

FedRepOpt: Gradient Re-parametrized Optimizers in Federated Learning

Kin Wai Lau^{1,2}, Yasar Abbas Ur Rehman¹, Pedro Porto Buarque de Gusmão³, Lai-Man Po², Lan Ma¹, and Yuyang Xie¹

¹ TCL AI LAB, Hong Kong, China

² Department of Electrical Engineering, City University of Hong Kong, Hong Kong, China {kinwailau6-c}@my.cityu.edu.hk

³ University of Surrey, United Kingdom

Abstract. Federated Learning (FL) has emerged as a privacy-preserving method for training machine learning models in a distributed manner on edge devices. However, on-device models face inherent computational power and memory limitations, potentially resulting in constrained gradient updates. As the model's size increases, the frequency of gradient updates on edge devices decreases, ultimately leading to suboptimal training outcomes during any particular FL round. This limits the feasibility of deploying advanced and large-scale models on edge devices, hindering the potential for performance enhancements. To address this issue, we propose FedRepOpt, a gradient re-parameterized optimizer for FL. The gradient re-parameterized method allows training a simple local model with a similar performance as a complex model by modifying the optimizer's gradients according to a set of model-specific hyperparameters obtained from the complex models. In this work, we focus on VGG-style and Ghost-style models in the FL environment. Extensive experiments demonstrate that models using FedRepOpt obtain a significant boost in performance of 16.7% and 11.4% compared to the RepGhost-style and RepVGG-style networks, while also demonstrating a faster convergence time of 11.7% and 57.4% compared to their complex structure. Codes are available at <https://github.com/StevenLauHKHK/FedRepOpt>.

Keywords: Federated Learning · Reparameterization · CNN

1 Introduction

Collaborative on-device training of deep learning models promises enormous potential for the future Internet of Things (IoT) applications in smart homes, health care, robotics, autonomous driving, environmental monitoring, and finance [16, 18, 19, 22, 23, 32]. Among these collaborative on-device training schemes, Federated Learning (FL) has garnered significant attention from research and the industrial community. It is particularly valued for enabling collaborative learning of feature representations from real-world data while preserving data privacy [25, 35]. In FL, distributed devices (clients) collaboratively train a common deep-learning model on their local data under the orchestration of a central

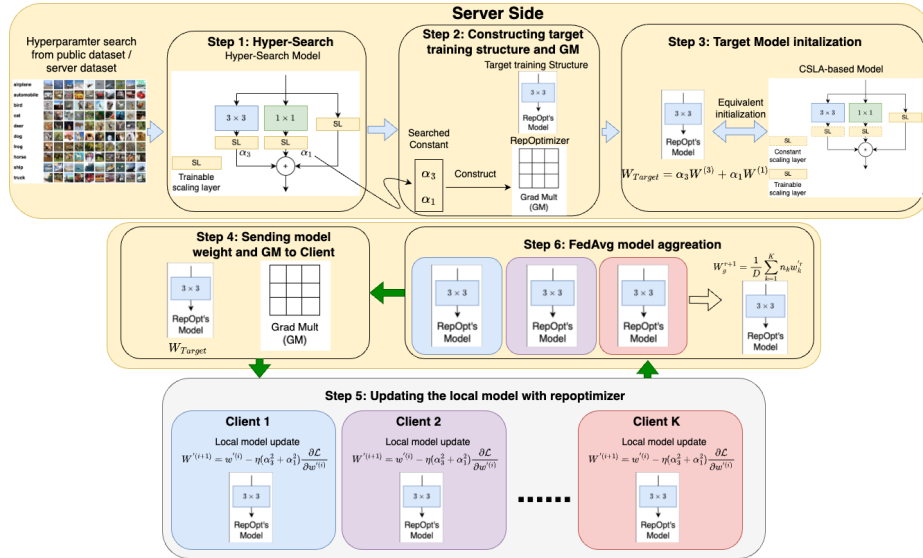


Fig. 1: Overview of federated reoptimizer framework. It comprises six steps training pipeline: (1) Server performs the model-specific hyperparameter search (i.e., α_3 and α_1) using the HS dataset from public dataset or server dataset. (2) Convert the parallel branch CSLA structure by a single operator and the equivalent training dynamic constant hyper-parameter called gradient multiplier (GM). (3) Initialize the target training structure with the equivalent CSLA structure. (4) Sending the initialized model and GM to each client (5) During the client training, the gradient is multiplied with a constant scalar GM. After training, the updated models are sent to the server. (6) The server aggregates all the client models to obtain a new global model. These steps repeat until the model converges.

server [34]. Although local training can help reduce communication costs, current practices overlook the existing computational and memory constraints of real devices, which in practice leads to restriction of the types of architectures that can be trained and the number of partial model updates a device can perform, ultimately limiting the scope and benefits of the FL training.

