

Federated Learning Over Images: Vertical Decompositions and Pre-Trained Backbones Are Difficult to Beat

Erdong Hu¹, Yuxin Tang¹, Anastasios Kyrillidis, Chris Jermaine
Rice University

{eh51, yuxin.tang, anastasios, cmj4}@rice.edu

Abstract

We carefully evaluate a number of algorithms for learning in a federated environment, and test their utility for a variety of image classification tasks. We consider many issues that have not been adequately considered before: whether learning over data sets that do not have diverse sets of images affects the results; whether to use a pre-trained feature extraction “backbone”; how to evaluate learner performance (we argue that classification accuracy is not enough), among others. Overall, across a wide variety of settings, we find that vertically decomposing a neural network seems to give the best results, and outperforms more standard reconciliation-used methods.

1. Introduction

There has been a recent influx of work aimed at developing and evaluating new algorithms for Federated Learning (FL) [2, 6, 20, 29], particularly for image classification [3, 14, 18, 27, 28, 31, 35, 50, 53]. Most papers are primarily concerned with the development of innovative algorithms, and less concerned with the design of appropriate benchmarks and especially, appropriate baselines to test against. This paper, in contrast, is concerned with benchmarking: *Which existing methods for FL work best, and under what conditions do they work, or not work?* As such, our primary contribution is a set of carefully-designed experiments, rather than the introduction of a new FL algorithm. Our benchmarks are designed to address the following concerns regarding the design of FL benchmarks and baselines.

(1) *Learning algorithms are typically evaluated on a few diverse data sets, rather than a large variety of more focused data sets.* The most common benchmark for evaluating federated image classification is CIFAR-100 [25]. This dataset includes whales, chairs, dinosaurs, and so on. This is a very

diverse set of classes, and we are concerned that few FL tasks will involve differentiating dinosaurs from rabbits.

For example, imagine that the members of a bird watching club taking pictures with their smartphones; club members label the pictures with the bird species, and FL is used to build a classifier. Or, a set of companies who are typically competitors—and hence cannot exchange data—want to work together to classify pictures of industrial drill bits based on whether they are going to fail. Both of these deployments involve narrow domains, and the classification problems involve fine-grained differentiation [13, 48] among members of a narrow category. True, not all applications of FL will be narrow, but they are likely to be far narrower than classifying whales versus chairs. We argue that, while evaluating on a broad data set—such as CIFAR-100—is useful, it is also necessary to evaluate FL algorithms on a variety of more narrow data sets.

(2) *Most evaluations focus on final accuracy or the number of communication rounds required for convergence.* FL algorithms are often evaluated by reporting final accuracies (after convergence), or by plotting test accuracy as a function of the number of epochs, or the number of communication rounds. However, final accuracy, at least in isolation, is not really a useful metric. After all, the simplest “FL” algorithm is classical, data parallel learning [33, 41, 52]. That is, run distributed gradient descent; compute gradients locally over mini-batches in a fully synchronous way, then do an all-reduce. As this is functionally equivalent to centralized learning, it is invariant to data distribution, and is likely to be the most accurate method in terms of final accuracy. However, it is inefficient in a federated environment.

Likewise, considering accuracy as a function of communication rounds ignores the computation or communication cost of each round. Communication and computation costs can vary across methods, and are typically far more important than the number of rounds. A method that can quickly achieve high accuracy with a relatively large number of very inexpensive communication rounds—where a small fraction of the model is communicated at each round—is proba-

¹Equal contribution.

bly preferred to one that uses few rounds, but must transmit a huge amount of data. Similarly, for computation, given similar accuracies, an algorithm that performs one forward and one backward gradient descent pass is preferred over one that performs two additional forward passes, as extra overhead [28], even though both operate in the same number of communication rounds.

(3) *Pre-trained feature extractors need to be a standard baseline.* Training the feature-extraction backbone that is used in most deep image processing networks—that is, the series of convolutional and pooling layers and without the final classification layer—is not easy. It is difficult to get the training process right, and training a good backbone requires a lot of computation.

Researchers have recently suggested using pre-trained backbones, especially for few-shot training [7, 16]. In a pre-trained backbone, a pre-trained feature extractor such as a convolutional neural network (CNN) [15, 19, 39] or vision transformer [10] is used without modification to embed an image in a high-dimensional space. Learning is then performed on the resulting vectors, and not on the original images. The goal is to leverage an existing, well-trained model on the new learning problem. Given the resource constraints in many FL scenarios, it is unclear why attempting to fully train such a backbone from scratch in a federated environment would be the first approach. Instead, we suggest simply taking a standard, pre-trained backbone (such as a ResNet [15], or a DenseNet [19], or both used together) and directly using that backbone as a feature extractor, without training.

The benefit is that each training image only needs to be pushed through the backbone one time to obtain a compact set of features, and then that set can be used during training. As the backbone need not to be communicated or repeatedly used to process the input images, both CPU/GPU cycles and communication are saved during federated training. Furthermore, final accuracies might actually be better, as this approach sidesteps the difficulty of training the backbone.

(4) *Off-the-shelf performance is really what matters, but it is often neglected in research studies.* In a centralized environment, it is feasible to train and re-train many times, as various parameters (learning rate, proximal weight, exact neural architecture, etc.) are tuned. In FL, this is much less feasible. Most FL scenarios imagine a resource-constrained environment (possibly involving edge devices), where one cannot ignore the cost of tuning. FL is typically happening on someone else’s hardware, and so running an algorithm many times during parameter tuning is likely to make participation in training far less palatable.

Further, application-specific tuning may simply be impossible. A device may enter the federated training and then

drop out shortly after, never to be seen again. One cannot access a missing device’s data to perform validation, and clearly, one cannot retrain over data that cannot be accessed.

As such, we argue that “off-the-shelf” performance is more important than in centralized learning. That is, it becomes important to settle on one set of universal parameters that tend to work in most deployments, such that the learning algorithm can be universally deployed without further tuning. At the very least, this is an important use case that should be covered in most experiments.

Our Contributions. Our contributions are as follows.

