

# On the effectiveness of partial variance reduction in federated learning with heterogeneous data

Bo Li\*    Mikkel N. Schmidt    Tommy S. Alstrøm  
 Technical University of Denmark  
 {bli, mns, tsal}@dtu.dk

Sebastian U. Stich  
 CISPA<sup>‡</sup>  
 stich@cispa.de

## Abstract

Data heterogeneity across clients is a key challenge in federated learning. Prior works address this by either aligning client and server models or using control variates to correct client model drift. Although these methods achieve fast convergence in convex or simple non-convex problems, the performance in over-parameterized models such as deep neural networks is lacking. In this paper, we first revisit the widely used FedAvg algorithm in a deep neural network to understand how data heterogeneity influences the gradient updates across the neural network layers. We observe that while the feature extraction layers are learned efficiently by FedAvg, the substantial diversity of the final classification layers across clients impedes the performance. Motivated by this, we propose to correct model drift by variance reduction only on the final layers. We demonstrate that this significantly outperforms existing benchmarks at a similar or lower communication cost. We furthermore provide proof for the convergence rate of our algorithm.

## 1. Introduction

Federated learning (FL) is emerging as an essential distributed learning paradigm in large-scale machine learning. Unlike in traditional machine learning, where a model is trained on the collected centralized data, in federated learning, each client (e.g. phones and institutions) learns a model with its local data. A centralized model is then obtained by aggregating the updates from all participating clients without ever requesting the client data, thereby ensuring a certain level of user privacy [13, 17]. Such an algorithm is especially beneficial for tasks where the data is sensitive, e.g. chemical hazards detection and diseases diagnosis [33].

Two primary challenges in federated learning are i) han-

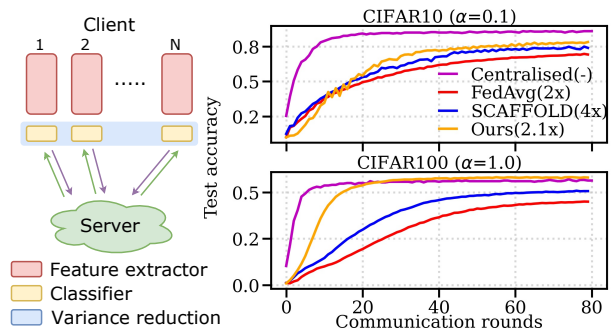


Figure 1. Our proposed FedPVR framework with the performance (communicated parameters per round client $\leftrightarrow$ server). Smaller  $\alpha$  corresponds to higher data heterogeneity. Our method achieves a better speedup than existing approaches by transmitting a slightly larger number of parameters than FedAvg.

dling data heterogeneity across clients [13] and ii) limiting the cost of communication between the server and clients [10]. In this setting, FedAvg [17] is one of the most widely used schemes: A server broadcasts its model to clients, which then update the model using their local data in a series of steps before sending their individual model to the server, where the models are aggregated by averaging the parameters. The process is repeated for multiple communication rounds. While it has shown great success in many applications, it tends to achieve subpar accuracy and convergence when the data are heterogeneous [14, 24, 31].

The slow and sometimes unstable convergence of FedAvg can be caused by client drift [14] brought on by data heterogeneity. Numerous efforts have been made to improve FedAvg’s performance in this setting. Prior works attempt to mitigate client drift by penalizing the distance between a client model and the server model [20, 31] or by performing variance reduction techniques while updating client models [1, 14, 32]. These works demonstrate fast convergence on convex problems or for simple neural networks; however, their performance on deep neural

\* Work done while at CISPA

<sup>‡</sup> CISPA Helmholtz Center for Information Security

networks, which are state-of-the-art for many centralized learning tasks [11,34], has yet to be well explored. Adapting techniques that perform well on convex problems to neural networks is non-trivial [7] due to their “intriguing properties” [38] such as over-parametrization and permutation symmetries.

To overcome the above issues, we revisit the FedAvg algorithm with a deep neural network (VGG-11 [34]) under the assumption of data heterogeneity and full client participation. Specifically, we investigate which layers in a neural network are mostly influenced by data heterogeneity. We define *drift diversity*, which measures the diversity of the directions and scales of the averaged gradients across clients per communication round. We observe that in the non-IID scenario, the deeper layers, especially the final classification layer, have the highest diversity across clients compared to an IID setting. This indicates that FedAvg learns good feature representations even in the non-IID scenario [5] and that the significant variation of the deeper layers across clients is a primary cause of FedAvg’s subpar performance.

Based on the above observations, we propose to align the classification layers across clients using variance reduction. Specifically, we estimate the average updating direction of the classifiers (the last several fully connected layers) at the client  $c_i$  and server level  $c$  and use their difference as a control variate [14] to reduce the variance of the classifiers across clients. We analyze our proposed algorithm and derive a convergence rate bound.

We perform experiments on the popular federated learning benchmark datasets CIFAR10 [19] and CIFAR100 [19] using two types of neural networks, VGG-11 [34] and ResNet-8 [11], and different levels of data heterogeneity across clients. We experimentally show that we require fewer communication rounds compared to the existing methods [14,17,31] to achieve the same accuracy while transmitting a similar or slightly larger number of parameters between server and clients than FedAvg (see Fig. 1). With a (large) fixed number of communication rounds, our method achieves on-par or better top-1 accuracy, and in some settings it even outperforms centralized learning. Using conformal prediction [3], we show how performance can be improved further using adaptive prediction sets.

We show that applying variance reduction on the last layers increases the diversity of the feature extraction layers. This diversity in the feature extraction layers may give each client more freedom to learn richer feature representations, and the uniformity in the classifier then ensures a less biased decision. We summarize our contributions here:

- We present our algorithm for partial variance-reduced federated learning (FedPVR). We experimentally demonstrate that the key to the success of our algorithm is the diversity between the feature extraction layers and the alignment between the classifiers.

- We prove the convergence rate in the convex settings and non-convex settings, precisely characterize its weak dependence on data-heterogeneity measures and show that FedPVR provably converges as fast as the centralized SGD baseline in most practical relevant cases.
- We experimentally show that our algorithm is more communication efficient than previous works across various levels of data heterogeneity, datasets, and neural network architectures. In some cases where data heterogeneity exists, the proposed algorithm even performs slightly better than centralized learning.