In practical FL, a diverse range of clients coexist, including those equipped with high-end devices as well as those utilizing low-end devices. A naive solution to address the computational constraints of the low-end devices would be to train a desired complex model only on selected high-end devices and use a structural compression method, such as re-parameterization [9], to enable model deployment in both low-end devices and high-end devices during *inference*. However, this solution would (1) restrict the learning to the data available on high-end devices and (2) require different architectures for training and inference.

Instead of modifying the model architecture, a better approach could focus on *changing the gradients of the model used during the training process* as shown in Figure 1. Gradient re-parameterization (GR) [6], also known as RepOpt, facilitates the training and inference of the compressed version of the multi-branch

models without any loss of accuracy. This is accomplished by manipulating the gradients of the compressed model according to hyper-parameters that are specific to the corresponding multi-branch model. These hyper-parameters are commonly obtained by training the multi-branch model on an auxiliary Hyper-Search (HS) dataset. This in turn mitigates the requirement of separate models during FL training and inference, while maintaining similar performance to the multi-branch non-compressed version of the models. Additionally, it allows FL systems to simultaneously train deep learning models on both low-end and high-end devices achieving full utility of the available pool of data and further counterbalance the effects of computational constraints of the client’s devices.

In this work, we propose FedRepOpt as a novel application for the RepOpt models [6], i.e., RepOpt-VGG and RepOpt-GhostNet, in FL. Our key findings suggest that (1) FedRepOpt-based models in vanilla FL offer better performance than their plain-style counterparts (VGG and GhostNet) and similar performance to their non-re-parameterized multi-branch counterparts both in IID and Non-IID settings, paving the way for the deployment of large and complex models in FL while being computationally less expensive. (2) Similar to the centralized settings, FedRepOpt-based models are agnostic to the type of the HS datasets for finding the *hyperparameters* and provide similar results with both IID and Non-IID versions of the HS datasets. This, in turn, eases the burden on the server in cases where HS datasets are not publicly available. (3) Interestingly, in FL, the learning pattern of the FedRepOpt-models and their complex non-re-parameterized counterparts are nearly similar. (4) In contrast to plain-style models, the large FedRepOpt-based models are less affected by the local SGD momentum, offering additional performance improvement with less computational complexity. The main contributions of this work are as follows:

1. We conduct the first systematic study of training RepOpt-based models in FL settings. This establishes a baseline for naïvely implementing various RepOpt-based models using FL; shedding light on the importance of RepOpt-based models in FL.
2. We propose FedRepOpt and validate its effectiveness under different FL configurations. We demonstrate that FedReOpt can achieve a significant boost in performance while incurring low computational footprints.
3. We further demonstrate that FedReOpt is not only adaptive to cross-silo settings but also effective in cross-device settings, even when dealing with non-iid cases.

2 Related Work

2.1 Reparameterization

Structural Reparameterization (SR) [7–9] aims to *simplify complex multi-branch networks into single-branch architectures during inference by applying linear transformations to the model’s parameters without sacrificing performance*. These approaches facilitate the utilization of complex network structures for effective

feature representation learning during the training phase, followed by deploying a streamlined and parameter-efficient network during the inference stage. For instance, RepVGG [9] proposed a multi-branch CNN architecture by introducing extra 1×1 convolutional layers in parallel with the original VGG network's 3×3 convolutional layers, enabling the acquisition of more expressive feature representations during training. After the training, the extra 1×1 layers were integrated with the 3×3 layers through the linear convolution transformations, resulting in a VGG-like model that incurs no additional parameters or computational cost during inference. Empirical evidence demonstrated that RepVGG surpassed the performance of the original VGG model without introducing any inference-time overhead.