- We suggest several rules-of-thumb governing benchmark and baseline development for FL over images, and use those rules to devise an extensive FL image classification benchmark, and use that benchmark.
- We show that using pre-trained features and other model reduction tools are necessary for practical FL.
- We highlight some surprising behaviors of current FL algorithms, including the effect of client population size, that to our knowledge are not sufficiently explored in current literature.
- Our source code is publicly available at <https://github.com/huerdong/FedVert-Experiments>

2. Background

The common setting for FL is the following: There are total N devices participating in the optimization. To protect data privacy (or due to communication constraints), devices are not allowed to share local data samples to others. Each device has its own local objective function f_i parameterized by its own model w_i . The coefficient $p_i = n_i/n$ is the number of data samples n_i on each device, divided by the unified total data sample $n = \sum_{i=1}^N n_i$. There will be a central server that collects all the parameters w_i uploaded by each device and aggregate them into a single model w .

The common formulation of FL is to optimize a global optimization function $L(w)$:

$$\min_w L(w) = \sum_{i=1}^N p_i L(w_i; \{x_{\text{loc}}, y_{\text{loc}}\}), \quad (1)$$

where $L(w_i; \{x_{\text{loc}}, y_{\text{loc}}\})$ represents the local objective for device i , based on local data $\{x_{\text{loc}}, y_{\text{loc}}\}$. The natural algorithm that arises from this objective is FedAvg [37] which directly applies stochastic gradient descent (SGD) to learn w_i ’s. In the most basic case, where FedAvg performs only one local gradient step, the above formulation becomes equivalent to data parallel SGD [33, 41, 52].

In real-world applications of FL, local data distributions are often skewed towards specific classes: i.e., the natural

distribution of real data can be very skewed, with certain classes only existing in a very small segment of clients' devices. In these heavily non-i.i.d. scenarios, the learned w_i in Eq. (1) can diverge in drastically different directions and, in the most extreme cases, diverge in opposite directions. When FL is performed in this scenario, the averaging can result in a model drastically different from any of the client models and has degraded performance as a consequence; a phenomenon often called client drift [54]. It is this concern that has motivated the development of a number of FL methods beyond FedAvg [21, 28, 30, 43], a few of which will be described in the next section.

3. Federated Learning Methods Tested

3.1. The Methods

In our experiments, we consider the following FL methods, chosen as a representative set of the state-of-the-art.

FedAvg. Federated averaging [37] is the most simple averaging algorithm and is the basis of the others, whereby client models are trained locally by (stochastic) gradient descent. These local models are aggregated by directly averaging their weights, or equivalently averaging their computed gradients using this averaged gradient for one global round of gradient descent.

FedProx. FedProx [30] is a modification on FedAvg that introduces an additional ℓ_2 -norm term to the local loss function, that takes the difference between the previously communicated global model and the current model. The aim is to avoid client drift by biasing the training towards an established model. Again, the clients then compute local models whose weights are averaged. Given the model weights w_0 from the previous communication round, the FedProx modification to the FedAvg objective is:

$$\min_{w_i} \left\{ L(w_i; \{x_{\text{loc}}, y_{\text{loc}}\}) + \frac{\mu}{2} \|w_i - w_0\|_2^2 \right\}, \quad i \in [N]$$

where μ is a regularization parameter.

MOON. MOON (Model-Contrastive) [28] uses a similar principle as FedProx by comparing the current model with the previous model. However, it utilizes a cross-entropy loss, rather than ℓ_2 -norm loss, as in:

$$\min_{w_i} \left\{ L(w_i; \{x_{\text{loc}}, y_{\text{loc}}\}) - \mu \log \frac{e^{\text{sim}(z, z_{\text{glob}})/\tau}}{e^{\text{sim}(z, z_{\text{glob}})/\tau} + e^{\text{sim}(z, z_{\text{prev}})/\tau}} \right\}$$

where $\text{sim}()$ is the cosine similarity function, and $z, z_{\text{glob}}, z_{\text{prev}}$ are representations of the data generated by the respective model (respectively: learned, current global, previous local). τ is a temperature hyperparameter.

FedAdam. While the above algorithms modify the local training process of FedAvg, FedAdam [40] modifies the aggregation step. Once the averaged weights are computed

via local gradient descent and their average is computed, it is used to compute a global gradient surrogate to be used in the Adam algorithm [22].

FedNova. FedNova [43] introduces a client gradient re-weighting scheme when accumulating the models parameters/gradients and generalizes FedAvg and FedProx as specific parameterizations. However, we use the specific formulation specified in [43], where these weights are normalizations of the local gradients.

IST. We consider the previous set of learning algorithms to be “reconciliation-used”; in that, they attempt to train a copy of the model at each site. In contrast, IST, or *independent subnet training* [51] does not utilize averaging or some other method to try to reconcile versions of the model trained at the various site. Instead, IST decomposes the model and sends non-overlapping parts of the model to different sites. In one training round, all neurons (or activations) in a neural network are randomly partitioned to the active sites. A weight is only sent to a site if it connects two neurons that have been assigned to the site. Thus, each site is assigned only a very small subnetwork, which is trained locally for a number of gradient descent steps. At the end of local training, the updated weights are shuffled (with or without the aid of a central server) and the process is repeated. Since subnetworks are independent, there is no averaging or other form of reconciliation needed. There are many variants of this approach in the literature, including FjORD [17], HeteroFL [9], LotteryFL [26], FedSelect [4], FedRolex [1], Federated Dropout [5], PVT [47], and the independent subnetwork training approach in the non-FL setting [11, 32, 44, 45, 51]. It is this latter version of the approach (for MLPs) that we apply here.

ISTProx. This is IST, but with a proximal term added to the loss function.

3.2. Analytic Comparison

The focus of our experiments will be on comparing each of the methods based on their ability to produce high test accuracy with a low cost. In the paper, we will define cost in terms of floating point operations (FLOPs) required to reach a given accuracy, and in terms of floating point numbers transferred. As such, we begin by considering the number of FLOPS required and the floating point numbers that must be transferred to complete one gradient descent step for each of the methods.

In our analysis, we assume that we are using a pre-trained backbone to transform each input image into a vector, and we are learning a MLP with two weight matrices. The MLP accepts an n_1 -dimensional input. It has a hidden layer of n_2 neurons, and has n_3 outputs. Assume a batch size of b at each of s sites. Since each image only goes through the pre-trained backbone once, the FLOP cost

of the pre-trained backbone is negligible when amortized over many gradient descent steps. We will consider only the FLOPs associated with the weights, as the cost associated with updating weights dominates the cost associated with the activation layers.