## 2. Related work

### 2.1. Federated learning

Federated learning (FL) is a fast-growing field [13,40]. We mainly describe FL methods in non-IID settings where the data is distributed heterogeneously across clients. Among the existing approaches, FedAvg [25] is the *de facto* optimization technique. Despite its solid empirical performances in IID settings [13,25], it tends to achieve a subpar accuracy-communication trade-off in non-IID scenarios.

Many works attempt to tackle FL when data is heterogeneous across clients [1,9,14,31,39,42]. FedProx [31] proposes a temperature parameter and proximal regularization term to control the divergence between client and server models. However, the proximal term does not bring the alignment between the global and local optimal points [1]. Similarly, some works control the update direction by introducing client-dependent control variate [1,14,17,27,32] that is also communicated between the server and clients. They have achieved a much faster convergence rate, but their performance in a non-convex setup, especially in deep neural networks, such as ResNet [11] and VGG [34], is not well explored. Besides, they suffer from a higher communication cost due to the transmission of the extra control variates, which may be a critical issue for resources-limited IoT mobile devices [10]. Among these methods, SCAF-FOLD [14] is the most closely related method to ours, and we give a more detailed comparison in section 3 and 5.

Another line of work develops FL algorithms based on the characteristics, such as expressive feature representations [28] of neural networks. Collins et al. [5] show that FedAvg is powerful in learning common data representations from clients’ data. FedBabu [29], TCT [41], and CCVR [24] propose to improve FL performance by fine-tuning the classifiers with a standalone dataset or features that are simulated based on the client models. However, preparing a standalone dataset/features that represents the data distribution across clients is challenging as this usually requires domain knowledge and may raise privacy concerns.

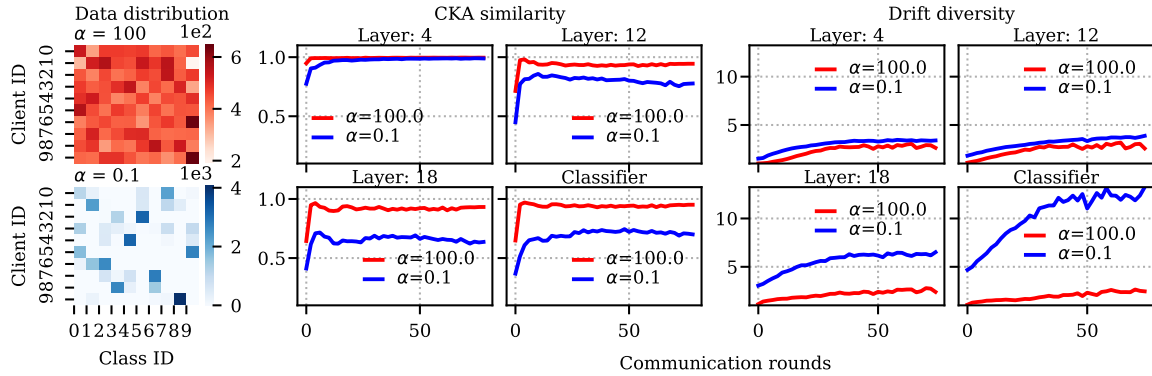


Figure 2. Data distribution (number of images per client per class) with different levels of heterogeneity, client CKA similarity, and the drift diversity of each layer in VGG-11 (20 layers) with FedAvg. **Deep layers in an over-parameterised neural network have higher disagreement and variance when the clients are heterogeneous using FedAvg.**

Moon [20] encourages the similarity of the representations across different client models by using contrastive loss [4] but with the cost of three full-size models in memory on each client, which may limit its applicability in resource-limited devices.

Other works focus on reducing the communication cost by compressing the transmitted gradients [2, 8, 21, 26, 36]. They can reduce the communication bandwidth by adjusting the number of bits sent per iteration. These works are complementary to ours and can be easily integrated into our method to save communication costs.

## 2.2. Variance reduction

Stochastic variance reduction (SVR), such as SVRG [12], SAGA [6], and their variants, use control variate to reduce the variance of traditional stochastic gradient descent (SGD). These methods can remarkably achieve a linear convergence rate for strongly convex optimization problems compared to the sub-linear rate of SGD. Many federated learning algorithms, such as SCAFFOLD [14] and DANE [32], have adapted the idea of variance reduction for the whole model and achieved good convergence on convex problems. However, as [7] demonstrated, naively applying variance reduction techniques gives no actual variance reduction and tends to result in a slower convergence in deep neural networks. This suggests that adapting SVR techniques in deep neural networks for FL requires a more careful design.

## 2.3. Conformal prediction

Conformal prediction is a general framework that computes a prediction set guaranteed to include the true class with a high user-determined probability [3, 30]. It requires no retraining of the models and achieves a finite-sum coverage guarantee [3]. As FL algorithms can hardly perform as well as centralized learning [24] when the data heterogene-

Table 1. Notations used in this paper

$R, r$	Number of communication rounds and round index
$K, k$	Number of local steps, local step index
$N, i$	Number of clients, client index
$\mathbf{y}_{i,k}^r$	client model $i$ at step $k$ and round $r$
$\mathbf{x}^r$	server model at round $r$
$\mathbf{c}_i^r, \mathbf{c}^r$	client and server control variate

ity is high, we can integrate conformal prediction in FL to improve the empirical coverage by slightly increasing the predictive set size. This can be beneficial in sensitive use cases such as detecting chemical hazards, where it is better to give a prediction set that contains the correct class than producing a single but wrong prediction.

## 3. Method

### 3.1. Problem statement

Given  $N$  clients with full participation, we formalise the problem as minimizing the average of the stochastic functions with access to stochastic samples in Eq. 1 where  $\mathbf{x}$  is the model parameters and  $f_i$  represents the loss function at client  $i$  with dataset  $\mathcal{D}_i$ ,

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left( f(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{x}) \right), \quad (1)$$

where  $f_i(\mathbf{x}) := \mathbb{E}_{\mathcal{D}_i} [f_i(\mathbf{x}; \mathcal{D}_i)]$ .