Similarly, [7] proposed the Diverse Branch Block (DBB) method, which enhanced the representation capacity of an individual convolutional layer by combining multiple branches with diverse complexities to enrich the feature representation. During training, the DBB block contains a sequence of convolutions, multi-scale convolutions, and average pooling layers. The convolutional blocks are converted into a single convolutional layer during inference by utilizing the linear transformation properties of convolution. Inspired by the success of RepVGG and DBB, [8] introduced RepLKNet, which adopted a similar strategy by utilizing parallel large and small kernels during training. After the training process, the small kernels were merged into the large kernel, empowering it to capture both global and local information, thereby yielding improvements in model performance.

While SR techniques enhance model performance during inference without incurring extra parameters or computational costs, it is important to note that they do require additional training costs that cannot be overlooked. In the recent work, [6] introduced RepOptimizers (RepOpt), a gradient parameterization technique that addresses the training costs associated with Structural Reparameterization (SR). RepOpt migrates the additional 1×1 convolution layer into an optimizer during the training process, achieving comparable performance to SR models without incurring any extra training costs. This is achieved by modifying the gradients based on a set of model-specific hyperparameters.

2.2 Federated Visual Representation Learning

Federated Learning (FL) [21] has drawn much attention in recent years due to privacy concerns about the user's data transmitted and kept in the centralized server for training. In FL, data remains on the client's side while the server only manages the models collected from the clients. It aims to learn a global model in a decentralized fashion and obtain an accuracy similar (or better) compared to the model trained in a centralized fashion.

FedAvg [21] is a classical FL algorithm that iteratively trains a global model by training the local models at the client and aggregating (averaging) them at the server. Although it is easy to implement and guarantees the convergence of the model, this simple aggregation strategy performs poorly in realistic scenarios where the data are not Independent and Identically Distributed (non-

iid). The accuracy of the learned global model is much lower than the model trained with centralized data. To tackle this issue, the existing methodologies have explored better global aggregation [11, 24, 30] and local training [1, 13]. For instance, [11] proposes local-model training loss as a weighting coefficient for aggregation, allowing each participant’s contribution to be weighted based on their local performance. Additionally, [24] introduces adaptive federated optimization approaches that leverage knowledge from past iterations by employing separate gradient-based optimization on the server side. Besides, [13] proposes a variation reduction method to reduce the client drift between the central server and each local client during the local updates. It further reduces the communication rounds between the server and local clients.

The limited bandwidth of the internet connection between edge devices and global servers is a bottleneck that can lead to longer training times when synchronizing parameters. Recent research in this area has [2, 14, 26] proposed to quantize model updates into fewer bits, while others [10, 12, 20, 28] proposed to sparsify local updates by filtering certain values. Another approach in [4] proposed partial parameter synchronization, as many parameters stabilize before the final model convergence, especially in the early stages. In contrast to these approaches, our method introduces the gradient re-parameterized technique to reduce the communication round and further boost the performance. To the best of our knowledge, no work has studied the effects of the Gradient reparameterization [6] in FL. We believe the work proposed here will provide the initial foundation for searching the replicas of even large-scale models with similar performance but significantly lower parameter costs.

3 Preliminaries

3.1 Constant Scale Linear Addition (CSLA)

The CSLA are linear blocks that contain multi-branch linear trainable operators (e.g., convolution or scale layers) and constant scales, as shown in Figure 1. According to [6], the CSLA blocks can be transformed into a single trainable operator significantly reducing the model parameters. The equivalent training dynamics of CSLA blocks are then realized by multiplying the gradients of the single trainable operator with constant multipliers which are derived from the constant scales of the CSLA blocks. These constant multipliers are called Grad Mult. From the perspective of FL, training a CSLA-equivalent compressed model has a lower computational cost and maintains accuracy similar to the CSLA model.

The CSLA-based models [6], such as CSLA-RepVGG, follow a similar structure as models that can be structurally re-parameterized [9]. For instance, as shown in Figure 1 (step 3), the CSLA-based model (CSLA-RepVGG) contains a 3×3 , a 1×1 convolutional layer, and a scaling layer in parallel. As mentioned in [6], the CSLA block can be transformed into a 3×3 kernel. The equivalency of the transformed 3×3 kernel and the CSLA block can be maintained by the

modified gradients in the optimizer (RepOpt) during the local model updating. The details of the RepOpt are explained in the following subsection.

