

Federated Visual Classification with Real-World Data Distribution

Tzu-Ming Harry Hsu^{1,2*}, Hang Qi², and Matthew Brown²

¹ MIT CSAIL

stmharry@mit.edu

² Google Research

{stmharry, hangqi, mtbr}@google.com

Abstract. Federated Learning enables visual models to be trained on-device, bringing advantages for user privacy (data need never leave the device), but challenges in terms of data diversity and quality. Whilst typical models in the datacenter are trained using data that are independent and identically distributed (IID), data at source are typically far from IID. Furthermore, differing quantities of data are typically available at each device (imbalance). In this work, we characterize the effect these real-world data distributions have on distributed learning, using as a benchmark the standard Federated Averaging (FedAvg) algorithm. To do so, we introduce two new large-scale datasets for species and landmark classification, with realistic per-user data splits that simulate real-world edge learning scenarios. We also develop two new algorithms (FedVC, FedIR) that intelligently resample and reweight over the client pool, bringing large improvements in accuracy and stability in training. The datasets are made available online.

1 Introduction

Federated learning (FL) is a privacy-preserving framework, originally introduced by McMahan et al. [19], for training models from decentralized user data residing on devices at the edge. Models are trained iteratively across many federated rounds. For each round, every participating device (a.k.a. *client*), receives an initial model from a central server, performs stochastic gradient descent (SGD) on its local training data and sends back the gradients. The server then aggregates all gradients from the participating clients and updates the starting model. FL preserves user privacy in that the raw data used for training models never leave the devices throughout the process. In addition, differential privacy [20] can be applied for a theoretically bounded guarantee that no information about individuals can be derived from the aggregated values on the central server.

Federated learning is an active area of research with a number of open questions [16, 13] remaining to be answered. A particular challenge is the distribution of data at user devices. Whilst in centralized training, data can be assumed to be

* work done whilst at Google.

independent and identically distributed (IID), this assumption is unlikely to hold in federated settings. Decentralized training data on end-user devices will vary due to user-specific habits, preferences, geographic locations, etc. Furthermore, in contrast to the streamed batches from a central data store in the data center, devices participating in an FL round will have differing amounts of data available for training.

In this work, we study the effect these heterogeneous client data distributions have on learning visual models in a federated setting, and propose novel techniques for more effective and efficient federated learning. We focus in particular on two types of distribution shift: **Non-Identical Class Distribution**, meaning that the distribution of visual classes at each device is different, and **Imbalanced Client Sizes**, meaning that the number of data available for training at each device varies. Our key contributions are:

- **We analyze the effect of learning with per-user data** in real-world datasets, in addition to carefully controlled setups with parametric (Dirichlet) and natural (geographic) distributions.
- **We propose two new algorithms** to mitigate per-client distribution shift and imbalance, substantially improving classification accuracy and stability.
- **We provide new large-scale datasets** with per-user data for two classification problems (natural world and landmark recognition) to the community.

Ours is the first work to our knowledge that attempts to train large-scale visual classification models for real-world problems in a federated setting. We expect that more is to be done to achieve robust performance in this and related settings, and are making our datasets available to the community to enable future research in this area³.

2 Related Work

Synthetic Client Data Several authors have explored the **FedAvg** algorithm on synthetic non-identical client data partitions generated from image classification datasets. McMahan et al. [19] synthesize pathological non-identical user splits from the MNIST dataset, sorting training examples by class labels and partitioning into shards such that each client is assigned 2 shards. They demonstrate that **FedAvg** on non-identical clients still converges to 99% accuracy, though taking more rounds than identically distributed clients. In a similar sort-and-partition manner, [35,26] use extreme partitions of the CIFAR-10 dataset to form a population consisting of 10 clients in total. In contrast to these pathological data splits, Yurochkin et al. [33] and Hsu et al. [11] synthesize more diverse non-identical datasets with Dirichlet priors.

Realistic Datasets Other authors look at more realistic data distributions at the client. For example, Caldas et al. [2] use the Extended MNIST dataset [3] split

³ https://github.com/google-research/google-research/tree/master/federated_vision_datasets




Fig. 1: **iNaturalist Species Distribution.** Visualized here are the distributions of Douglas-Fir and Red Maple in the continental US within iNaturalist. In a federated learning context, visual categories vary with location, and users in different locations will have very different training data distributions.

over the writers of the digits and the CelebA dataset [17] split by the celebrity on the picture. The Shakespeare and Stack Overflow datasets [7] contain natural per-user splits of textual data using roles and online user ids, respectively. Luo et al. [18] propose a dataset containing 900 images from 26 street-level cameras, which they use to train object detectors. These datasets are however limited in size, and are not representative of data captured on user devices in a federated learning context. Our work aims to address these limitations (see Section 4).

Variance reduction methods have been used in the federated learning literature to correct for the distribution shift caused by heterogeneous client data. Sahu et al. [24] introduce a proximal term to client objectives for bounded variance. Karimireddy et al. [14] propose to use control variates for correcting client updates drift. Importance sampling is a classic technique for variance reduction in Monte Carlo methods [12,9] and has been used widely in domain adaption literature for countering covariate and target shift [23,34,21]. In this work, we adopt a similar idea of importance reweighting in a novel federated setting resulting in augmented client objectives. Different from the classic setting where samples are drawn from one proposal distribution which has the same support as the target, heterogeneous federated clients form multiple proposal distributions, each of which has partially common support with the target.