Begin with FedAvg. In this case, to push a batch through, we have approximately $2b(n_1n_2 + n_2n_3)$ FLOPs for the forward pass and approximately $2b(n_1n_2 + 2n_2n_3)$ for the backward pass, for a total of $4n_1n_2b + 6n_2n_3b$ FLOPs at each site. Coordination requires $2s(n_1n_2 + n_2n_3)$ floating point numbers be sent (the model needs to be distributed to all sites, and then the updated model needs to be collected from all sites). The number of FLOPs incurred during averaging is typically not significant, as we are simply average s copies of the two weight matrices, having n_1n_2 and n_2n_3 floating point numbers each, and this may be done using a distributed algorithm.

Adding a proximal term in the case of FedProx may have little effect, as this changes the backward pass at each site by adding FLOP computations that scales linearly with the weight matrices, but does not scale with the batch size b . Thus, for any reasonable batch size, it is insignificant. If multiple gradient update steps are performed at each site, it also requires storing an older copy of the weight matrices at each site.

FedAdam is a federated version of the Adam algorithm and as such, it requires storing four additional matrices—typically at a coordinator—two of which are associated with each of the MLP’s weight matrices. It must approximate the gradient and run the Adam algorithm at the coordinator during each update. However, this does not affect the memory, computation, or communication requirements at each site, and so any differences may not be of practical importance.

FedNova requires a normalization step at each site during the local gradient computation, but this only requires multiplying each gradient by a scalar, which is likely insignificant and so it is also, for practical purposes, computationally no more expensive than FedAvg. MOON is more expensive than FedAvg in terms of FLOPs, as it needs to compute multiple representations of each training iteration. The obvious MOON implementation will require $4bn_1n_2$ additional FLOPs compared to FedAvg at each site. In addition, MOON needs to locally maintain two additional copies of the weight matrices.

Next, consider IST. At each site, n_2 is effectively cut by a fraction s , because the neurons in the hidden later is partitioned s ways and only the weights associated with neurons assigned to a site are actually sent or used there. Hence, the total number of FLOPs is $(4n_1n_2b + 6n_2n_3b)/s$, and coordination requires that only $2(n_1n_2 + n_2n_3)$ floating point numbers be sent and received; each weight is assigned to exactly one site. There are no FLOPs associated with coordination, as there is no averaging done; weights are sim-

ply updated and sent around. Due to the partitioning, the memory usage at each site is on average smaller than that of FedAvg by a factor of s . Like FedAvg, adding a proximal term has little effect on either FLOPs or bytes transferred.

4. Experimental Setup

4.1. Pre-Trained Backbones

All of the methods we tested (except for one baseline) utilize pre-trained backbones. That is, we take a pre-trained feature image classifiers: ResNet101 [15] and DenseNet121 [19], and remove the classification layers to produce two “backbones.” When processing an image, we simply concatenate the feature vectors produced by the two backbones to produce a 3072 dimensional vector. We also evaluate FedProx, but without a pre-trained backbone. This version of FedProx (called “FedFull”) uses a ResNet18 instead of a single-hidden layer MLP.

4.2. Data Sets and Hyperparameters

Overall, we conduct experiments on a diverse and broad range of image classification data sets: (1) CUB-200-2011 (Birds) [42], (2) Stanford Cars (Cars) [23], (3) VGGFlowers (Flwrs) [38], (4) Aircraft (Aircraft) [36], (5) Describable Textures (Textre) [8], and (6) CIFAR 100 (CFAR) [24]. The only processing we apply on these data sets, besides using the pre-trained feature extractors, is the standard cropping and resizing of image data. For the Aircraft data set, there are multiple levels of specificity of labels - manufacturer, family, variant - we chose variant for the finest-grained classes.

We use the Birds data set to tune the methods. After extensive experimentation on Birds, we chose a learning rate of 0.01, as we found this to give stable convergence in all algorithms, except FedAdam, where we choose a learning rate of 0.03. For all algorithms except for IST, ISTProx, and FedFull, tuning led us to a single-hidden-layer, multi-layer perceptron (MLP) with 1,000 hidden neurons sitting on top of the feature extractor. FedFull was trained with its “out of the box” configuration. For IST and ISTProx, we find scaling the hidden neurons based on the number of concurrent training sites made sense (3000 neurons for ten sites, 6000 for 20, 9000 for 30, 18000 for 60, and so on). This ensures that the client models are not too small.

A number of other hyperparameters were also optimized using the Birds data set. For the various hyperparameters in each algorithm, we applied a grid search on a set list of parameter values, as described in the next subsection. We consider the number of local iterations of stochastic gradient descent as a hyperparameter with either 1, 5, or 25 local training batches, before a communication round. We use a batch size of 32 images. For FedProx, FedNova, and IST-Prox, we tuned the additional proximal hyperparameter μ . For MOON, we have a similar weight parameter μ (the co-

(a) Final accuracies						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	73.5	66.7	95.9	48.8	74.3	76.7
FedProx	74.0	65.7	95.8	48.9	74.7	75.1
MOON	74.3	66.0	95.6	48.9	74.6	75.3
FedAdam	67.6	51.1	92.4	41.9	72.8	70.1
FedNova	73.3	66.8	95.6	49.0	74.7	74.7
FedFull	5.2	2.8	35.1	5.4	40.3	7.8
IST	74.6	67.8	96.4	50.6	74.6	76.2
ISTProx	74.8	66.8	96.5	50.4	74.7	76.1

(b) Communication (GB) to threshold acc.						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	9	6	4	13	3	10
FedProx	13	FAIL	6	21	5	15
MOON	13	43	6	21	4	16
FedAdam	8	FAIL	2	FAIL	1	7
FedNova	57	138	17	84	12	68
FedFull	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
IST	3	288	1	3	1	3
ISTProx	3	FAIL	2	3	1	3

(c) GFLOPs to threshold acc.						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	904	642	445	1334	275	979
FedProx	53	FAIL	23	86	18	62
MOON	110	FAIL	46	178	34	121
FedAdam	FAIL	FAIL	170	FAIL	152	750
FedNova	1140	2768	333	1680	232	1355
FedFull	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
IST	11	1152	6	14	3	12
ISTProx	11	FAIL	6	14	3	13

Table 1: i.i.d. results, 100 sites, 10% participation.

efficient of the contrastive term) and τ , the temperature of the contrastive term. FedAdam has a central learning rate as well, which we set to 0.01. Exact values for all of our hyperparameters are given in the appendix. We stress that, while we performed careful tuning on the Birds data set, all hyperparameters were re-used blindly and without modification, on all other data sets.