3.2 RepOptimizer (RepOpt)

The core idea of RepOpt [6] is to shift the structural priors of the model into an optimizer. As explained in subsection 3.1, we replace the multi-branch linear convolutional operators in the CSLA block with a single operator W' and modify the gradients by multiplying them by a constant multiplier, called Grad Mult. We observe that the output of the CSLA block and the W' with a modified optimizer are always identical in any number of rounds in the FL settings (see. Figure 2) given the same training data (i.e., $Y_{CSLA} = Y_{GM}$). The details of the equivalency proof can be found in [6].

To ensure that both the outputs of the CSLA-based model and its RepOpt-based counterparts are equivalent in FL, two rules are required to follow: First, the RepOpt’s model W' should be initialized with the equivalent initial parameters as the CSLA model (i.e., $W' \leftarrow \alpha_3 W_3 + \alpha_1 W_1$). Second, the gradients of the RepOpt’s model should be multiplied by $(\alpha_3^2 + \alpha_1^2)$ during the weight update in each iteration. In this work, we use the regular SGD as an optimizer. The weight W' of the RepOpt-based model (CSLA-equivalent) can then be updated by $W'^{(i+1)} \leftarrow W'^{(i)} - \lambda(\alpha_3^2 + \alpha_1^2) \frac{\partial L}{\partial W'^{(i)}}$, where L is the objective function and λ is the learning rate.

4 Method

In this section, we provide the details of our systematic study on training a CSLA and its RepOpt counterparts in FL. We consider RepVGG [27], RepGhostNet [5], and their CSLA counterparts CSLA-VGG and CSLA-GhostNet [6] as our baseline. These baselines are compared with the RepOpt-based counterparts, i.e., RepOpt-VGG and RepOpt-GhostNet.

4.1 FedRepOpt

We consider Z partitions $\{d_z\}_{z=1}^Z$ of dataset D to compose Z decentralized clients with n_z samples on each local data set [6]. We follow a six-step process to train the RepOpt-based version of the CSLA-based models as shown in Figure 1. The server first performs the *model-specific* hyperparameter search of the CSLA-based model using the HS dataset (i.e., finding α_3, α_1 using the CSLA-VGG model), which can be a public dataset or the dataset available at the server. The server then initiates the FL training by transmitting the structural reparameterized version of the CSLA-based model and the RepOpt to the clients. We called this model here as FedRepOpt-model. For simplicity, here, we assume the FedReOpt-VGG version of the RepOpt-VGG model [6].

At each communication round r of FL, the server randomly selects K clients participating in the training and initializes their local models with the global

Algorithm 1 *FedRepOpt*: Let us consider the server randomly selecting K clients at the given round. The clients train the SSL model with L layers for E local epochs on its dataset z_k with n_k number of samples. The FL optimization lasts R rounds.

Input: $K, R, n_k, d_k, \eta, E, L, HS, CSLA$

Output: w_g^R

Central server does:

- 1: $\alpha_3, \alpha_1, W', \text{RepOpt} = \mathbf{Train}$ (CSLA, HS)
- 2: $w_g = W'$
- 3: **for** $r = 1, \dots, R$ **do**
- 4: Server randomly selects K clients.
- 5: **for** $k = 1, \dots, K$ **do**
- 6: $w_k^r, n_k = \mathbf{TrainLocally}(k, w_g^r, E, \alpha_3, \alpha_1)$
- 7: **end for**
- 8: **Aggregation:**
- 9: $w_g^{r+1} = \frac{1}{D} \sum_{k=1}^K n_k w_k^r$
- 10: **end for**

TrainLocally ($k, w_g^r, E, \alpha_3, \alpha_1$):

- 1: Initialize $w_k^r = w_g^r$
 - 2: **for** $e = 1, \dots, E$ **do**
 - 3: $w_k^{(i+1)r} = w_k^{(i)r} - \eta(\alpha_3^2 + \alpha_1^2) \frac{\partial \mathcal{L}_k}{\partial w_k^{(i)r}}$
 - 4: **end for**
 - 5: Upload w_k^r, n_k to the server.
-

model weights w_g^r . Then, each decentralized client k learns the local model using its local data d_k with RepOpt. The local model weight update can be written as follows.