3 Federated Visual Classification Problems

Many problems in visual classification involve data that vary around the globe [5,8]. This means that the distribution of data visible to a given user device will vary, sometimes substantially. For example, user observations in the citizen scientist app iNaturalist will depend on the underlying species distribution in that region (see Figure 1). Many other factors could potentially influence the data present on a device, including the interests of the user, their photography habits, etc. For this study we choose two problems with an underlying geographic variation to illustrate the general problem of non-identical user data, *Natural Species Classification* and *Landmark Recognition*.

Natural Species Classification We create a dataset and classification problem based on the iNaturalist 2017 Challenge [28], where images are contributed by a




Fig. 2: **iNaturalist Distribution.** In (a) we show the re-balancing of the original iNaturalist-2017 dataset. In (b) and (c) we show class and example counts vs clients for our 5 iNaturalist partitionings with varying levels of class distribution shift and size imbalance. The client count is different in each partitioning.

community of citizen scientists around the globe. Domain experts take pictures of natural species and provide annotations during field trips. Fine-grained visual classifiers could potentially be trained in a federated fashion with this community of citizen scientists without transferring images.

Landmark Recognition We study the problem of visual landmark recognition based on the 2019 Landmark Recognition Challenge [30], where the images are taken and uploaded by Wikipedia contributors. It resembles a scenario where smartphone users take photos of natural and architectural landmarks (e.g., famous buildings, monuments, mountains, etc.) while traveling. Landmark recognition models could potentially be trained via federated learning without uploading or storing private user photos at a centralized party.

Both datasets have data partitioning per user, enabling us to study a realistic federated learning scenario where labeled images were provided by the user and learning proceeds on-device. For experimentation in lab, we use a simulation engine for federated learning algorithms, similar to TensorFlow Federated [6].

4 Datasets

In the following section, we describe in detail the datasets we develop and analyze key distributional statistics as a function of user and geo-location. We plan to make these datasets available to the community.

4.1 iNaturalist-User-120k and iNaturalist-Geo Splits

iNaturalist-2017 [28] is a large scale fine-grained visual classification dataset comprised of images of natural species taken by citizen scientists. It has 579,184 training examples and 95,986 test examples covering over 5,000 classes. Images in this dataset are each associated with a fine-grained species label, a longitude-latitude coordinate where the picture was originally taken, and authorship information.

The iNaturalist-2017 training set has a very long-tailed distribution over classes as shown in Figure 2a, while the test set is relatively uniform over classes.

While studying learning robustly with differing training and test distributions is a topic for research [29] in itself, in our federated learning benchmark, we create class-balanced training and test sets with uniform distributions. This allows us to focus on distribution variations and imbalance at the *client level*, without correcting for overall domain shift between training and test sets.

To equalize the number of examples across classes, we first sort all class labels by their count and truncate tail classes with less than 100 training examples. This is then followed by subsampling per-class until all remaining classes each have 100 examples. This results in a balanced training set consisting of 1,203 classes and 120,300 examples. We use this class-balanced iNaturalist subset for the remainder of the paper.

The iNaturalist-2017 dataset includes user ids, which we use to partition the balanced training set into a population of 9,275 clients. We refer to this partitioning as **iNaturalist-User-120k**. This contributor partitioning resembles a realistic scenario where images are collected per-user.

In addition, for study the federated learning algorithms with client populations of varying levels of deviation from the global distribution, we also generate a *wide range* of populations by partitioning the dataset at varying levels of granularity according to the geographic locations.

To utilize the geo-location tags, we leverage the S2 grid system, which defines a hierarchical partitioning of the planet surface. We perform an adaptive partitioning similar to [31]. Specifically, every S2 cell is recursively subdivided into four finer-level cells until no single cell contains more than N_{\max} examples. Cells ending up with less than N_{\min} examples are discarded. With this scheme, we are able to control the granularity of the resulting S2 cells such that a smaller N_{\max} results in a larger client count. We use $N_{\max} \in \{30k, 3k, 1k, 100\}$, $N_{\min} = 0.01N_{\max}$ and refer to the resulting data partitionings as **iNaturalist-Geo-{30k, 3k, 1k, 100}**, respectively. Rank statistics of our geo- and per-user data splits are shown in Figures 2b and 2c.

4.2 Landmarks-User-160k

Google Landmarks Dataset V2 (GLD-v2) [30] is a large scale image dataset for landmark recognition and retrieval, consisting of 5 million images with each attributed to one of over 280,000 authors. The full dataset is noisy: images with the same label could depict landmark exteriors, historical artifacts, paintings or sculptures inside a building. For benchmarking federated learning algorithms on a well-defined image classification problem, we use the cleaned subset (GLD-v2-clean), which is a half the size of the full dataset. In this set, images are discarded if the computed local geometric features from which cannot be matched to at least two other images with the same label [22].