### 3.2. Motivation

When the data  $\{\mathcal{D}_i\}$  are heterogeneous across clients, FedAvg suffers from *client drift* [14], where the average of the local optimal  $\bar{\mathbf{x}}^* = \frac{1}{N} \sum_{i \in N} \mathbf{x}_i^*$  is far from the global optimal  $\mathbf{x}^*$ . To understand what causes client drift, specifically which layers in a neural network are influenced most by the data heterogeneity, we perform a simple experiment using FedAvg and CIFAR10 datasets on a VGG-11. The detailed experimental setup can be found in section 4.

In an over-parameterized model, it is difficult to directly calculate client drift  $\|\bar{\mathbf{x}}^* - \mathbf{x}^*\|^2$  as it is challenging to obtain the global optimum  $\mathbf{x}^*$ . We instead hypothesize that we can represent the influence of data heterogeneity on the model by measuring 1) drift diversity and 2) client model similarity. Drift diversity reflects the diversity in the amount each client model deviates from the server model after an update round.

**Definition 1** (Drift diversity). *We define the drift diversity across  $N$  clients at round  $r$  as:*

$$\xi^r := \frac{\sum_{i=1}^N \|\mathbf{m}_i^r\|^2}{\|\sum_{i=1}^N \mathbf{m}_i^r\|^2} \quad \mathbf{m}_i^r = \mathbf{y}_{i,K}^r - \mathbf{x}^{r-1} \quad (2)$$

Drift diversity  $\xi$  is high when all the clients update their models in different directions, i.e., when dot products between client updates  $\mathbf{m}_i$  are small. When each client performs  $K$  steps of vanilla SGD updates,  $\xi$  depends on the directions and amplitude of the gradients over  $N$  clients and is equivalent to  $\frac{\sum_{i=1}^N \|\sum_k g_i(\mathbf{y}_{i,k})\|^2}{\|\sum_{i=1}^N \sum_k g_i(\mathbf{y}_{i,k})\|^2}$ , where  $g_i(\mathbf{y}_{i,k})$  is the stochastic mini-batch gradient.

After updating client models, we quantify the client model similarity using centred kernel alignment (CKA) [18] computed on a test dataset. CKA is a widely used permutation invariant metric for measuring the similarity between feature representations in neural networks [18, 24, 28].

Fig. 2 shows the movement of  $\xi$  and CKA across different levels of data heterogeneity using FedAvg. We observe that the similarity and diversity of the early layers (e.g. layer index 4 and 12) are with a higher agreement between the IID ( $\alpha = 100.0$ ) and non-IID ( $\alpha = 0.1$ ) experiments, which indicates that FedAvg can still learn and extract good feature representations even when it is trained with non-IID data. The lower similarity on the deeper layers, especially the classifiers, suggests that these layers are strongly biased towards their local data distribution. When we only look at the model that is trained with  $\alpha = 0.1$ , we see the highest diversity and variance on the classifiers across clients compared to the rest of the layers. Based on the above observations, we propose to align the classifiers across clients using variance reduction. We deploy client and server control variates to control the updating directions of the classifiers.

### 3.3. Classifier variance reduction

Our proposed algorithm (Alg. I) consists of three parts: i) client updating (Eq. 5-6) ii) client control variate updating, (Eq. 7), and iii) server updating (Eq. 8-9)

We first define a vector  $\mathbf{p} \in \mathbb{R}^d$  that contains 0 or 1 with  $v$  non-zero elements ( $v \ll d$ ) in Eq. 3. We recover SCAFFOLD with  $\mathbf{p} = \mathbf{1}$  and recover FedAvg with  $\mathbf{p} = \mathbf{0}$ . For the set of indices  $j$  where  $\mathbf{p}_j = 1$  ( $S_{\text{svr}}$  from Eq. 4), we update the corresponding weights  $\mathbf{y}_{i,S_{\text{svr}}}$  with variance reduction such that we maintain a state for each client ( $\mathbf{c}_i \in \mathbb{R}^v$ ) and

for the server ( $\mathbf{c} \in \mathbb{R}^v$ ) in Eq. 5. For the rest of the indices  $S_{\text{sgd}}$  from Eq. 4, we update the corresponding weights  $\mathbf{y}_{i,S_{\text{sgd}}}$  with SGD in Eq. 6. As the server variate  $\mathbf{c}$  is an average of  $\mathbf{c}_i$  across clients, we can safely initialise them as  $\mathbf{0}$ .

In each communication round, each client receives a copy of the server model  $\mathbf{x}$  and the server control variate  $\mathbf{c}$ . They then perform  $K$  model updating steps (see Eq. 5-6 for one step) using cross-entropy as the loss function. Once this is finished, we calculate the updated client control variate  $\mathbf{c}_i$  using Eq. 7. The server then receives the updated  $\mathbf{c}_i$  and  $\mathbf{y}_i$  from all the clients for aggregation (Eq. 8-9). This completes one communication round.

$$\mathbf{p} := \{0, 1\}^d, \quad v = \sum \mathbf{p} \quad (3)$$

$$S_{\text{svr}} := \{j : \mathbf{p}_j = 1\}, \quad S_{\text{sgd}} := \{j : \mathbf{p}_j = 0\} \quad (4)$$

$$\mathbf{y}_{i,S_{\text{svr}}} \leftarrow \mathbf{y}_{i,S_{\text{svr}}} - \eta_l (g_i(\mathbf{y}_i)_{S_{\text{svr}}} - \mathbf{c}_i + \mathbf{c}) \quad (5)$$

$$\mathbf{y}_{i,S_{\text{sgd}}} \leftarrow \mathbf{y}_{i,S_{\text{sgd}}} - \eta_l g_i(\mathbf{y}_i)_{S_{\text{sgd}}} \quad (6)$$

$$\mathbf{c}_i \leftarrow \mathbf{c}_i - \mathbf{c} + \frac{1}{K\eta_l} (\mathbf{x}_{S_{\text{svr}}} - \mathbf{y}_{i,S_{\text{svr}}}) \quad (7)$$

$$\mathbf{x} \leftarrow (1 - \eta_g)\mathbf{x} + \frac{1}{N} \sum_{i \in N} \mathbf{y}_i \quad (8)$$

$$\mathbf{c} \leftarrow \frac{1}{N} \sum_{i \in N} \mathbf{c}_i \quad (9)$$

---

#### Algorithm I Partial variance reduction (FedPVR)

---

**server:** initialise the server model  $\mathbf{x}$ , the control variate  $\mathbf{c}$ , and global step size  $\eta_g$