$$w_k^{(i+1)r} = w_k^{(i)r} - \eta(\alpha_3^2 + \alpha_1^2) \frac{\partial \mathcal{L}_k}{\partial w_k^{(i)r}} \quad (1)$$

The server then receives the local models $\{w_k^r\}_{k=1}^K$ and aggregates them based on the aggregation strategy, such as FedAvg [21], to generate a new global model w_g^{r+1} for the next $r + 1$ round.

$$w_g^{r+1} = \frac{1}{D} \sum_{k=1}^K d_k w_k^r \quad (2)$$

As shown in Figure 1, the step 4 to step 6 are repeated until model convergence. The pseudo-code for the FedRepOpt is shown in Algorithm 1. We evaluate the performance of the global model at the end of each round.

5 Experiments

5.1 Datasets

We conducted our experiments with Tiny ImageNet for FL training and evaluation and CIFAR-100 as an HS dataset to find *model-specific* hyperparameters. To simulate a realistic FL environment, we generate the IID/Non-IID variants of Tiny ImageNet based on actual class labels using the Dirichlet coefficient α , where the lower value exhibits greater heterogeneity. The dataset is randomly partitioned into 10/100 shards for *cross-silo/cross-device* settings to mimic the settings of having 10/100 disjoint clients participating in FL.

5.2 Architecture and Implementation

We consider the FL version of the two architectures proposed in the original work on RepOpt [6]. We consider Fed-RepVGG [27], Fed-RepGhostNet [5], and their CSLA counterparts Fed-CSLA-VGG and Fed-CSLA-GhostNet [6] as our baseline. It should be noted that RepVGG and RepGhostNet have different but equivalent model structures during the training (Tr) and inference (Inf) stages. This results in 2 different but equivalent architectures for RepVGG and RepGhostNet models, i.e., Fed-RepVGG-Tr, Fed-RepVGG-Inf, Fed-RepGhostNet-Tr, and Fed-RepGhostNet-Inf. The Inf stage model is obtained by utilizing the reparameterization technique [6] to merge the multi-branch blocks of the model into single-branch blocks. To facilitate a fair comparison between the conventional single-branch design and FedRepOpt-based models, we also train the Inf stage models from scratch. We compare Fed-RepX-Tr, Fed-RepX-Inf, and Fed-CSLA-based models with FedRepOpt-based models in the following experiments, where 'X' denotes either VGG or GhostNet. The details of the network architecture can be found in the supplementary material.

We developed the FedRepOpt on top of the Flower [3] federated learning framework. Unless otherwise specified, we keep the same settings as proposed by [6] for local training and evaluation. For the measurement of the training speed TPR (time/round), we test all the models on six NVIDIA RTX3090 using a batch size of 32. We first train one round to warm up the hardware, followed by 10 rounds to record the average training time per round.

5.3 FL Training

For cross-silo FL, unless otherwise specified, the local training on each client lasts for $E = 1$ local epochs per round. We set the total number of rounds R to 240 to ensure that each client acquires sufficient participation during the FL training phase. The selection of E and R is based on our empirical observations. We set the momentum of 0 throughout the FL training and use FedAvg [21] as an aggregation strategy to combine all the client models at the server. After each round, the aggregated global model is first evaluated on the test dataset at the server before being transmitted to the clients. For cross-device FL, the local

Table 1: Tiny ImageNet Accuracy and training speed on Non-IID and IID Setting. The search dataset, in this case, was CIFAR100. TPR represents the training time per round (Lower is better). 'Tr' and 'Inf' stand for training stage architecture and inference stage architecture, respectively.