For creating a dataset for federated learning with natural user identities, we partition the GLD-v2-clean subset according to the authorship attribute. In addition, we mitigate the long tail while maintaining realism by requiring every landmark to have at least 30 images and be visited by at least 10 users, meanwhile requiring every user to have contributed at least 30 images that depict 5 or more

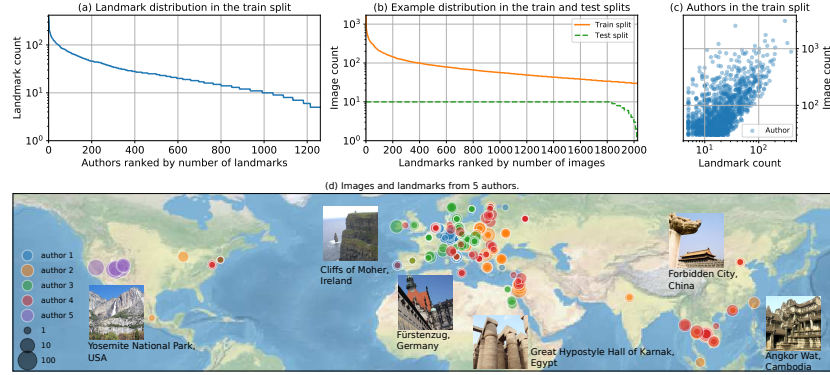


Fig. 3: **Landmarks-User-160k Distribution**. Images are partitioned according to the authorship attribute from the GLD-v2 dataset. Filtering is applied to mitigate long tail in the train split.

landmarks. The resulting dataset has 164,172 images of 2,028 landmarks from 1,262 users, which we refer to as the train split of **Landmarks-User-160k**.

Following the dataset curation in [30], the test split is created from the leftover images in GLD-v2-clean whose authors do not overlap with those in the train split. The test split contains 19,526 images and is well-balanced among classes. 1,835 of the landmarks have exactly 10 test images, and there is a short tail for the rest of the landmarks due to insufficient samples (Figure 3).

5 Methods

The datasets described above contain significant distribution variations among clients, which presents considerable challenges for efficient federated learning [16, 13]. In the following, we describe our baseline approach of Federated Averaging algorithm (**FedAvg**) (Section 5.1) and two new algorithms intended to specifically address the non-identical class distributions and imbalanced client sizes present in the data (Sections 5.2 and 5.3 respectively).

5.1 Federated Averaging and Server Momentum

A standard algorithm [19] for FL, and the baseline approach used in this work, is Federated Averaging (**FedAvg**). See Algorithm 1. For every federated round, K clients (the *report goal*) are randomly selected with uniform probability from a pool of active clients. Selected clients, indexed by k , download the same starting model θ_t from a central server and perform local SGD optimization, minimizing an empirical loss $L(b)$ over local mini-batches b with learning rate η , for E epochs before sending the accumulated model update $\Delta\theta_t^k$ back to the server. The server then averages the updates from the reporting clients $\bar{g}_t = \sum_{k=1}^K \frac{n_k}{n} \Delta\theta_t^k$ with

Algorithm 1: FedAvg, FedIR, and FedVC.

Server training loop:
Initialize θ_0
for each round $t = 0, 1, \dots$ **do**
 Subset of K clients $\leftarrow \text{SelectClients}(K)$
 for each client $k = 1, 2, \dots, K$ **do in parallel**
 $\Delta\theta_t^k \leftarrow \text{ClientUpdate}(k, \theta_t)$
 $\bar{g}_t \leftarrow \text{AggregateClient}(\{\Delta\theta_t^k\}_{k=1}^K)$
 $\theta_{t+1} \leftarrow \theta_t - \gamma \bar{g}_t$

SelectClients(K):
return K clients sampled uniformly ▷ with probability $\propto n_i$ for client i

ClientUpdate(k, θ_t):
 $\theta \leftarrow \theta_t$
for each local mini-batch b over E epochs **do** ▷ over S steps
 $\theta \leftarrow \theta - \eta \nabla L(b; \theta)$ ▷ $\nabla \tilde{L}(b; \theta)$ in Eq.3
return $\Delta\theta \leftarrow \theta_t - \theta$ to server

AggregateClient($\{\Delta\theta_t^k\}_{k=1}^K$):
return $\sum_{k=1}^K \frac{n_k}{n} \Delta\theta_t^k$, where $n = \sum_{k=1}^K n_k$ ▷ $\frac{1}{K} \sum_{k=1}^K \Delta\theta_t^k$

weights proportional to the sizes of clients' local data and finishes the federated round by applying aggregated updates to the starting model $\theta_{t+1} \leftarrow \theta_t - \gamma \bar{g}_t$, where γ is the server learning rate. Given this framework, alternative optimizers can be applied. FedAvgM [11] has been shown to improve robustness to non-identically distributed client data. It uses a momentum optimizer on the server with the update rule $\theta_{t+1} \leftarrow \theta_t - \gamma v_t$, where $v_t \leftarrow \beta v_{t-1} + \bar{g}_t$ is the exponentially weighted moving average of the model updates with powers of β .

5.2 Importance Reweighted Client Objectives

Now we address the non-identical class distribution shift in federated clients. Importance reweighting are commonly used for learning from datasets distributed differently from a target distribution. Given that the distribution variation among clients is an inherent characteristic of FL, we propose the following scheme.