4.3. Evaluation Metrics

Choosing an evaluation metric is among the most difficult aspects of evaluating FL methods, but does not seem to have received much attention in the literature. Most papers focus on final accuracy: *What is the accuracy that a FL method exhibits on a test set after training?* Unfortunately, final accuracy as a metric makes little sense in isolation from other considerations. After all, the method that likely gives the very best final accuracy is simple dis-

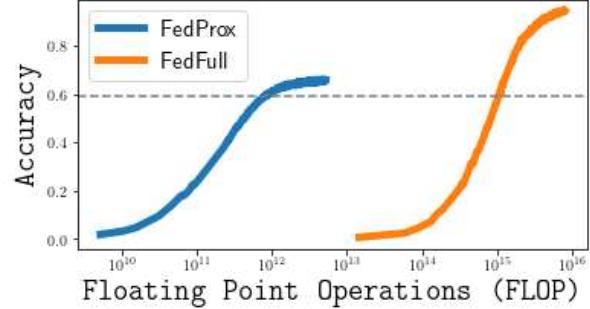


Figure 1: Accuracy as a function of time for FedProx and FedFull on the Stanford Cars data set.

tributed mini-batch training, where each site computes gradients over a small subset of its local data; those gradients are summed and a central site updates the model. This is algorithmically equivalent to centralized training, but it is typically not tested, because it is generally assumed that the convergence would be too “slow,” for some implicit definition of “slow.” Thus, a more carefully-defined evaluation framework (beyond simple accuracy) is needed.

Hence, we consider three metrics: *final accuracy*, *computational efficiency*, and *communication efficiency* [34,46]. Of the three, we assert that final accuracy should be considered the *least important*; it is most useful as a way to contextualize the other two metrics (e.g., “*what final accuracy can I achieve in G gigaFLOPs of computation?*”).

4.4. Tuning and Evaluation

For a given method, using the Birds data set, we first perform a grid search over the various hyperparameters, and for each hyperparameter set, we run the method for a “long time,” where a “long time” allows a maximum of 10^{14} floating point operations (or 100 teraFLOPs, where floating point operations are summed across all devices participating in federated training) and five terabytes of data communicated (assuming single-precision floating point computations). Note that 100 teraFLOPs corresponds to around 14 hours of CPU time on a single mobile phone, or perhaps several hundred hours of processor time on a limited IoT device. Once a method exceeds 100 teraFLOPs or five terabytes of communication, we take its very best observed test accuracy, averaged over ten communication rounds. The best average observed over all hyperparameter settings is taken as the method’s “final accuracy”.

Once we have the method’s final accuracy, we then perform a grid search over all hyperparameters to find the two sets of parameter settings for the method that (a) minimize the number of FLOPs to reach 90% of the final accuracy, and (b) minimize the number of GB of communication necessary to reach 90% of the final accuracy. Now we have

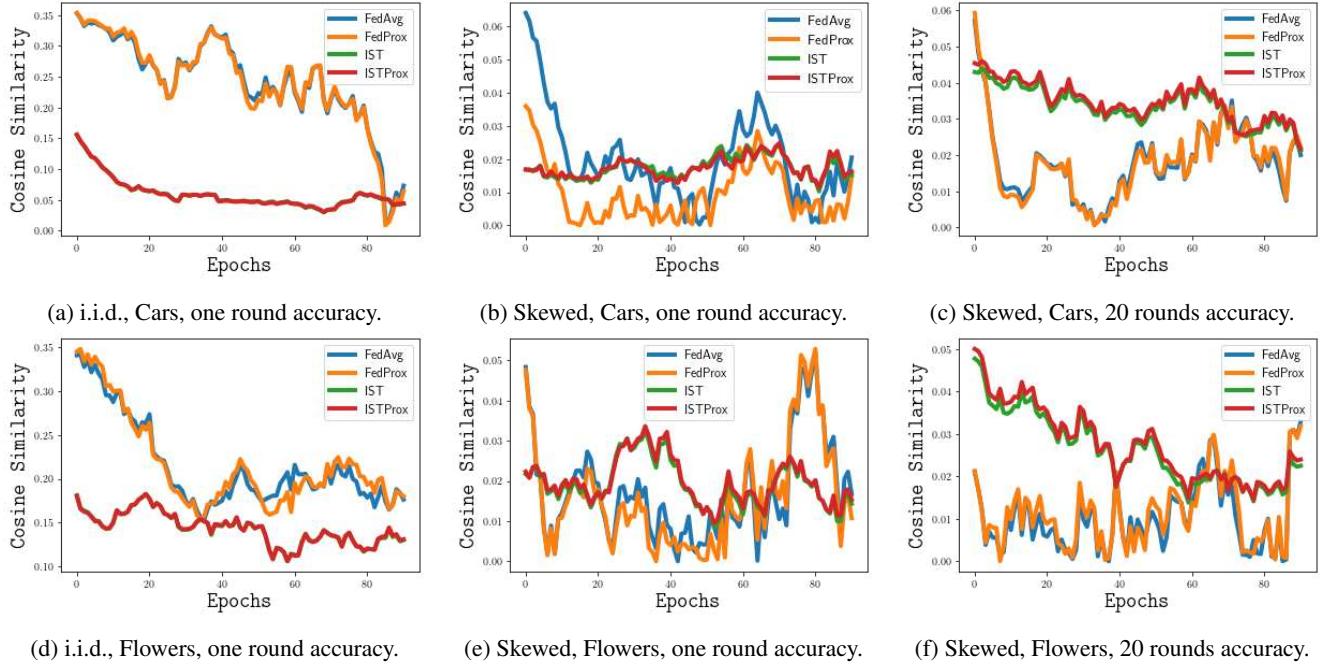


Figure 2: Comparison of similarity of the FL descent direction to a centrally-computed descent direction. One means the directions are identical; 0 represents de-correlation. At left is i.i.d. data; center is skewed data, similarity computed over one descent step; right is skewed data, similarity computed over 20 descent steps.