**client:** initialise control variate  $\mathbf{c}_i$  and local step size  $\eta_l$

**mask:**  $\mathbf{p} := \{0, 1\}^d$ ,  $S_{\text{sgd}} := \{j : \mathbf{p}_j = 0\}$ ,  $S_{\text{svr}} := \{j : \mathbf{p}_j = 1\}$

1: **procedure** MODEL UPDATING

2:   **for**  $r = 1 \rightarrow R$  **do**

3:     **communicate**  $\mathbf{x}$  and  $\mathbf{c}$  to all clients  $i \in [N]$

4:     **for** On client  $i \in [N]$  in parallel **do**

5:        $\mathbf{y}_i \leftarrow \mathbf{x}$

6:       **for**  $k = 1 \rightarrow K$  **do**

7:          compute minibatch gradient  $g_i(\mathbf{y}_i)$

8:           $\mathbf{y}_{i,S_{\text{sgd}}} \leftarrow \mathbf{y}_{i,S_{\text{sgd}}} - \eta_l g_i(\mathbf{y}_i)_{S_{\text{sgd}}}$

9:           $\mathbf{y}_{i,S_{\text{svr}}} \leftarrow \mathbf{y}_{i,S_{\text{svr}}} - \eta_l (g_i(\mathbf{y}_i)_{S_{\text{svr}}} - \mathbf{c}_i + \mathbf{c})$

10:       **end for**

11:        $\mathbf{c}_i \leftarrow \mathbf{c}_i - \mathbf{c} + \frac{1}{K\eta_l} (\mathbf{x}_{S_{\text{svr}}} - \mathbf{y}_{i,S_{\text{svr}}})$

12:       **communicate**  $\mathbf{y}_i, \mathbf{c}_i$

13:     **end for**

14:      $\mathbf{x} \leftarrow (1 - \eta_g)\mathbf{x} + \frac{1}{N} \sum_{i \in N} \mathbf{y}_i$

15:      $\mathbf{c} \leftarrow \frac{1}{N} \sum_{i \in N} \mathbf{c}_i$

16:   **end for**

17: **end procedure**

---

In terms of implementation, we can simply assume the control variate for the block of weights that is updated with SGD as  $\mathbf{0}$  and implement line 8 and 9 in one step

**Ours vs SCAFFOLD** [14] While our work is similar to SCAFFOLD in the use of variance reduction, there are some fundamental differences. We both communicate control variates between the clients and server, but our control

variate ( $2v \leq 0.1d$ ) is significantly smaller than the one in SCAFFOLD (2d). This 2x decrease in bits can be critical for some low-power IoT devices as the communication may consume more energy [10]. From the application point of view, SCAFFOLD achieved great success in convex or simple two layers problems. However, adapting the techniques that work well from convex problems to over-parameterized models is non-trivial [38], and naively adapting variance reduction techniques on deep neural networks gives little or no convergence speedup [7]. Therefore, the significant improvement achieved by our method gives essential and non-trivial insight into what matters when tackling data heterogeneity in FL in over-parameterized models.

### 3.4. Convergence rate

We state the convergence rate in this section. We assume functions  $\{f_i\}$  are  $\beta$ -smooth following [16, 35]. We then assume  $g_i(\mathbf{x}) := \nabla f_i(\mathbf{x}; \mathcal{D}_i)$  is an unbiased stochastic gradient of  $f_i$  with variance bounded by  $\sigma^2$ . We assume strongly convexity ( $\mu > 0$ ) and general convexity ( $\mu = 0$ ) for some of the results following [14]. Furthermore, we also make assumptions about the heterogeneity of the functions.

For convex functions, we assume the heterogeneity of the function  $\{f_i\}$  at the optimal point  $\mathbf{x}^*$  (such a point always exists for a strongly convex function) following [15, 16].

**Assumption 1** ( $\zeta$ -heterogeneity). We define a measure of variance at the optimum  $\mathbf{x}^*$  given  $N$  clients as :

$$\zeta^2 := \frac{1}{N} \sum_{i=1}^N \mathbb{E} \|\nabla f_i(\mathbf{x}^*)\|^2. \quad (10)$$

For the non-convex functions, such an unique optimal point  $\mathbf{x}^*$  does not necessarily exist, so we generalize Assumption 1 to Assumption 2.

**Assumption 2** ( $\hat{\zeta}$ -heterogeneity). We assume there exists constant  $\hat{\zeta}$  such that  $\forall \mathbf{x} \in \mathbb{R}^d$

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E} \|\nabla f_i(\mathbf{x})\|^2 \leq \hat{\zeta}^2. \quad (11)$$

Given the mask  $\mathbf{p}$  as defined in Eq. 3, we know  $\|\mathbf{p} \odot \mathbf{x}\| \leq \|\mathbf{x}\|$ . Therefore, we have the following propositions.

**Proposition 1** (Implication of Assumption 1). Given the mask  $\mathbf{p}$ , we define the heterogeneity of the block of weights that are not variance reduced at the optimum  $\mathbf{x}^*$  as:

$$\zeta_{1-p}^2 := \frac{1}{N} \sum_{i=1}^N \|\mathbf{p} \odot \nabla f_i(\mathbf{x}^*)\|^2, \quad (12)$$

If Assumption 1 holds, then it also holds that:

$$\zeta_{1-p}^2 \leq \zeta^2. \quad (13)$$

In Proposition 1,  $\zeta_{1-p}^2 = \zeta^2$  if  $\mathbf{p} = \mathbf{0}$  and  $\zeta_{1-p}^2 = 0$  if  $\mathbf{p} = \mathbf{1}$ . If  $\mathbf{p} \neq \mathbf{0}$  and  $\mathbf{p} \neq \mathbf{1}$ , as the heterogeneity of the shallow weights is lower than the deeper weights [41], we have  $\zeta_{1-p}^2 \leq \zeta^2$ . Similarly, we can validate Proposition 2.