Consider a target distribution $p(x, y)$ of images x and class labels y on which a model is supposed to perform well (e.g. a validation dataset known to the central server), and a predefined loss function $\ell(x, y)$. The objective of learning is to minimize the expected loss $\mathbb{E}_p[\ell(x, y)]$ with respect to the target distribution p . SGD in the centralized setting achieves this by minimizing an empirical loss on mini-batches of IID training examples from the same distribution, which are absent in the federated setting. Instead, training examples on a federated client k are sampled from a client-specific distribution $q_k(x, y)$. This implies that the empirical loss being optimized on every client is a *biased* estimator of the loss with respect to the target distribution, since $\mathbb{E}_{q_k}[\ell(x, y)] \neq \mathbb{E}_p[\ell(x, y)]$.

We propose an importance reweighting scheme, denoted FedIR, that applies importance weights $w_k(x, y)$ to every client's local objective as $\tilde{\ell}(x, y) =$

$\ell(x, y)w_k(x, y)$, where $w_k(x, y) = \frac{p(x, y)}{q_k(x, y)}$. With the importance weights in place, an unbiased estimator of loss with respect to the target distribution can be obtained using training examples from the client distribution

$$\mathbb{E}_p[\ell(x, y)] = \sum_{x, y} \frac{\ell(x, y)p(x, y)}{q_k(x, y)} q_k(x, y) = \mathbb{E}_{q_k} \left[\ell(x, y) \frac{p(x, y)}{q_k(x, y)} \right]. \quad (1)$$

Assuming that all clients share the same conditional distribution of images given a class label as the target, i.e. $p(x|y) \approx q_k(x|y) \forall k$, the importance weights can be computed on every client directly from the class probability ratio

$$w_k(x, y) = \frac{p(x, y)}{q_k(x, y)} = \frac{p(y)p(x|y)}{q_k(y)q_k(x|y)} \approx \frac{p(y)}{q_k(y)}. \quad (2)$$

Note that this computation does not sabotage the privacy-preserving property of federated learning. The denominator $q_k(y)$ is private information available locally at and never leaves client k , whereas the numerator $p(y)$ does not contain private information about clients and can be transmitted from the central server with minimal communication cost: C scalars in total for C classes.

Since scaling the loss also changes the effective learning rate in the SGD optimization, in practice, we use self-normalized weights when computing loss over a mini-batch b as

$$\tilde{L}(b) = \frac{\sum_{(x, y) \in b} \ell(x, y)w_k(x, y)}{\sum_{(x, y) \in b} w_k(x, y)}. \quad (3)$$

This corresponds to the self-normalized importance sampling in the statistics literature [9]. **FedIR** does not change server optimization loops and can be applied together with other methods, such as **FedAvgM**. See Algorithm 1.

5.3 Splitting Imbalanced Clients with Virtual Clients

The number of training examples in users' devices vary in the real world. Imbalanced clients can cause challenges for both optimization and engineering practice. Previous empirical studies [19, 11] suggest that the number of local epochs E at every client has crucial effects on the convergence of **FedAvg**. A larger E implies more optimization steps towards local objectives being taken, which leads to slow convergence or divergence due to increased variance. Imbalanced clients suffer from this optimization challenge even when E is small. Specifically, a client with a large number of training examples takes significantly more local optimization steps than another with fewer training examples. This difference in steps is proportional to the difference in the number of training examples. In addition, a client with an overly large training dataset will take a long time to compute updates, creating a bottleneck in the federated learning round. Such clients would be abandoned by a FL production system in practice, if failing to report back to the central server within a certain time window [1].

Table 1: **Training Dataset Statistics.** Note that while CIFAR-10/100 and iNaturalist datasets each have different partitionings with different levels of identicalness, the underlying data pool is unchanged and thus sharing the same centralized learning baselines.

Dataset	Clients		Classes	Examples	Centralized Accuracy
	Count	Size Imbalance	Count	Count	
Synthetic					
CIFAR-10	100	✗	10	50,000	86.16%
CIFAR-100	100	✗	100	50,000	55.21%
iNaturalist Geo Splits	11 to 3606	✓	1,203	120,300	57.90%
Real-World					
iNaturalist-User-120k	9,275	✓	1,203	120,300	57.90%
Landmarks-User-160k	1,262	✓	2,028	164,172	67.05%

We hence propose a new *Virtual Client* (FedVC) scheme to overcome both issues. The idea is to conceptually split large clients into multiple smaller ones, and repeat small clients multiple times such that all *virtual* clients are of similar sizes. To realize this, we fix the number of training examples used for a federated learning round to be N_{VC} for every client, resulting in exactly $S = N_{VC}/B$ optimization steps taken at every client given a mini-batch size B . Concretely, consider a client k with a local dataset \mathcal{D}_k with size $n_k = |\mathcal{D}_k|$. A random subset consisting of N_{VC} examples is uniformly resampled from \mathcal{D}_k for every round the client is selected. This resampling is conducted without replacement when $n_k \geq N_{VC}$; with replacement otherwise. In addition, to avoid underutilizing training examples from large clients, the probability that any client is selected for a round is set to be proportional to the client size n_k , in contrast to uniform as in FedAvg. Key changes are outlined in Algorithm 1. It is clear that FedVC is equivalent to FedAvg when all clients are of the same size.