three sets of hyperparameters: one optimized for accuracy, one for gigaFLOPs, and one for communication. *The reason that we use three sets of hyperparameters is that the goals of maximizing accuracy, minimizing FLOPs and minimizing communication are often in direct opposition to one another, and these resources can often be traded (e.g., more local iterations and FLOPs for less communication).*

Next, we consider each data set in sequence. We first run each method to 100 teraFLOPs or five terabytes of communication on the data set, using the accuracy-optimized parameters. We then compute the maximum accuracy that any method obtained on the data set, and set 90% of this accuracy as a target. Then, for each method and each data set, we compute two additional numbers: the number of gigaFLOPs required to reach the target accuracy for five consecutive rounds (using the FLOP-optimized settings), and the number of gigabytes of communication required to reach the target accuracy (using the communication-optimized settings). If a method is unable to reach the target accuracy within the 100 teraFLOPs/five terabytes limit, it is said to have failed (indicated as FAIL in tables).

We estimate the amount of FLOPs used during training using the FlopCountAnalysis tool in the Python *fvcore* [12]. Similarly, for our communication costs we estimate the size of a model in bytes with the Python *torchinfo* [49] package.

The reason for this evaluation strategy is simple. In practice, the preferred FL method is likely to be the one that

reaches a high level of accuracy with little cost. Hence, we (somewhat arbitrarily, but precisely) define “high level of accuracy” as 90% of the peak accuracy observed over all methods for a given data set. Given this target, evaluation is then a matter of counting floating point operations and bytes to reach that target.

Note that other metrics are possible: power usage and elapsed wall clock time are two of the most obvious. However, both of these will be closely related to FLOPs and communication required, and hence we stick to these two.

5. Identically Distributed Data

5.1. Setup

In our first set of experiments, the images are randomly partitioned across the sites, so that each site has a randomly-selected (with replacement) subset of the overall set of images in the training data set. After a round of communication finishes, ten sites are randomly selected as being “on the network” and those sites participate in the present round of training. After local training finishes and all of the results have been communicated, another set of ten sites are randomly selected and the process continues.

For this simple “i.i.d.” setup, our experiments will answer a few key questions. First and foremost, should we be doing full backpropagation through a convolutional network, or should we be exclusively using a pre-trained back-

(a) Final accuracies						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	61.5	56.4	88.9	40.2	68.1	61.4
FedProx	58.4	50.1	85.4	35.6	66.8	60.5
MOON	57.7	48.0	86.9	37.6	66.2	58.6
FedAdam	49.4	16.9	81.0	28.9	54.0	40.5
FedNova	58.4	47.1	83.6	37.0	64.3	60.3
IST	64.3	50.8	90.7	39.8	69.5	49.4
ISTProx	64.8	50.3	91.3	40.6	68.2	56.7

(b) Communication (GB) to threshold acc.						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	131	62	51	79	55	59
FedProx	250	FAIL	124	FAIL	71	54
MOON	FAIL	FAIL	193	FAIL	143	74
FedAdam	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
FedNova	743	FAIL	430	608	309	160
IST	79	FAIL	37	119	69	FAIL
ISTProx	72	FAIL	59	134	72	219

(c) GFLOPs to theshold acc						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	13121	6245	5150	7892	5461	5884
FedProx	1000	FAIL	495	FAIL	283	218
MOON	FAIL	FAIL	581	1243	653	1207
FedAdam	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
FedNova	14872	FAIL	8614	12172	6181	3200
IST	315	FAIL	148	475	277	FAIL
ISTProx	287	FAIL	238	535	288	875

Table 2: Skewed data results, 100 sites, 10% participation.

bone? Second, when one considers computation and communication (rather than just final accuracy) do the methods compare differently than when the only metric is accuracy? And finally, does the data set used matter? Results are shown in Table 1. For each data set, the best result is indicated in bold.

5.2. Discussion

There are a few clear results. First, except for FedAdam and FedFull, the methods generally have very similar final accuracy, with the two IST-based methods being slightly higher. Thus, one key finding is that for data without skewness across sites, *if one is only interested in final accuracy, it does not seem to matter much which method is used*, and this result holds across data sets. The two exceptions to this are FedAdam and FedFull. FedAdam does not reach the same final accuracy as the other methods, but it generally does quite a good job reaching high accuracy with relatively little communication. This corresponds to the fact that it tends to converge in relatively few communication rounds.

FedFull is a very interesting case. Note that it has poor final accuracy—only 3% in the case of the Cars data set. As a result, it is labeled as “FAIL” in every case, because it can never reach 90% of the peak accuracy observed for the other methods on a given data set. One might ask, *Why is FedFull such a poor choice?* The answer is simple: *as we put a 100 teraFLOP/five terabyte “timeout” on the computation, FedFull timed out in every case.* That is, it *always* hit the 100 teraFLOP/five terabyte limit. Were FedFull allowed infinite computation and infinite communication, it is actually the preferred method, as the pre-trained backbones limit accuracy. For example, consider Figure 1, which plots accuracy as a function of FLOPs for both FedFull and FedProx on Cars. FedFull eventually reaches 90% accuracy, but it takes nearly *four orders of magnitude* more FLOPs to top out compared to FedProx. This FLOP usage gap is apparent in Fig. 1 where even the first communication round of FedFull requires more FLOPs than the amount required for the entire FedProx training. As a result, we would argue that full backpropagation is effectively unusable in resource-constrained FL. Note, however, that our FedFull was trained using an initialization consisting of random weights. Perhaps a hybrid method makes sense, where one begins with FedProx (with a pre-trained backbone) and then once the MLP begins to reach peak accuracy, we allow backpropagation over the backbone to begin. This issue deserves further investigation. That said, given our maximum FLOP count and communication, FedFull is not competitive and in the interest of space, we will not consider it further.