Model	Param (M) TPR (sec)		Tiny ImageNet			
			Non-IID		IID	
			Top-1	Top-5	Top-1	Top-5
Fed-RepGhost-Tr 0.5x	1.289	57	13.45 ± 0.48	35.61 ± 0.83	14.24 ± 0.40	36.22 ± 0.59
Fed-RepGhost-Inf 0.5x	1.284	52	13.10 ± 0.20	35.32 ± 0.29	12.93 ± 0.40	34.01 ± 0.79
Fed-CSLA-GhostNet 0.5x	1.286	54	14.93 ± 0.29	37.69 ± 0.33	16.17 ± 0.59	39.04 ± 1.17
FedRepOpt-GhostNet 0.5x	1.284	51	15.29 ± 0.85	38.16 ± 0.93	16.19 ± 0.47	39.13 ± 0.56
Fed-RepVGG-B1-Tr	55.8	74	35.30 ± 0.49	65.09 ± 0.79	40.69 ± 0.33	68.29 ± 0.35
Fed-RepVGG-B1-Inf	50.2	46	33.45 ± 0.11	63.66 ± 0.48	37.79 ± 0.12	65.89 ± 0.39
Fed-CSLA-VGG-B1	55.7	78	36.97 ± 0.22	66.31 ± 0.31	42.09 ± 0.12	69.33 ± 0.06
FedRepOpt-VGG-B1	50.2	47	37.27 ± 0.34	66.37 ± 0.09	42.15 ± 0.03	69.35 ± 0.12

training lasts for $E = 5$ epochs per round and we simulate the FL training for $R = 1000$ rounds. During each round, the server randomly selects 10% of the client to participate in the collaborative model training. All training schemes are implemented with PyTorch and Flower [3].

5.4 FedRepOpt-based vs Fed-Rep-based models

To verify the effectiveness of FedRepOpt, we first compare the performance of FedRepOpt-based models against the baseline Fed-Rep-based models. We report this comparison in terms of parameter size (M), training time per round (TPR), and Top1% and Top5 % accuracy in the IID and Non-IID settings as shown in Table 1.

First, FedRepOpt-VGG and FedRepOpt-GhostNet incur similar parameter counts as plain-style Fed-RepVGG-Inf and Fed-RepGhostNet-Inf, which results in 10% and 0.38% parameter savings compared to the multi-branch Fed-RepVGG-Tr and Fed-RepGhost-Tr respectively. A similar conclusion holds when comparing the parameters count of FedRepOpt-VGG and FedRepOpt-GhostNet with Fed-CSLA-VGG and Fed-CSLA-GhostNet. Regarding training speed, we found that the FedRepOpt version of VGG and GhostNet models are 36% and 10.5% faster than the multi-branch Fed-Rep-Tr-based counterparts. Compared to their CSLA-based counterparts, the FedRepOpt version of VGG and GhostNet shows an improvement in the training speed of 39% and 5.55%, respectively. In terms of accuracy, FedRepOpt-based models outperform multi-branch Fed-Rep-Tr-based models and single-branch Fed-Rep-Inf-based models. It is interesting to observe that the findings differ from the centralized training discussed in [6], where Rep-based models (i.e., RepVGG-Tr and RepGhost-Tr) achieve comparable results to RepOpt-based (i.e., RepOpt-VGG and RepOpt-Ghost) and CSLA models (i.e., CSLA-VGG and CSLA-Ghost). We conjecture that the batch normalization layer in the multi-branch of Rep-based models adversely affects federated learning, as highlighted in [15, 31]. This is attributed to the mismatch between the local and global statistical parameters in batch normalization layers, which leads to gradient deviation between the local and global models.




Fig. 2: Learning behavior of CSLA-based models and their reoptimized versions. (a) Ghost Model IID (b) VGG Model (IID) (c) Ghost Model NIID (d) VGG Model NIID

Second, our proposed methods, FedRepOpt-VGG and FedRepOpt-GhostNet achieve similar performance as Fed-CSLA-VGG and Fed-CSLA-GhostNet respectively. These results show that the equivalency holds between CSLA-based models and their RepOpt-based counterparts even in FL.

The fact that FedRepOpt-based models provide similar results as the multi-branch Fed-CSLA-based models in FL with fewer parameter counts makes them the natural choice for FL training.

5.5 Learning Behavior of RepOpt and CSLA models

To verify that FedRepOpt-based models follow the same learning behavior as Fed-CSLA-based models like the centralized training in [6], we plot the global models' accuracy for each round as shown in Figure 2. Surprisingly, we found that the FL learning curve of Fed-CSLA-VGG and Fed-CSLA-GhostNet is similar to the learning of FedRepOpt-VGG and FedRepOpt-GhostNet networks respectively both in IID and Non-IID settings. We conjecture that the reason for such similar performance of FedRepOpt compared to Fed-CSLA is because FedRepOpt naturally follows the rules of initialization and iteration as outlined by [6] for training the centralized versions of FedRepOpt and Fed-CSLA.