6 Experiments

We now present an empirical study using the datasets and methods of Sections 4 and 5. We start by analyzing the classification performance as a function of non-identical data distribution (Section 6.1), using the CIFAR10/100 datasets. Next we show how *Importance Reweighting* can improve performance in the more non-identical cases (Section 6.2). With real user data, clients are also imbalanced, we show how this can be mitigated with *Federated Virtual Clients* in Section 5.3. Finally we present a set of benchmark results with the per-user splits of iNaturalist and Landmark datasets (Section 6.4). A summary of the datasets used is provided in Table 1. Implementation details are deferred to Section 6.5.

Metrics When using the same dataset, the performance of a model trained with federated learning algorithms is inherently upper bounded by that of a model trained in the centralized fashion. We evaluate the *relative accuracy*, defined as $\text{Acc}_{\text{federated}}/\text{Acc}_{\text{centralized}}$, and compare this metric under different types of

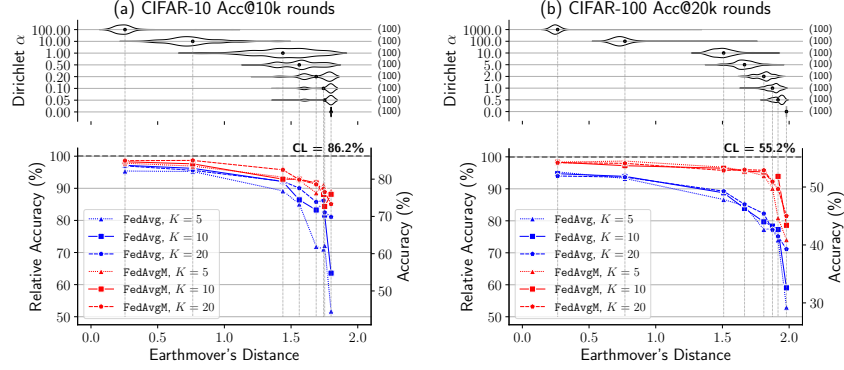


Fig. 4: **Relative Accuracy v.s. Non-identicalness.** Federated learning experiments are performed on (a) CIFAR-10 and (b) CIFAR-100 using local epoch $E = 1$. The top row demonstrates the distributions of EMD of clients with different data partitionings. Total client counts are annotated to the right, and the weighted average of all client EMD is marked. Data is increasingly non-identical to the right. The dashed line indicates the centralized learning performance. The best accuracies over a grid of hyperparameters are reported (see Appendix ??).

budgets. The centralized training baseline uses the same configurations and hyperparameters for a fair comparison.

6.1 Classification Accuracy vs Distribution Non-Identicalness

Our experiments use CIFAR10/100 datasets to characterize classification accuracy with a continuous range of distribution non-identicalness. We follow the protocol described in [11] such that the class distribution of every client is sampled from a Dirichlet distribution with varying concentration parameter α .

We measure distribution non-identicalness using an average *Earthmover's Distance* (EMD) metric. Specifically, we take the discrete class distribution \mathbf{q}_i for every client, and define the population's class distribution as $\mathbf{p} = \sum_i \frac{n_i}{n} \mathbf{q}_i$, where $n = \sum_i n_i$ counts training samples from all clients. The non-identicalness of a dataset is then computed as the weighted average of distances between clients and the population: $\sum_i \frac{n_i}{n} \text{Dist}(\mathbf{q}_i, \mathbf{p})$. $\text{Dist}(\cdot, \cdot)$ is a distance metric between two distributions, which we, in particular, use $\text{EMD}(\mathbf{q}, \mathbf{p}) \equiv \|\mathbf{q} - \mathbf{p}\|_1$, bounded between $[0, 2]$.

Figures 4a and 4b show the trend in classification accuracy as a function of distribution non-identicalness (average EMD difference). We are able to approach centralized learning accuracy with data on the identical end. A substantial drop around an EMD of 1.7 to 2.0 is observed in both datasets. Applying momentum on the server, FedAvgM significantly improves the convergence under heterogeneity




Fig. 5: **Comparing Base Methods with and without FedIR.** Accuracy shown at 2.5k communication rounds. Centralized learning accuracy marked with dashed lines.

Table 2: **Accuracy of Federated Virtual Client on iNaturalist.** Acc@round denotes the accuracy at a FL communication round. Acc@batch denotes the batch count accumulated over the largest clients per round, and is a proxy for a fixed time budget.

Data	Method	FedVC	K	Acc@Round(%)			Acc@Batch(%)		
				1k	2.5k	5k	10k	25k	50k
Geo-3k	FedAvg	✗	10	47.0	47.9	48.7	37.8	44.4	46.5
	FedAvgM	✗	10	47.2	50.4	45.0	42.5	47.1	44.9
	FedAvg	✓	10	37.4	46.2	52.8	46.2	53.1	55.5
	FedAvgM	✓	10	49.7	54.8	56.7	54.8	56.7	57.1
User-120k	FedAvg	✗	10	34.7	39.7	41.3	37.8	39.8	42.9
	FedAvgM	✗	10	31.9	39.2	41.3	32.3	41.6	43.4
	FedAvg	✓	10	31.3	39.7	43.9	39.7	48.9	52.8
	FedAvgM	✓	10	37.9	43.7	49.1	43.7	47.4	54.6
Centralized				57.9					

conditions for all datasets. Using more clients per round (larger report goal K) is also beneficial for training but has diminishing returns.