The final observation we make is the general superiority of IST, over each of the six data sets. It is perhaps not surprising that IST is efficient in terms of communication and in terms of computation. Each site obtains only a fraction of the MLP that is being trained, saving communication *and* computation at each site. However, it is surprising that it is *also* generally the most accurate method, in terms of final accuracy (or close to it) in every case. We will examine this finding more fully, later in the paper.

6. Effect of Data Skewness

6.1. Setup

How might these results change if data are not uniformly distributed across sites? To answer this question, we modify our experiments as follows. Rather than assigning each training image to each site with uniform probability, we sample the class probabilities for each site from a Dirichlet(0.01) distribution, and rerun our experiments (with 100 sites and 10 sites active in each round). This creates a relatively extreme—yet still realistic—amount of skew. For the six data sets we tested, the three-tuple (number of classes in data set, min number of classes at any site, max number of classes at any site) is: for Birds (200, 7,

(a) 1000 sites						
Change in acc., 2% → 6% participation						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	0.1	-0.8	0	0	-1.1	1.3
FedProx	0.1	-2	0.4	0.8	1	1.1
MOON	0.3	2.9	0.5	0.5	-0.2	1.4
FedAdam	12.8	16.8	7	12.5	12.1	12.4
FedNova	0.6	-0.5	-0.5	0	0.6	0.8
	1.2	7.3	0.5	2.4	0.7	4.9
ISTProx	0.4	6.3	0.4	1.8	7.6	4.1

(b) 100 sites						
Change in acc., 10% → 30% participation						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	-0.5	-0.8	0.6	-0.1	1.2	4
FedProx	-0.4	-2	3.7	2.1	-0.6	0
MOON	1.5	2.9	-0.5	-0.3	-0.1	2.8
FedAdam	2	14.9	-0.9	0.6	0.4	5
FedNova	0.3	-0.5	4.3	-0.7	-0.6	-0.3
IST	2.3	7.3	-1.7	0.1	0.1	6.2
ISTProx	3	6.3	-0.9	3	1.3	2.2

Table 3: Observed change in accuracy, computed as “(% correct with fewer sites) - (% correct with more sites)”, when active sites increases. Red indicates negative change.

20), Aircraft (100, 2, 15), Stanford Cars (196, 4, 20), Flowers (102, 3, 15), Textures (47, 1, 9), and CIFAR100 (100, 1, 13). This matches the realistic case where each site has access to only a few of the classes, and from this limited local information, an accurate global model must be constructed. Results are shown in Table 2.

6.2. Discussion

Probably the single most interesting finding is that for a highly-skewed situation, there seems to be divergence in performance, depending upon the data set. The two IST methods seem to perform best overall, with either IST or ISTProx being the superior option in terms of final accuracy, FLOPs, or communication for four of the six data sets. However, IST does relatively poorly on the venerable CFAR data set. CFAR is a somewhat strange benchmark, as the images are tiny (32×32) and widely varied within the data set, and both IST variants have problems with this. Nevertheless, it is a good stress test. Note that in the two cases where IST does not do the best, it is either simple FedAvg or FedProx that is the best performer. One of the most surprising findings is that FedAvg dominates all other methods in all metrics (accuracy, FLOPs, communication) on Cars.

Overall, across the two experiments (i.i.d. and skewed), we find that adding a proximal term to FedAvg (to obtain FedProx) seems to have little effect on accuracy, but *does*

(a) 1000 sites						
Change in comm., 2% → 6% participation						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	3.02	2.41	2.53	4.28	2.25	2.07
FedProx	2.76	NA	2.34	2.98	2	2.12
MOON	2.78	0.88	2.52	2.91	2.28	2.11
FedAdam	NA	NA	0	NA	NA	NA
FedNova	2.6	3.81	2.43	3.73	2.77	3.12
IST	1.41	0	1.07	1.95	1.03	NA
ISTProx	1.29	NA	1.15	2.2	0.48	NA

(b) 100 sites						
Change in comm., 10% → 30% participation						
	Birds	Cars	Flwrs	Aircraft	Textre	CFAR
FedAvg	0.65	NA	1.97	0.24	1.24	0
FedProx	∞	3.13	2.82	6.43	1.75	2.41
MOON	∞	NA	1.27	NA	5.31	7.39
FedAdam	NA	NA	1.29	NA	2.23	2.89
FedNova	0.3	-0.5	4.3	-0.7	-0.6	-0.3
IST	∞	NA	1.09	∞	2.73	6.44
ISTProx	NA	NA	NA	NA	NA	NA

Table 4: Multiplicative change in communication when tripling the number of active sites. Red text indicates more than a $3\times$ increase; black indicates less than a $3\times$ increase. “NA” indicates the case when there was a FAIL in both bases. ∞ indicates a change from non-fail to a FAIL, and 0 indicates a change from FAIL to non-fail.

speedup convergence, both in terms of reducing communication and FLOPs. However, IST seems to be the best performer overall, allowing for significant reductions in communication and FLOPs. The addition of a proximal term *does not* seem to help IST much. The other methods tested all seem to under-perform.

6.3. FL, Approximating Centralized Learning

How does the training process for a method (such as FedAvg and FedProx) that trains the entire model at each site and then reconciles the various versions of the model, compare with IST, which avoids this through decomposition?

To examine, we centrally train a model from initialization through convergence using min-batch gradient descent, obtaining model parameters M_i for the i th training epoch. Then, using FL, we start training in a federated setting from M_i for some number of communication rounds, to obtain new model parameters \hat{M}_i . We refer to the set of model parameters we *would* have obtained after processing exactly the same set of data centrally, via a series of mini-batches, as M_i^* . Then the cosine similarity of the directions both methods move, defined as $(\hat{M}_i - M_i) \cdot (M_i^* - M_i) / (\|\hat{M}_i - M_i\| \times \|M_i^* - M_i\|)$,

shows the utility of FL as an approximation for centralized learning. Cosine similarity close to one means FL and centralized training move in the same direction. Zero implies de-correlation.

Some results obtained in this way are plotted in Figure 2, where we consider FedAvg, FedProx, IST, and ISTProx. These plots depict this cosine similarity as a function of the training epochs. We plot running averages to smooth the data. The top three plots are for Cars, the bottom three Flwrs. The results for identically distributed data are at left, whereas skewed data are in the center, and at the right.