5.6 Effects of using Momentum

It has been shown in [17] that SGD momentum can have adverse effects on FL training. One can observe in Table 2 that the plain-style model with large parameters like Fed-RepVGG-Inf is indeed highly affected by momentum compared to its multi-branch counterpart, i.e., Fed-RepVGG-Tr. We conjecture that the increase in momentum causes these large models in FL to quickly overfit on the local data resulting in a further increase in the client drift during model aggregation at the server. Interestingly, we found that FedRepOpt-VGG and FedRepOpt-GhostNet are less affected by momentum both in IID and non-IID settings. We conjecture the reason for this is due to the regularization effects provided by Grad Mult in the RepOpt.

5.7 Effects of Data Heterogeneity

Data heterogeneity is inherent in FL and it significantly affects the performance of the final global model. Table 3 shows the performance of VGG-style and

Table 2: Tiny ImageNet accuracy w/o momentum on IID and Non-IID setting with/without momentum

Model	w/ momentum		w/o momentum	
	Non-IID	IID	Non-IID	IID
Fed-RepGhost-Tr 0.5x	26.72 ± 0.35	29.10 ± 0.50	13.45 ± 0.48	14.24 ± 0.40
Fed-RepGhost-Inf 0.5x	26.07 ± 0.40	27.47 ± 0.33	13.10 ± 0.20	12.93 ± 0.40
Fed-CSLA-Ghost 0.5x	28.31 ± 0.33	31.61 ± 0.87	14.93 ± 0.29	16.71 ± 0.59
FedRepOpt-Ghost 0.5x	28.30 ± 0.30	30.57 ± 0.69	15.29 ± 0.85	16.19 ± 0.47
Fed-RepVGG-B1-Tr	38.76 ± 0.61	50.41 ± 0.32	35.30 ± 0.11	40.69 ± 0.33
Fed-RepVGG-B1-Inf	15.93 ± 1.45	15.79 ± 1.54	33.44 ± 0.11	37.79 ± 0.12
Fed-CSLA-VGG-B1	45.34 ± 0.24	51.80 ± 0.10	36.97 ± 0.22	42.09 ± 0.12
FedRepOpt-VGG-B1	45.33 ± 0.31	50.93 ± 0.70	37.27 ± 0.34	42.15 ± 0.03

Table 3: Accuracy on Tiny ImageNet (Non-IID). Higher values of α denote lower levels of heterogeneity.

Model	α				
	0.1	0.3	0.5	0.7	0.9
Fed-RepGhost-Tr 0.5x	13.45 ± 0.48	14.61 ± 0.47	14.52 ± 1.10	14.59 ± 0.59	15.39 ± 0.37
Fed-RepGhost-Inf 0.5x	13.10 ± 0.20	14.27 ± 0.23	13.63 ± 0.68	14.88 ± 0.22	15.04 ± 0.21
Fed-CSLA-Ghost 0.5x	14.93 ± 0.29	15.93 ± 0.81	16.37 ± 0.41	16.45 ± 0.65	16.51 ± 0.27
FedRepOpt-GhostNet 0.5x	15.29 ± 0.85	15.49 ± 0.58	15.64 ± 0.45	16.38 ± 0.71	16.44 ± 0.29
Fed-RepVGG-B1-Tr	35.30 ± 0.49	39.40 ± 0.42	39.30 ± 0.13	40.42 ± 0.60	40.66 ± 0.23
Fed-RepVGG-B1-Inf	33.44 ± 0.11	37.62 ± 0.56	37.27 ± 0.26	38.13 ± 0.36	38.21 ± 0.31
Fed-CSLA-VGG-B1	36.97 ± 0.22	40.59 ± 0.34	41.14 ± 0.19	41.43 ± 0.29	41.15 ± 0.49
FedRepOpt-VGG-B1	37.27 ± 0.34	40.81 ± 0.15	40.74 ± 0.33	40.93 ± 0.24	41.53 ± 0.27

GhostNet-style models with varying levels of data heterogeneity. One can see from Table 3, that the performance of the Fed-CSLA-based models and the FedRepOpt models with various levels of heterogeneity are nearly similar. Additionally, compared to the single-stream version (i.e., Fed-RepGhostNet-Inf and Fed-RepVGG-B1-Inf) of the Fed-Rep-based models, the FedRepOpt versions (i.e., FedRepOpt-GhostNet and FedRepOpt-VGG) of the Fed-CSLA models show better and more stable performance with increasing levels of data heterogeneity.

