
Dirichlet Labels Skew	CIFAR10	1	48.68%	47.70%	48.50%	48.37%
		10	62.39%	62.77%	70.44%	71.50%
		30	67.10%	67.01%	75.54%	74.83%
	MedMNIST	1	76.27%	76.25%	81.62%	81.27%
		10	85.54%	84.92%	88.74%	89.40%
		30	86.92%	86.93%	90.49%	90.64%
Pathological Labels Skew	CIFAR10	1	40.52%	35.20%	47.42%	45.60%
		10	53.52%	60.57%	64.05%	64.33%
		30	48.09%	59.62%	61.58%	67.09%
	MedMNIST	1	64.42%	60.93%	72.27%	70.52%
		10	65.83%	57.13%	78.81%	71.23%
		30	80.36%	74.65%	83.41%	84.95%
# best performance			13	5	11	7

Communication perspective: it seems that, with the same amount of epochs, less communication achieves better results. For example, in each setting (apart from the labels quantity skew case), after ten rounds of ten epochs, i.e. 100 epochs, both FedAvg and FedCurv have better accuracy than 100 rounds of one epoch. This means that, perhaps counter-intuitively, training locally before performing aggregation can boost the model’s accuracy. This seems to indicate that pursuing local optimizations can lead to better approximations of the local optima. Why this is the case is an interesting avenue for future investigation.

6. Conclusions

In this paper, we experimented with two federated Learning algorithms in five different non-IID settings. In our experiments, neither of the two algorithms outperforms the other in all the partitioning strategies. However, somewhat unexpectedly, FedAvg produced better models in a majority of non-IID settings despite competing with an algorithm that was explicitly developed to improve in this scenario. Interestingly, both algorithms seem to perform better when the number of epochs per round is increased (which also has the benefit of reducing the communication cost). This is, to the best of our knowledge, a new observation, and we aim to

Table 6

Comparison between FedAvg and FedCurv in the covariate shift setting.

Datasets	Epochs	10 rounds		100 rounds	
		FedAvg	FedCurv	FedAvg	FedCurv
CIFAR10	1	46.82%	43.55%	48.90%	45.32%
	10	62.57%	63.28%	66.42%	68.90%
	30	69.04%	64.56%	73.14%	71.28%
MedMNIST	1	75.07%	77.70%	79.34%	82.32%
	10	84.46%	84.59%	87.16%	88.18%
	30	85.26%	84.99%	89.41%	89.45%
# best performance		3	3	2	4

investigate it in the future. Among the datasets we tested, the ones implementing the quantity and pathological labels skews are those posing the hardest challenges to the algorithms. Also, as expected, the quantity skew appears to be the less challenging type of skew. In future work, we aim to collect further datasets and expand the number of FL algorithms we test to provide a comprehensive picture of the state of the art of FL in non-IID settings.

Acknowledgements

This has been partially supported by the *European PILOT* project, which has received funding from the EuroHPC JU under grant agreement No.101034126. The JU receives support from the European Union’s Horizon 2020 research and innovation programme and Spain, Italy, Switzerland, Germany, France, Greece, Sweden, Croatia and Turkey.

References

- [1] P. Voigt, A. v. d. Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 1st ed., Springer Publishing Company, Incorporated, 2017.
- [2] J. R. Zech, M. A. Badgeley, M. Liu, A. B. Costa, J. J. Titano, E. K. Oermann, Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study, *PLOS Medicine* 15 (2018) 1–17. doi:10.1371/journal.pmed.1002683.
- [3] B. McMahan, E. Moore, D. Ramage, S. Hampson, B. A. y Arcas, Communication-efficient learning of deep networks from decentralized data, in: A. Singh, X. J. Zhu (Eds.), *Proc. of the 20th Intl. Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, volume 54 of *Proc. of Machine Learning Research*, PMLR, 2017, pp. 1273–1282. URL: <http://proceedings.mlr.press/v54/mcmahan17a.html>.
- [4] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. A. Bonawitz, Z. Charles, G. Cormode, R. Cummings, R. G. L. D’Oliveira, H. Eichner, S. E. Rouayheb,