FedAvg and FedProx are shown to approximate the centralized gradient descent direction well early on, but the similarity drops after initial training. This makes sense, as initially, with a random initialization, learning is “easy” as the model is far from optimal. In the i.i.d. case FedAvg and FedProx perform well. But in the skewed case, after the initial training period, similarity to the centralized training direction drops. IST seems more well-behaved, with less variance. Especially considering how well the methods approximate the direction of centralized training through twenty rounds on skewed data (the right column), the IST-based methods seem to do better. This is especially pronounced for the Flwrs data set through the first 20 epochs, where IST averages a similarity of 0.035, compared to 0.01 for the averaging-based methods. Perhaps this is due to the high variance of the directions computed by FedAvg and FedProx—as estimators for the centrally-predicted direction—compared to IST. Examining why this is deserves further investigation in future work.

7. Varying Sites and Connectivity

7.1. Setup

All of our experiments thus far have focused on the simple case of 100 sites, where at each communication round, 10% of the sites are connected. We now ask: how do things change when the number of sites changes and/or the number of sites connected at each communication round changes? Keeping all of the other settings the same, we try 1000 sites, with either 2% or 6% of the sites connected at each communication round, and 100 sites, with either 10% of 30% of the sites connected at each round. Due to space constraints, we consider accuracy and communication here, and provide further results of these experiments in the appendix. In Table 3, we show the absolute change in percentage accuracy moving from 2% to 6% connectivity, and from 10% to 30% connectivity. In Table 4, we show the multiplicative change in communication required. Note that, as we increase the number of active sites by 3 \times , we might expect a 3 \times increase in communication if the same number of communication rounds are required. If fewer rounds are required, we would see an increase that is less than 3 \times .

7.2. Discussion

FedAdam, IST, and ISTProx are generally helped a bit by increasing participation. But for the other methods, the results obtained by increasing participation rates are mixed, with a surprising number of red values in both Table 3 and Table 4, indicating a decrease in accuracy or communication efficiency via the addition of extra active sites. In addition, it seems that the increase from 2% participation to 6% participation is much more helpful than 10% to 30%.

The lack of clear benefit to increasing the rate of participation is not surprising in retrospect. Aside from FedFull, each of the methods was able to run to convergence in the allotted FLOP/communication budget. The central question becomes: can we run to convergence in fewer rounds by doing more work in each round? This seems very unlikely, as this is akin to increasing the batch size in centralized learning, and expecting a significant decrease in the number of batches required for convergence. There is typically *some* decrease in the number of batches with increasing batch size, but it often is not significant. In reconciliation-based methods, this may be particularly inefficient as local training rounds cause the active sites to diverge and simply makes it more difficult to reconcile them. IST may have another advantage in that regard, in that more active sites does not increase communication for IST, it simply means that each site gets less of the full model. As long as each site can tolerate local learning using a smaller model (due to partitioning more ways), the learning may benefit from seeing a greater variety of data in each round. This certainly seems to be the case in when moving from 2% to 6%, which benefits IST.

8. Conclusion

We have designed a set of experiments to evaluate FL methods. Among the issues considered were what were the appropriate evaluation metrics (accuracy alone makes little sense), whether to use a pre-trained backbone, and the effect of connectivity on the efficiency of the learning. We find that overall, the method of decomposition of a neural network into independent subnetworks seems to be the best option. This has the benefit of decreasing communication and computation compared to “reconciliation-based” methods that train a copy of the model at each site, as the model is sharded, rather than broadcasted.

Acknowledgements. We would like thank the anonymous reviewers for their comments on the submitted version of the paper. Work presented in this paper has been supported by an NIH CTSA, award No. UL1TR003167 and by the NSF under grant Nos. 1918651, 1910803, 2008240, and 2131294.

References

- [1] Samiul Alam, Luyang Liu, Ming Yan, and Mi Zhang. Fedrolex: Model-heterogeneous federated learning with rolling sub-model extraction. *arXiv preprint arXiv:2212.01548*, 2022. 3
- [2] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, et al. Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 1:374–388, 2019. 1
- [3] Qi Chang, Hui Qu, Yikai Zhang, Mert Sabuncu, Chao Chen, Tong Zhang, and Dimitris N Metaxas. Synthetic learning: Learn from distributed asynchronous discriminator gan without sharing medical image data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13856–13866, 2020. 1
- [4] Zachary Charles, Kallista Bonawitz, Stanislav Chiknavaryan, Brendan McMahan, et al. Federated select: A primitive for communication-and memory-efficient federated learning. *arXiv preprint arXiv:2208.09432*, 2022. 3
- [5] Gary Cheng, Zachary Charles, Zachary Garrett, and Keith Rush. Does federated dropout actually work? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3387–3395, 2022. 3
- [6] Li Chou, Zichang Liu, Zhuang Wang, and Anshumali Shrivastava. Efficient and less centralized federated learning. In *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part I* 21, pages 772–787. Springer, 2021. 1
- [7] Arkabandhu Chowdhury, Mingchao Jiang, Swarat Chaudhuri, and Chris Jermaine. Few-shot image classification: Just use a library of pre-trained feature extractors and a simple classifier. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9445–9454, October 2021. 2
- [8] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014. 4
- [9] Enmao Diao, Jie Ding, and Vahid Tarokh. HeteroFL: Computation and communication efficient federated learning for heterogeneous clients. In *International Conference on Learning Representations*. 3
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 2
- [11] Chen Dun, Cameron R Wolfe, Christopher M Jermaine, and Anastasios Kyrillidis. ResIST: Layer-wise decomposition of resnets for distributed training. In *Uncertainty in Artificial Intelligence*, pages 610–620. PMLR, 2022. 3
- [12] FAIR. facebookresearch/fvcore: Collection of common code that's shared among different research projects in fair computer vision team. <https://github.com/facebookresearch/fvcore>. Accessed: 2022-10-11. 6
- [13] Michael Fleischman and Eduard Hovy. Fine grained classification of named entities. In *COLING 2002: The 19th International Conference on Computational Linguistics*, 2002. 1
- [14] Chaoyang He, Alay Dilipbhai Shah, Zhenheng Tang, Di Fan1Adarshan Nayyar Sivashunmugam, Keerti Bhogaraju, Mita Shimpi, Li Shen, Xiaowen Chu, Mahdi Soltanolkotabi, and Salman Avestimehr. Fedcv: A federated learning framework for diverse computer vision tasks, 2021. 1
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2, 4
- [16] Stefan Hinterstoisser, Vincent Lepetit, Paul Wohlhart, and Kurt Konolige. On pre-trained image features and synthetic images for deep learning. *CoRR*, abs/1710.10710, 2017. 2
- [17] Samuel Horváth, Stefanos Laskaridis, Mario Almeida, Ilias Leontiadis, Stylianos Venieris, and Nicholas Lane. Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 12876–12889. Curran Associates, Inc., 2021. 3
- [18] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Federated visual classification with real-world data distribution. In *European Conference on Computer Vision*, pages 76–92. Springer, 2020. 1
- [19] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2, 4
- [20] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021. 1
- [21] 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. 3
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 3