6.2 Importance Reweighting

Importance Reweighting is proposed for addressing the per-client distribution shift. We evaluate FedIR with both FedAvg and FedAvgM on both two datasets with natural user splits: iNaturalist-User-120k and Landmarks-User-160k.

For Landmarks, we experiment with two different training schemes: (a) fine-tuning the entire network (*all layers*) end to end, (b) only training the last *two layers* while freezing the network backbone. We set the local epochs to $E = 5$ and experiment with report goals $K = \{10, 50, 100\}$, respectively.

The result in Figure 5 shows a consistent improvement on the Landmarks-User-160k dataset over the FedAvg baseline. While FedAvgM gives the most significant improvements in all runs, FedIR further improves the convergence speed and accuracy especially when the report goal is small (Figure 7).

Landmarks-User-160k (EMD = 1.94) has more skewed data distribution than iNaturalist-User-120k (EMD = 1.83) and benefits more from FedIR.

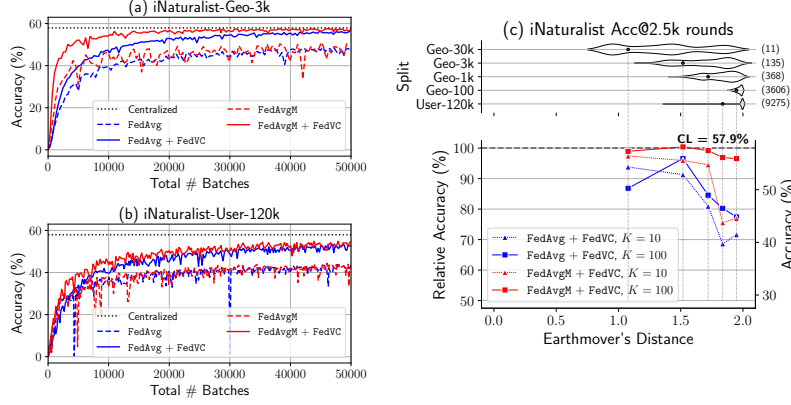


Fig. 6: **Learning with Federated Virtual Clients.** Curves on the left are learned on the iNaturalist geo-partitioning Geo-3k and user split User-120k each with 135 clients and 9275 clients. Experiments on multiple iNaturalist partitionings are shown on the right, plotting relative accuracy at 2.5k communication rounds to mean EMD. Centralized learning achieves a 57.9% accuracy.

6.3 Federated Virtual Clients

We apply the Virtual Clients scheme (**FedVC**) to both **FedAvg** and **FedAvgM** and evaluate its efficacy using iNaturalist user and geo-location datasets, each of which contains significantly imbalanced clients. In the experiments, 10 clients are selected for every federated round. We use a mini-batch size $B = 64$ and set the virtual client size $N_{VC} = 256$.

Figure 6 demonstrates the efficiency and accuracy improvements gained via **FedVC** when clients are imbalanced. The convergence of vanilla **FedAvg** suffers when clients perform excessive local optimization steps. In iNaturalist-Geo-3k, for example, clients can take up to 46 (i.e. $3000/64$) local steps before reporting to the server. To show that **FedVC** utilizes data efficiently, we report accuracy at fixed batch budgets in addition to fixed round budgets. Batch budget is calculated by summing the number of local batches taken for the largest client per round. As shown in Table 2, **FedVC** consistently yields superior accuracy on both **FedAvg** and **FedAvgM**. Learning curves in Figure 6 show that **FedVC** also decreases the learning volatility and stabilizes learning.

iNaturalist per-user and geo-location datasets reflect varying degrees of non-identicalness. Figure 6c, though noisier, exhibits a similar trend compared to Figure 4. The performance degrades as the degree of non-identicalness, characterized by EMD, increases.

6.4 Federated Visual Classification Benchmarks

Having shown that our proposed modifications to **FedAvg** indeed lead to a speedup in learning on both iNaturalist and Landmarks, we wish to also provide some

Table 3: **iNaturalist-User-120k accuracy**. Numbers reported at fixed communication rounds. K denotes the report goal per round.