**Proposition 2** (Implication of Assumption 2). Given the mask  $\mathbf{p}$ , we assume there exists constant  $\hat{\zeta}_{1-p}$  such that  $\forall \mathbf{x} \in \mathbb{R}^d$ , the heterogeneity of the block of weights that are not variance reduced:

$$\frac{1}{N} \sum_{i=1}^N \|\mathbf{p} \odot \nabla f_i(\mathbf{x})\|^2 \leq \hat{\zeta}_{1-p}, \quad (14)$$

If Assumption 2 holds, then it also holds that:

$$\hat{\zeta}_{1-p}^2 \leq \hat{\zeta}^2. \quad (15)$$

**Theorem 1.** For any  $\beta$ -smooth function  $\{f_i\}$ , the output of FedPVR has expected error smaller than  $\epsilon$  for  $\eta_g = \sqrt{N}$  and some values of  $\eta_l$ ,  $R$  satisfying:

- **Strongly convex:**  $\eta_l \leq \min\left(\frac{1}{80K\eta_g\beta}, \frac{26}{20\mu K\eta_g}\right)$ ,

$$R = \tilde{\mathcal{O}}\left(\frac{\sigma^2}{\mu N K \epsilon} + \frac{\zeta_{1-p}^2}{\mu \epsilon} + \frac{\beta}{\mu}\right), \quad (16)$$

- **General convex:**  $\eta_l \leq \frac{1}{80K\eta_g\beta}$ ,

$$R = \mathcal{O}\left(\frac{\sigma^2 D}{K N \epsilon^2} + \frac{\zeta_{1-p}^2 D}{\epsilon^2} + \frac{\beta D}{\epsilon} + F\right), \quad (17)$$

- **Non-convex:**  $\eta_l \leq \frac{1}{26K\eta_g\beta}$ , and  $R \geq 1$ , then:

$$R = \mathcal{O}\left(\frac{\beta \sigma^2 F}{K N \epsilon^2} + \frac{\beta \hat{\zeta}_{1-p}^2 F}{N \epsilon^2} + \frac{\beta F}{\epsilon}\right), \quad (18)$$

Where  $D := \|\mathbf{x}^0 - \mathbf{x}^*\|^2$  and  $F := f(\mathbf{x}^0) - f^*$ .

Given the above assumptions, the convergence rate is given in Theorem 1. When  $\mathbf{p} = \mathbf{1}$ , we recover SCAFFOLD convergence guarantee as  $\zeta_{1-p}^2 = 0$ ,  $\hat{\zeta}_{1-p}^2 = 0$ . In the strongly convex case, the effect of the heterogeneity of the block of weights that are not variance reduced  $\zeta_{1-p}^2$  becomes negligible if  $\tilde{\mathcal{O}}\left(\frac{\zeta_{1-p}^2}{\epsilon}\right)$  is sufficiently smaller than  $\tilde{\mathcal{O}}\left(\frac{\sigma^2}{N K \epsilon}\right)$ . In such case, our rate is  $\frac{\sigma^2}{N K \epsilon} + \frac{1}{\mu}$ , which recovers the SCAFFOLD in the strongly convex without sampling and further matches that of SGD (with mini-batch size  $K$  on each worker). We also recover the FedAvg rate\* at simple IID case. See Appendix. B for the full proof.

\*FedAvg at strongly convex case has the rate  $R = \tilde{\mathcal{O}}\left(\frac{\sigma^2}{\mu K N \epsilon} + \frac{\sqrt{\beta} G}{\mu \sqrt{\epsilon}} + \frac{\beta}{\mu}\right)$  with  $G$  measures the gradient dissimilarity. At simple IID case,  $G=0$  [14].



## 4. Experimental setup

We demonstrate the effectiveness of our approach with CIFAR10 [19] and CIFAR100 [19] on image classification tasks. We simulate the data heterogeneity scenario following [22] by partitioning the data according to the Dirichlet distribution with the concentration parameter  $\alpha$ . The smaller the  $\alpha$  is, the more imbalanced the data are distributed across clients. An example of the data distribution over multiple clients using the CIFAR10 dataset can be seen in Fig. 2. In our experiment, we use  $\alpha \in \{0.1, 0.5, 1.0\}$  as these are commonly used concentration parameters [22]. Each client has its local data, and this data is kept to be the same during all the communication rounds. We hold out the test dataset at the server for evaluating the classification performance of the server model. Following [22], we perform the same data augmentation for all the experiments.

We use two models: VGG-11 and ResNet-8 following [22]. We perform variance reduction for the last three layers in VGG-11 and the last layer in ResNet-8. We use 10 clients with full participation following [41] (close to cross-silo setup) and a batch size of 256. Each client performs 10 local epochs of model updating. We set the server learning rate  $\eta_g = 1$  for all the models [14]. We tune the clients learning rate from  $\{0.05, 0.1, 0.2, 0.3\}$  for each individual experiment. The learning rate schedule is experimentally chosen from constant, cosine decay [23], and multiple step decay [22]. We compare our method with the representative federated learning algorithms FedAvg [25], FedProx [31], SCAFFOLD [14], and FedDyn [1]. All the results are averaged over three repeated experiments with different random initialization. We leave 1% of the training data from each client out as the validation data to tune the hyperparameters (learning rate and schedule) per client. See Appendix C for additional experimental setups. The code is at [github.com/lyn1874/fedpvr](https://github.com/lyn1874/fedpvr).

## 5. Experimental results

We demonstrate the performance of our proposed approach in the FL setup with data heterogeneity in this section. We compare our method with the existing state-of-the-art algorithms on various datasets and deep neural networks. For the baseline approaches, we finetune the hyperparameters and only show the best performance we get. Our main findings are 1) we are more communication efficient than the baseline approaches, 2) conformal prediction is an effective tool to improve FL performance in high data heterogeneity scenarios, and 3) the benefit of the trade-off between diversity and uniformity for using deep neural networks in FL.