- D. Evans, J. Gardner, Z. Garrett, A. Gascón, B. Ghazi, P. B. Gibbons, M. Gruteser, Z. Harchaoui, C. He, L. He, Z. Huo, B. Hutchinson, J. Hsu, M. Jaggi, T. Javidi, G. Joshi, M. Khodak, J. Konečný, A. Korolova, F. Koushanfar, S. Koyejo, T. Lepoint, Y. Liu, P. Mittal, M. Mohri, R. Nock, A. Özgür, R. Pagh, H. Qi, D. Ramage, R. Raskar, M. Raykova, D. Song, W. Song, S. U. Stich, Z. Sun, A. T. Suresh, F. Tramèr, P. Vepakomma, J. Wang, L. Xiong, Z. Xu, Q. Yang, F. X. Yu, H. Yu, S. Zhao, Advances and open problems in federated learning, *Found. Trends Mach. Learn.* 14 (2021) 1–210. doi:10.1561/22000000083.
- [5] H. Zhu, J. Xu, S. Liu, Y. Jin, Federated learning on non-iid data: A survey, *Neurocomputing* 465 (2021) 371–390. doi:10.1016/j.neucom.2021.07.098.
- [6] X. Li, K. Huang, W. Yang, S. Wang, Z. Zhang, On the convergence of fedavg on non-iid data, *CoRR abs/1907.02189* (2019). URL: <http://arxiv.org/abs/1907.02189>. arXiv:1907.02189.
- [7] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, V. Smith, Federated optimization in heterogeneous networks, in: I. S. Dhillon, D. S. Papailiopoulos, V. Sze (Eds.), *Proc. of Machine Learning and Systems 2020, MLSys 2020*, Austin, TX, USA, March 2-4, 2020, mlsys.org, 2020.
- [8] J. Wang, Q. Liu, H. Liang, G. Joshi, H. V. Poor, Tackling the objective inconsistency problem in heterogeneous federated optimization, in: H. Larochelle, M. Ranzato, R. Hassel, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, December 6-12, 2020, virtual, 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/564127c03caab942e503ee6f810f54fd-Abstract.html>.
- [9] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, A. T. Suresh, SCAFFOLD: stochastic controlled averaging for federated learning, in: *Proc. of the 37th Intl. Conference on Machine Learning, ICML 2020*, 13-18 July 2020, Virtual Event, volume 119 of *Proc. of Machine Learning Research*, PMLR, 2020, pp. 5132–5143. URL: <http://proceedings.mlr.press/v119/karimireddy20a.html>.
- [10] N. Shoham, T. Avidor, A. Keren, N. Israel, D. Benditkis, L. Mor-Yosef, I. Zeitak, Overcoming forgetting in federated learning on non-iid data, *CoRR abs/1910.07796* (2019). arXiv:1910.07796.
- [11] I. J. Goodfellow, M. Mirza, X. Da, A. C. Courville, Y. Bengio, An empirical investigation of catastrophic forgetting in gradient-based neural networks, in: Y. Bengio, Y. LeCun (Eds.), *2nd International Conference on Learning Representations, ICLR 2014*, Banff, AB, Canada, April 14-16, 2014, *Conference Track Proceedings*, 2014. URL: <http://arxiv.org/abs/1312.6211>.
- [12] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (1998) 2278–2324. doi:10.1109/5.726791.
- [13] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images (2009) 32–33. URL: <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- [14] J. Yang, R. Shi, D. Wei, Z. Liu, L. Zhao, B. Ke, H. Pfister, B. Ni, Medmnist v2: A large-scale lightweight benchmark for 2d and 3d biomedical image classification, *CoRR abs/2110.14795* (2021). URL: <https://arxiv.org/abs/2110.14795>. arXiv:2110.14795.
- [15] L. Liu, F. Zhang, J. Xiao, C. Wu, Evaluation framework for large-scale federated learning, *CoRR abs/2003.01575* (2020). URL: <https://arxiv.org/abs/2003.01575>. arXiv:2003.01575.
- [16] M. Polato, R. Esposito, M. Aldinucci, Boosting the federation: Cross-silo federated learning without gradient descent, in: *International Joint Conference on Neural Networks, IJCNN*