Method	FedVC	FedIR	K	Accuracy@Rounds(%)		
				1k	2.5k	5k
FedAvg	✓	✗	10	31.3	39.7	43.9
FedAvg	✓	✗	100	36.9	46.5	51.4
FedAvg	✓	✓	10	30.1	41.3	47.5
FedAvg	✓	✓	100	35.5	44.8	49.8
FedAvgM	✓	✗	10	37.9	43.7	49.1
FedAvgM	✓	✗	100	53.0	56.1	57.2
FedAvgM	✓	✓	10	38.4	42.1	47.0
FedAvgM	✓	✓	100	51.3	54.3	56.2
Centralized				57.9		




Fig. 7: **Landmarks-User-160k Learning Curves**. Only the last two layers of the network are fine-tuned. FedIR is also shown due to its ability to address skewed training distribution as presented in this dataset.

benchmark results on natural user partitioning with reasonable operating points. We hope that these datasets can be used for understanding real-world federated visual classification, and act as benchmarks for future improvements.

iNaturalist-User-120k The iNaturalist-User-120k data has 9,275 clients and 120k examples, containing 1,203 species classes. We use report goals $K = \{10, 100\}$. FedVC samples $N_{VC} = 256$ examples per client. A summary of the benchmark results is shown in Table 3.

Notice that FedAvgM with FedVC and a large report goal of $K = 100$ has a 57.2% accuracy, almost reaching the same level as in centralized learning (57.9%). With that said, there is still plenty of room to improve performance with small reporting clients and round budgets. Being able to learn fast with a limited pool of clients is one of the critical research areas for practical visual FL.

Landmarks-User-160k The Landmarks-User-160k dataset comprises 164,172 images for 2,028 landmarks, divided among 1,262 clients. We follow the setup in Section 6.2 where we experiment with either training the whole model or fine-tuning the last two layers. Report goal $K = \{10, 50, 100\}$ are used.

Similarly, FedAvgM with the $K = 100$ is able to achieve 65.9% accuracy at 5k communication rounds, which is just 1.2% off from centralized learning. Interestingly, when we train only the last two layers with FL, the accuracy is as well not far off from centralized learning (39.8% compared to 40.3%)

Table 4: **Landmarks-User-160k Accuracy.**

Method	FedIR	K	Accuracy@Rounds(%)					
			Two layers			All layers		
			1k	2.5k	5k	1k	2.5k	5k
FedAvg	X	10	4.2	14.6	24.6	18.2	38.1	49.7
FedAvg	X	50	4.5	16.5	26.0	20.9	42.0	53.3
FedAvg	X	100	4.9	16.5	26.3	21.9	42.3	53.4
FedAvg	✓	10	6.3	17.4	26.6	19.6	38.5	51.7
FedAvg	✓	50	7.4	19.7	28.8	26.0	45.2	55.0
FedAvg	✓	100	7.2	20.1	29.0	26.5	45.7	55.2
FedAvgM	X	10	23.0	30.1	30.8	29.4	44.1	53.7
FedAvgM	X	50	29.9	36.4	38.6	55.2	62.0	64.8
FedAvgM	X	100	31.9	37.4	39.6	56.3	63.4	65.0
FedAvgM	✓	10	26.5	32.1	31.3	27.9	45.1	53.5
FedAvgM	✓	50	31.6	37.5	38.9	53.1	61.6	63.2
FedAvgM	✓	100	33.7	38.3	39.8	57.7	64.1	65.9
Centralized			40.27			67.05		

6.5 Implementation Details

We use MobileNetV2 [25] pre-trained on ImageNet [4] for both iNaturalist and Landmarks experiments; for the latter, a 64-dimensional bottleneck layer between the 1280-dimensional features and the softmax classifier. We replaced BatchNorm with GroupNorm [32] due to its superior stability for FL tasks [10]. During training, the image is randomly cropped then resized to a target input size of 299×299 (iNaturalist) or 224×224 (Landmarks) with scale and aspect ratio augmentation similar to [27]. A weight decay of 4×10^{-5} is applied. For CIFAR-10/100 experiments, we use a CNN similar to LeNet-5 [15] which has two 5×5 , 64-channel convolution layers, each precedes a 2×2 max-pooling layer, followed by two fully-connected layers with 384 and 192 channels respectively and finally a softmax linear classifier. This model is not the state-of-the-art on the CIFAR datasets, but is sufficient to show the relative performance for our investigation. Weight decay is set to 4×10^{-4} . Unless otherwise stated, we use a learning rate of 0.01 and momentum of 0.9 in FedAvgM, kept constant without decay for simplicity. The client batch size is 32 in Landmarks and 64 for others.

7 Conclusions

We have shown that large-scale visual classifiers can be trained using a privacy-preserving, federated approach, and highlighted the challenges that per-user data distributions pose for learning. We provide two new datasets and benchmarks, providing a platform for other explorations in this space. We expect others to improve on our results, particularly when the number of participating clients and round budget is small. There remain many challenges for Federated Learning that are beyond the scope of this paper: real world data may include domain shift, label noise, poor data quality and duplication. Model size, bandwidth and unreliable client connections also pose challenges in practice. We hope our work inspires further exploration in this area.

Acknowledgements We thank Andre Araujo, Grace Chu, Tobias Weyand, Bingyi Cao, Huizhong Chen, Tomer Meron, and Hartwig Adam for their valuable feedback and support.