### 5.1. Communication efficiency and accuracy

We first report the number of rounds required to achieve a certain level of Top 1% accuracy (66% for CIFAR10 and

44% for CIFAR100) in Table 2. An algorithm is more communication efficient if it requires less number of rounds to achieve the same accuracy and/or if it transmits fewer number of parameters between the clients and server. Compared to the baseline approaches, we require much fewer number of rounds for almost all types of data heterogeneity and models. We can achieve a speedup between 1.5 and 6.7 than FedAvg. We also observe that ResNet-8 tends to converge slower than VGG-11, which may be due to the aggregation of the Batch Normalization layers that are discrepant between the local data distribution [22].

We next compare the top-1 accuracy between centralized learning and federated learning algorithms. For the centralized learning experiment, we tune the learning rate from  $\{0.01, 0.05, 0.1\}$  and report the best test accuracy based on the validation dataset. We train the model for 800 epochs which is as same as the total number of epochs in the federated learning algorithms (80 communication rounds x 10 local epochs). The results are shown in Table 3. We also show the number of copies of the parameters that need to be transmitted between the server and clients (*e.g.* 2x means we communicate  $x$  and  $y_i$ )

Table 3 shows that our approach achieves a much better Top-1 accuracy compared to FedAvg with transmitting a similar or slightly bigger number of parameters between the server and client per round. Our method also achieves slightly better accuracy than centralized learning when the data is less heterogeneous (*e.g.*  $\alpha = 0.5$  for CIFAR10 and  $\alpha = 1.0$  for CIFAR100).

### 5.2. Conformal prediction

When the data heterogeneity is high across clients, it is difficult for a federated learning algorithm to match the centralized learning performance [24]. Therefore, we demonstrate the benefit of using simple post-processing conformal prediction to improve the model performance.

We examine the relationship between the empirical coverage and the average predictive set size for the server model after 80 communication rounds for each federated learning algorithm. The empirical coverage is the percentage of the data samples where the correct prediction is in the predictive set, and the average predictive size is the average of the length of the predictive sets over all the test images [3]. See Appendix C for more information about conformal prediction setup and results.

The results for when  $\alpha = 0.1$  for both datasets and architectures are shown in Fig. 3. We show that by slightly increasing the predictive set size, we can achieve a similar accuracy as the centralized performance. Besides, our approach tends to surpass the centralized top-1 performance similar to or faster than other approaches. In sensitive use cases such as chemical threat detection, conformal prediction is a valuable tool to achieve certified accuracy at the

Table 2. The required number of communication rounds (speedup compared to FedAvg) to achieve a certain level of top-1 accuracy (66% for the CIFAR10 dataset and 44% for the CIFAR100 dataset). Our method requires fewer rounds to achieve the same accuracy.

	CIFAR10 (66%)				CIFAR100 (44%)			
	$\alpha=0.1$		$\alpha=0.5$		$\alpha=0.1$		$\alpha=1.0$	
	VGG-11	ResNet-8	VGG-11	ResNet-8	VGG-11	ResNet-8	VGG-11	ResNet-8
FedAvg	No. rounds	No. rounds	No. rounds	No. rounds	No. rounds	No. rounds	No. rounds	No. rounds
FedProx	55(1.0x)	90(1.0x)	15(1.0x)	15(1.0x)	100+(1.0x)	100+(1.0x)	80(1.0x)	56(1.0x)
SCAFFOLD	52(1.1x)	75(1.2x)	16(0.9x)	20(0.8x)	100+(1.0x)	100+(1.0x)	80(1.0x)	59(0.9x)
FedDyn	39(1.4x)	57(1.6x)	14(1.0x)	9(1.7x)	80(> 1.3x)	<b>61(&gt; 1.6x)</b>	36(2.2x)	25(2.2x)
Ours	<b>27(2.0x)</b>	<b>50(1.8x)</b>	<b>9(1.6x)</b>	<b>5(3.0x)</b>	<b>37(&gt; 2.7x)</b>	66(> 1.5x)	<b>12(6.7x)</b>	<b>15(3.7x)</b>

Table 3. The top-1 accuracy (%) after running 80 communication rounds using different methods on CIFAR10 and CIFAR100, together with the number of communicated parameters between the client and the server. We train the centralised model for 800 epochs (= 80 rounds x 10 local epochs in FL). Higher accuracy is better, and we highlight the best accuracy in red colour.

	VGG-11				server $\leftrightarrow$ client	ResNet-8				
	CIFAR10		CIFAR100			CIFAR10		CIFAR100		
	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.1$	$\alpha = 1.0$		$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.1$	$\alpha = 1.0$	
Centralised	<b>87.5</b>		56.3		-	83.4		<b>56.8</b>		-
FedAvg	69.3	80.9	34.3	45.0	2x	64.9	79.1	38.8	47.0	2x
Fedprox	72.1	80.4	35.0	43.2	2x	66.1	77.9	42.0	47.2	2x
SCAFFOLD	74.1	83.5	43.4	50.6	4x	66.6	80.3	43.8	52.3	4x
FedDyn	77.4	80.1	43.8	45.2	2x	63.8	72.9	36.4	48.1	2x
Ours	78.2	84.9	43.5	<b>58.0</b>	2.1x	69.3	<b>83.6</b>	43.5	52.3	2.02x

cost of a slightly larger predictive set size.




Figure 3. Relation between average predictive size and empirical coverage when  $\alpha = 0.1$ . By slightly increasing the predictive set size, we can achieve a similar performance as the centralised model (Top-1 accuracy) even if the data are heterogeneously distributed across clients. Our method is similar to or faster than other approaches to surpass the centralised Top-1 accuracy.

### 5.3. Diversity and uniformity

We have shown that our algorithm achieves a better speedup and performance against the existing approaches with only lightweight modifications to FedAvg. We next investigate what factors lead to better accuracy. Specifi-




Figure 4. Drift diversity and learning curve for ResNet-8 on CIFAR100 with  $\alpha = 1.0$ . Compared to FedAvg, SCAFFOLD and our method can both improve the agreement between the classifiers. Compared to SCAFFOLD, our method results in a higher gradient diversity at the early stage of the communication, which tends to boost the learning speed as the curvature of the drift diversity seem to match the learning curve.

cally, we calculate the drift diversity  $\xi$  across clients after each communication round using Eq. 2 and average  $\xi$  across three runs. We show the result of using ResNet-8 and CIFAR100 with  $\alpha = 1.0$  in Fig. 4.