[23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 4

[24] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. 4

[25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 1

[26] Ang Li, Jingwei Sun, Binghui Wang, Lin Duan, Sicheng Li, Yiran Chen, and Hai Li. Lotteryfi: Empower edge intelligence with personalized and communication-efficient federated learning. In *2021 IEEE/ACM Symposium on Edge Computing (SEC)*, pages 68–79, 2021. 3

[27] Daiqing Li, Amlan Kar, Nishant Ravikumar, Alejandro F Frangi, and Sanja Fidler. Fed-sim: Federated simulation for medical imaging. arxiv e-prints, page. *arXiv preprint arXiv:2009.00668*, 2, 2020. 1

[28] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10713–10722, 2021. 1, 2, 3

[29] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020. 1

[30] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In I. Dhillon, D. Papailiopoulos, and V. Sze, editors, *Proceedings of Machine Learning and Systems*, volume 2, pages 429–450, 2020. 3

[31] Wenqi Li, Fausto Milletari, Daguang Xu, Nicola Rieke, Jonny Hancox, Wentao Zhu, Maximilian Baust, Yan Cheng, Sébastien Ourselin, M Jorge Cardoso, et al. Privacy-preserving federated brain tumour segmentation. In *International workshop on machine learning in medical imaging*, pages 133–141. Springer, 2019. 1

[32] Fangshuo Liao and Anastasios Kyrillidis. On the convergence of shallow neural network training with randomly masked neurons. *CoRR*, abs/2112.02668, 2021. 3

[33] Tao Lin, Sebastian U. Stich, Kumar Kshitij Patel, and Martin Jaggi. Don’t use large mini-batches, use local sgd. In *International Conference on Learning Representations*, 2020. 1, 2

[34] Xiulong Liu, Jiannong Cao, Yanni Yang, Wenyu Qu, Xibin Zhao, Keqiu Li, and Didi Yao. Fast rfid sensory data collection: Trade-off between computation and communication costs. *IEEE/ACM Trans. Netw.*, 27(3):1179–1191, jun 2019. 5

[35] Yang Liu, Anbu Huang, Yun Luo, He Huang, Youzhi Liu, Yuanyuan Chen, Lican Feng, Tianjian Chen, Han Yu, and Qiang Yang. Fedvision: An online visual object detection platform powered by federated learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(08):13172–13179, Apr. 2020. 1

[36] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *CoRR*, abs/1306.5151, 2013. 4

[37] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017. 2, 3

[38] Maria-Elena Nilsback and Andrew Zisserman. A visual vocabulary for flower classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1447–1454, 2006. 4

[39] Keiron O’Shea and Ryan Nash. An introduction to convolutional neural networks. *CoRR*, abs/1511.08458, 2015. 2

[40] Sashank J. Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021. 3

[41] Sebastian U. Stich. Local SGD converges fast and communicates little. In *International Conference on Learning Representations*, 2019. 1, 2

[42] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Caltech-ucsd birds-200-2011 (cub-200-2011). Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 4

[43] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020. 3

[44] Qihan Wang, Chen Dun, Fangshuo Liao, Chris Jermaine, and Anastasios Kyrillidis. LOFT: Finding lottery tickets through filter-wise training. *arXiv preprint arXiv:2210.16169*, 2022. 3

[45] Cameron R Wolfe, Jingkang Yang, Arindam Chowdhury, Chen Dun, Artun Bayer, Santiago Segarra, and Anastasios Kyrillidis. GIST: Distributed training for large-scale graph convolutional networks. *arXiv preprint arXiv:2102.10424*, 2021. 3

[46] Tianbao Yang. Trading computation for communication: Distributed stochastic dual coordinate ascent. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. 5

[47] Tien-Ju Yang, Dhruv Guliani, Françoise Beaufays, and Giovanni Motta. Partial variable training for efficient on-device federated learning. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4348–4352. IEEE, 2022. 3

[48] Ze Yang, Tiange Luo, Dong Wang, Zhiqiang Hu, Jun Gao, and Liwei Wang. Learning to navigate for fine-grained classification. *CoRR*, abs/1809.00287, 2018. 1

[49] Tyler Yep. Tylyerpytorchinfo: View model summaries in pytorch! <https://github.com/TylerYep/torchinfo>. Accessed: 2022-10-11. 6

[50] Peihua Yu and Yunfeng Liu. Federated object detection: Optimizing object detection model with federated learning. In

Proceedings of the 3rd International Conference on Vision, Image and Signal Processing, pages 1–6, 2019. 1

- [51] Binhang Yuan, Cameron R. Wolfe, Chen Dun, Yuxin Tang, Anastasios Kyrillidis, and Chris Jermaine. Distributed learning of fully connected neural networks using independent subnet training. *Proc. VLDB Endow.*, 15(8):1581–1590, apr 2022. 3
- [52] Jian Zhang, Christopher De Sa, Ioannis Mitliagkas, and Christopher Ré. Parallel SGD: When does averaging help? *arXiv preprint arXiv:1606.07365*, page arXiv:1606.07365, June 2016. 1, 2
- [53] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018. 1
- [54] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *CoRR*, abs/1806.00582, 2018. 3