- 2022, Padua, Italy, July 18-23, 2022, IEEE, 2022.
- [17] A. Lazarevic, Z. Obradovic, Boosting algorithms for parallel and distributed learning, *Distributed Parallel Databases* 11 (2002) 203–229. doi:10.1023/A:1013992203485.
 - [18] J. Cooper, L. Reyzin, Improved algorithms for distributed boosting, in: *55th Annual Allerton Conference on Communication, Control, and Computing*, Allerton 2017, Monticello, IL, USA, October 3-6, 2017, IEEE, 2017, pp. 806–813. doi:10.1109/ALLERTON.2017.8262822.
 - [19] S. Caldas, P. Wu, T. Li, J. Konečný, H. B. McMahan, V. Smith, A. Talwalkar, LEAF: A benchmark for federated settings, *CoRR abs/1812.01097* (2018). URL: <http://arxiv.org/abs/1812.01097>. arXiv:1812.01097.
 - [20] G. Cohen, S. Afshar, J. Tapson, A. van Schaik, EMNIST: extending MNIST to handwritten letters, in: *2017 International Joint Conference on Neural Networks, IJCNN 2017*, Anchorage, AK, USA, May 14-19, 2017, IEEE, 2017, pp. 2921–2926. doi:10.1109/IJCNN.2017.7966217.
 - [21] W. Shakespeare, *The complete works of William Shakespeare*, Race Point Publishing, 2014.
 - [22] S. Hu, Y. Li, X. Liu, Q. Li, Z. Wu, B. He, The OARF benchmark suite: Characterization and implications for federated learning systems, *CoRR abs/2006.07856* (2020). URL: <https://arxiv.org/abs/2006.07856>. arXiv:2006.07856.
 - [23] C. He, S. Li, J. So, M. Zhang, H. Wang, X. Wang, P. Vepakomma, A. Singh, H. Qiu, L. Shen, P. Zhao, Y. Kang, Y. Liu, R. Raskar, Q. Yang, M. Annavaram, S. Avestimehr, Fedml: A research library and benchmark for federated machine learning, *CoRR abs/2007.13518* (2020). URL: <https://arxiv.org/abs/2007.13518>. arXiv:2007.13518.
 - [24] S. J. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, H. B. McMahan, Adaptive federated optimization, in: *9th International Conference on Learning Representations, ICLR 2021*, Virtual Event, Austria, May 3-7, 2021, OpenReview.net, 2021. URL: <https://openreview.net/forum?id=LkFG3lB13U5>.
 - [25] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, S. Wermter, Continual lifelong learning with neural networks: A review, *Neural Networks* 113 (2019) 54–71. doi:10.1016/j.neunet.2019.01.012.
 - [26] J. Kirkpatrick, R. Pascanu, N. C. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, R. Hassel, Overcoming catastrophic forgetting in neural networks, *CoRR abs/1612.00796* (2016). arXiv:1612.00796.
 - [27] G. A. Reina, A. Gruzdev, P. Foley, O. Perepelkina, M. Sharma, I. Davidyuk, I. Trushkin, M. Radionov, A. Mokrov, D. Agapov, J. Martin, B. Edwards, M. J. Sheller, S. Pati, P. N. Moorthy, H. S. Wang, P. Shah, S. Bakas, Openfl: An open-source framework for federated learning, *CoRR abs/2105.06413* (2021). arXiv:2105.06413.
 - [28] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016*, Las Vegas, NV, USA, June 27-30, 2016, IEEE Computer Society, 2016, pp. 770–778. doi:10.1109/CVPR.2016.90.