Fig. 4 shows the drift diversity for different layers in ResNet-8 and the testing accuracy along the communication rounds. We observe that classifiers have the highest diver-

sity in FedAvg against other layers and methods. SCAFFOLD, which applies control variate on the entire model, can effectively reduce the disagreement of the directions and scales of the averaged gradient across clients. Our proposed algorithm that performs variance reduction only on the classifiers can reduce the diversity of the classifiers even further but increase the diversity of the feature extraction layers. This high diversity tends to boost the learning speed as the curvature of the diversity movement (Fig. 4 left) seems to match the learning curve (Fig. 4 right). Based on this observation, we hypothesize that this diversity along the feature extractor and the uniformity of the classifier is the main reason for our better speedup.

To test this hypothesis, we perform an experiment where we use variance reduction starting from different layers of a neural network. If the starting position of the use of variance reduction influences the learning speed, it indicates where in a neural network we need more diversity and where we need more uniformity. We here show the result of using VGG-11 on CIFAR100 with  $\alpha = 1.0$  as there are more layers in VGG-11. The result is shown in Fig. 5 where SVR:16  $\rightarrow$  20 is corresponding to our approach and SVR:0  $\rightarrow$  20 is corresponding to SCAFFOLD that applies variance reduction for the entire model. Results for using ResNet-8 is shown in Appendix. C.




Figure 5. Influence of using stochastic variance reduction(SVR) on layers that start from different positions in a neural network on the learning speed. SVR:0 $\rightarrow$ 20 applies variance reduction on the entire model (SCAFFOLD). SVR:16 $\rightarrow$ 20 applies variance reduction from the layer index 16 to 20 (ours). The later we apply variance reduction, the better performance speedup we obtain. However, no variance reduction (FedAvg) performs the worst here.

We see from Fig. 5 that the deeper in a neural network we apply variance reduction, the better learning speedup we can obtain. There is no clear performance difference between where to activate the variance reduction when the layer index is over 10. However, applying no variance reduction (FedAvg) achieves by far the worst performance. We believe that these experimental results indicate that in a distributed optimization framework, to boost the learning

speed of an over-parameterized model, we need some levels of diversity in the middle and early layers for learning richer feature representation and some degrees of uniformity in the classifiers for making a less biased decision.

## 6. Conclusion

In this work, we studied stochastic gradient descent learning for deep neural network classifiers in a federated learning setting, where each client updates its local model using stochastic gradient descent on local data. A central model is periodically updated (by averaging local model parameters) and broadcast to the clients under a communication bandwidth constraint. When data is homogeneous across clients, this procedure is comparable to centralized learning in terms of efficiency; however, when data is heterogeneous, learning is impeded. Our hypothesis for the primary reason for this is that when the local models are out of alignment, updating the central model by averaging is ineffective and sometimes even destructive.

Examining the diversity across clients of their local model updates and their learned feature representations, we found that the misalignment between models is much stronger in the last few neural network layers than in the rest of the network. This finding inspired us to experiment with aligning the local models using a partial variance reduction technique applied only on the last layers, which we named FedPVR. We found that this led to a substantial improvement in convergence speed compared to the competing federated learning methods. In some cases, our method even outperformed centralized learning. We derived a bound on the convergence rate of our proposed method, which matches the rates for SGD when the gradient diversity across clients is sufficiently low. Compared with FedAvg, the communication cost of our method is only marginally worse, as it requires transmitting control variates for the last layers.

We believe our FedPVR algorithm strikes a good balance between simplicity and efficiency, requiring only a minor modification to the established FedAvg method; however, in our further research, we plan to pursue more optimal methods for aligning and guiding the local learning algorithms, *e.g.* using adaptive procedures. Furthermore, the degree of over-parameterization in the neural network layers (*e.g.* feature extraction vs bottlenecks) may also play an important role, which we would like to understand better.

## Acknowledgements

The first three authors thank for financial support from the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 883390 (H2020-SU-SECU-2019 SERSing Project). BL thanks for the financial support from the Otto Mønsted Foundation.



## References

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N. Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. *CoRR*, abs/2111.04263, 2021. [1](#), [2](#), [6](#), [26](#)
- [2] Dan Alistarh, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD: randomized quantization for communication-optimal stochastic gradient descent. *CoRR*, abs/1610.02132, 2016. [3](#)
- [3] Anastasios Nikolas Angelopoulos, Stephen Bates, Michael Jordan, and Jitendra Malik. Uncertainty sets for image classifiers using conformal prediction. In *International Conference on Learning Representations*, 2021. [2](#), [3](#), [6](#)
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *CoRR*, abs/2002.05709, 2020. [3](#)
- [5] Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Fedavg with fine tuning: Local updates lead to representation learning, 2022. [2](#)
- [6] Aaron Defazio, Francis R. Bach, and Simon Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. *CoRR*, abs/1407.0202, 2014. [3](#)
- [7] Aaron Defazio and Léon Bottou. On the ineffectiveness of variance reduced optimization for deep learning. *CoRR*, abs/1812.04529, 2018. [2](#), [3](#), [5](#)
- [8] Aritra Dutta, El Houcine Bergou, Ahmed M. Abdelmoniem, Chen-Yu Ho, Atal Narayan Sahu, Marco Canini, and Panos Kalnis. On the discrepancy between the theoretical analysis and practical implementations of compressed communication for distributed deep learning. *CoRR*, abs/1911.08250, 2019. [3](#)
- [9] Liang Gao, Huazhu Fu, Li Li, Yingwen Chen, Ming Xu, and Cheng-Zhong Xu. Feddc: Federated learning with non-iid data via local drift decoupling and correction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10102–10111. IEEE, 2022. [2](#)
- [10] Malka N. Halgamuge, Moshe Zukerman, Kotagiri Ramamohanarao, and Hai Le Vu. An estimation of sensor energy consumption. *Progress in Electromagnetics Research B*, 12:259–295, 2009. [1](#), [2](#), [5](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. [2](#)
- [12] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, 2013. [3](#)
- [13] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista A. Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *CoRR*, abs/1912.04977, 2019. [1](#), [2](#)
- [14] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. SCAFFOLD: Stochastic controlled averaging for federated learning. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5132–5143. PMLR, 13–18 Jul 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [11](#), [12](#), [13](#), [21](#)
- [15] Ahmed Khaled, Konstantin Mishchenko, and Peter Richtárik. Better communication complexity for local SGD. *CoRR*, abs/1909.04746, 2019. [5](#)
- [16] Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified theory of decentralized SGD with changing topology and local updates. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5381–5393. PMLR, 13–18 Jul 2020. [5](#), [11](#), [13](#)
- [17] Jakub Konečný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *CoRR*, abs/1610.02527, 2016. [1](#), [2](#)
- [18] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey E. Hinton. Similarity of neural network representations revisited. *CoRR*, abs/1905.00414, 2019. [4](#)
- [19] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. [2](#), [6](#)
- [20] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. *CoRR*, abs/2103.16257, 2021. [1](#), [3](#)
- [21] Zhize Li, Dmitry Kovalev, Xun Qian, and Peter Richtárik. Acceleration for compressed gradient descent in distributed and federated optimization, 2020. [3](#)
- [22] Tao Lin, Lingjing Kong, Sebastian U. Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [6](#), [25](#), [26](#)
- [23] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with restarts. *CoRR*, abs/1608.03983, 2016. [6](#)
- [24] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *CoRR*, abs/2106.05001, 2021. [1](#), [2](#), [3](#), [4](#), [6](#)