References

1. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C.M., Konečný, J., Mazzocchi, S., McMahan, B., Overveldt, T.V., Petrou, D., Ramage, D., Roselander, J.: Towards federated learning at scale: System design. In: SysML 2019 (2019), <https://arxiv.org/abs/1902.01046> 8
2. Caldas, S., Wu, P., Li, T., Konečný, J., McMahan, H.B., Smith, V., Talwalkar, A.: LEAF: A benchmark for federated settings. arXiv preprint arXiv:1812.01097 (2018) 2
3. Cohen, G., Afshar, S., Tapson, J., Van Schaik, A.: Emnist: Extending mnist to handwritten letters. In: 2017 International Joint Conference on Neural Networks (IJCNN). pp. 2921–2926. IEEE (2017) 2
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 248–255. IEEE (2009) 14
5. Doersch, C., Singh, S., Gupta, A., Sivic, J., Efros, A.A.: What makes Paris look like Paris? Communications of the ACM **58**(12), 103–110 (2015) 3
6. Google: TensorFlow Federated (2019), <https://www.tensorflow.org/federated> 4
7. Google: TensorFlow Federated Datasets (2019), https://www.tensorflow.org/federated/api_docs/python/tff/simulation/datasets 3
8. Hays, J., Efros, A.A.: IM2GPS: Estimating geographic information from a single image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1–8. IEEE (2008) 3
9. Hesterberg, T.: Weighted average importance sampling and defensive mixture distributions. Technometrics **37**(2), 185–194 (1995) 3, 8
10. Hsieh, K., Phanishayee, A., Mutlu, O., Gibbons, P.B.: The non-IID data quagmire of decentralized machine learning. arXiv preprint arXiv:1910.00189 (2019) 14
11. Hsu, T.M.H., Qi, H., Brown, M.: Measuring the effects of non-identical data distribution for federated visual classification. arXiv preprint arXiv:1909.06335 (2019) 2, 7, 8, 10
12. Kahn, H., Marshall, A.W.: Methods of reducing sample size in Monte Carlo computations. Journal of the Operations Research Society of America **1**(5), 263–278 (1953) 3
13. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al.: Advances and open problems in federated learning. arXiv preprint arXiv:1912.04977 (2019) 1, 6
14. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S.J., Stich, S.U., Suresh, A.T.: Scaffold: Stochastic controlled averaging for on-device federated learning. arXiv preprint arXiv:1910.06378 (2019) 3
15. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE **86**(11), 2278–2324 (1998) 14
16. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated learning: Challenges, methods, and future directions. arXiv preprint arXiv:1908.07873 (2019) 1, 6
17. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of International Conference on Computer Vision (ICCV) (2015) 3
18. Luo, J., Wu, X., Luo, Y., Huang, A., Huang, Y., Liu, Y., Yang, Q.: Real-world image datasets for federated learning. arXiv preprint arXiv:1910.11089 (2019) 3
19. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Artificial Intelligence and Statistics. pp. 1273–1282 (2017) 1, 2, 6, 8

20. McMahan, H.B., Ramage, D., Talwar, K., Zhang, L.: Learning differentially private recurrent language models. In: International Conference on Learning Representations (ICLR) (2018) [1](#)
21. Ngiam, J., Peng, D., Vasudevan, V., Kornblith, S., Le, Q.V., Pang, R.: Domain adaptive transfer learning with specialist models. arXiv preprint arXiv:1811.07056 (2018) [3](#)
22. Ozaki, K., Yokoo, S.: Large-scale landmark retrieval/recognition under a noisy and diverse dataset. arXiv preprint arXiv:1906.04087 (2019) [5](#)
23. Saerens, M., Latinne, P., Decaestecker, C.: Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation* **14**(1), 21–41 (2002) [3](#)
24. Sahu, A.K., Li, T., Sanjabi, M., Zaheer, M., Talwalkar, A., Smith, V.: On the convergence of federated optimization in heterogeneous networks. arXiv preprint arXiv:1812.06127 (2018) [3](#)
25. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4510–4520 (2018) [14](#)
26. Sattler, F., Wiedemann, S., Müller, K.R., Samek, W.: Robust and communication-efficient federated learning from non-IID data. arXiv preprint arXiv:1903.02891 (2019) [2](#)
27. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1–9 (2015) [14](#)
28. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The iNaturalist species classification and detection dataset. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8769–8778 (2018) [3](#), [4](#)
29. Van Horn, G., Perona, P.: The devil is in the tails: Fine-grained classification in the wild. arXiv preprint arXiv:1709.01450 (2017) [5](#)
30. Weyand, T., Araujo, A., Cao, B., Sim, J.: Google Landmarks Dataset v2 - a large-scale benchmark for instance-level recognition and retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2020) [4](#), [5](#), [6](#)
31. Weyand, T., Kostrikov, I., Philbin, J.: Planet-photo geolocation with convolutional neural networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 37–55. Springer (2016) [5](#)
32. Wu, Y., He, K.: Group normalization. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 3–19 (2018) [14](#)
33. Yurochkin, M., Agarwal, M., Ghosh, S., Greenewald, K., Hoang, N., Khazaeni, Y.: Bayesian nonparametric federated learning of neural networks. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 7252–7261 (2019) [2](#)
34. Zhang, K., Schölkopf, B., Muandet, K., Wang, Z.: Domain adaptation under target and conditional shift. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 819–827 (2013) [3](#)
35. Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., Chandra, V.: Federated learning with non-IID data. arXiv preprint arXiv:1806.00582 (2018) [2](#)