- [25] H. Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. Federated learning of deep networks using model averaging. *CoRR*, abs/1602.05629, 2016. [2](#), [6](#)
- [26] Konstantin Mishchenko, Eduard Gorbunov, Martin Takáč, and Peter Richtárik. Distributed learning with compressed gradient differences. *CoRR*, abs/1901.09269, 2019. [3](#)
- [27] Konstantin Mishchenko, Grigory Malinovsky, Sebastian Stich, and Peter Richtárik. ProxSkip: Yes! local gradient steps provably lead to communication acceleration! finally! *International Conference on Machine Learning (ICML)*, 2022. [2](#)
- [28] Thao Nguyen, Maithra Raghu, and Simon Kornblith. Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth. *CoRR*, abs/2010.15327, 2020. [2](#), [4](#)
- [29] Jaehoon Oh, Sangmook Kim, and Se-Young Yun. Fedbabu: Towards enhanced representation for federated image classification. *CoRR*, abs/2106.06042, 2021. [2](#)
- [30] Yaniv Romano, Matteo Sesia, and Emmanuel J. Candès. Classification with valid and adaptive coverage. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA, 2020. Curran Associates Inc. [3](#)
- [31] Anit Kumar Sahu, Tian Li, Maziar Sanjabi, Manzil Zaheer, Ameet Talwalkar, and Virginia Smith. On the convergence of federated optimization in heterogeneous networks. *CoRR*, abs/1812.06127, 2018. [1](#), [2](#), [6](#), [26](#)
- [32] Ohad Shamir, Nathan Srebro, and Tong Zhang. Communication efficient distributed optimization using an approximate newton-type method. *CoRR*, abs/1312.7853, 2013. [1](#), [2](#), [3](#)
- [33] Micah J. Sheller, Brandon Edwards, G. Anthony Reina, Jason Martin, Sarthak Pati, Aikaterini Kotrotsou, Mikhail Milchenko, Weilin Xu, Daniel Marcus, Rivka R. Colen, and Spyridon Bakas. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1):12598, Jul 2020. [1](#)
- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. [2](#)
- [35] Sebastian U. Stich. Unified optimal analysis of the (stochastic) gradient method. *CoRR*, abs/1907.04232, 2019. [5](#), [11](#), [14](#)
- [36] Sebastian U. Stich, Jean-Baptiste Cordonnier, and Martin Jaggi. Sparsified SGD with memory. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 4452–4463. Curran Associates, Inc., 2018. [3](#)
- [37] Sebastian U. Stich and Sai Praneeth Karimireddy. The error-feedback framework: Better rates for SGD with delayed gradients and compressed communication. *CoRR*, abs/1909.05350, 2019. [15](#)
- [38] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. [2](#), [5](#)
- [39] F Varno, M Saghayi, L Rafiee, S Gupta, S Matwin, and M Havaei. Minimizing client drift in federated learning via adaptive bias estimation. *ArXiv*, abs/2204.13170, 2022. [2](#)
- [40] Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H. Brendan McMahan, Blaise Agüera y Arcas, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, Deepesh Data, Suhas N. Diggavi, Hubert Eichner, Advait Gadhikar, Zachary Garrett, Antonious M. Girgis, Filip Hanzely, Andrew Hard, Chaoyang He, Samuel Horváth, Zhouyuan Huo, Alex Ingerman, Martin Jaggi, Tara Javidi, Peter Kairouz, Satyen Kale, Sai Praneeth Karimireddy, Jakub Konečný, Sanmi Koyejo, Tian Li, Luyang Liu, Mehryar Mohri, Hang Qi, Sashank J. Reddi, Peter Richtárik, Karan Singhal, Virginia Smith, Mahdi Soltanolkotabi, Weikang Song, Ananda Theertha Suresh, Sebastian U. Stich, Ameet Talwalkar, Hongyi Wang, Blake E. Woodworth, Shanshan Wu, Felix X. Yu, Honglin Yuan, Manzil Zaheer, Mi Zhang, Tong Zhang, Chunxiang Zheng, Chen Zhu, and Wennan Zhu. A field guide to federated optimization. *CoRR*, abs/2107.06917, 2021. [2](#)
- [41] Yaodong Yu, Alexander Wei, Sai Praneeth Karimireddy, Yi Ma, and Michael I. Jordan. TCT: Convexifying federated learning using bootstrapped neural tangent kernels, 2022. [2](#), [5](#), [6](#)
- [42] Haoyu Zhao, Zhize Li, and Peter Richtárik. Fedpage: A fast local stochastic gradient method for communication-efficient federated learning. *CoRR*, abs/2108.04755, 2021. [2](#)