5.8 Effects of Local Epochs

The increase in the local epochs during FL training makes the model overfit on the local data distribution which can further exacerbate the issue of client drift [21] and reduce the final performance of the global model. One can see from Table 4 that increasing the local epochs deteriorates the performance of the large-scale models more severely compared to the small-scale models. We found that plain-style large models like Fed-RepVGG-Inf are largely affected by the increase in the local epochs compared to the Fed-RepVGG-Tr and Fed-CSLA-VGG models. Interestingly, we found that the performance of the FedRepOpt-based models is similar to Fed-CSLA-based models with increasing local epochs in FL. This further provides an advantage of FedRepOpt-based models over the plain-style models in FL settings.

Table 4: Tiny ImageNet accuracy on Non-IID setting with varying local epochs and rounds.

Model	E=1, R=240	E=5, R=48	E=10, R=24
Fed-RepGhost-Tr 0.5x	13.45 \pm 0.48	12.92 \pm 0.57	12.34 \pm 0.15
Fed-RepGhost-Inf 0.5x	13.10 \pm 0.20	12.68 \pm 0.29	11.76 \pm 0.27
Fed-CSLA-Ghost 0.5x	14.93 \pm 0.29	14.40 \pm 0.48	13.53 \pm 0.21
FedRepOpt-GhostNet 0.5x	15.29 \pm 0.85	14.15 \pm 0.19	13.69 \pm 0.53
Fed-RepVGG-B1-Tr	35.30 \pm 0.11	31.47 \pm 0.51	29.00 \pm 0.13
Fed-RepVGG-B1-Inf	33.44 \pm 0.11	26.82 \pm 0.10	23.53 \pm 0.32
Fed-CSLA-VGG-B1	36.97 \pm 0.22	34.38 \pm 0.52	32.53 \pm 0.48
FedRepOpt-VGG-B1	37.27 \pm 0.34	34.13 \pm 0.43	32.41 \pm 0.24

Table 5: Tiny ImageNet accuracy and training speed on cross-device Non-IID setting. TPR represents the training time per round.

Model	TPR (sec)	Top-1	Top-5
Fed-RepGhost-Tr 0.5x	44	18.54 \pm 0.44	43.40 \pm 0.46
Fed-RepGhost-Inf 0.5x	36	18.34 \pm 0.44	43.47 \pm 0.68
Fed-CSLA-Ghost 0.5x	46	21.07 \pm 0.41	47.10 \pm 0.47
FedRepOpt-RepGhost 0.5x	37	21.11 \pm 0.09	46.71 \pm 0.17
Fed-RepVGG-B1-Tr	105	42.15 \pm 0.53	69.20 \pm 0.74
Fed-RepVGG-B1-Inf	84	41.71 \pm 0.07	68.33 \pm 0.20
Fed-CSLA-VGG-B1	101	43.98 \pm 0.10	71.01 \pm 0.16
FedRepOpt-VGG-B1	94	43.36 \pm 0.73	70.64 \pm 0.45

5.9 FedRepOpt in Cross-Device Setting

Due to partial device participation along with heterogeneous data distribution, cross-device FL settings are more challenging than cross-silo FL settings. One can observe from Table 5 that FedRepOpt-GhostNet and FedRepOpt-VGG achieve an average Top-1 accuracy of 21.11% and 43.36% respectively which is nearly equivalent to their corresponding Fed-CSLA versions. Additionally, we find that FedRep-GhostNet and FedRep-VGG are 1.18 \times and 1.11 \times faster than Fed-RepGhost-Tr and FedRep-VGG-Tr respectively. Furthermore, FedRepOpt-GhostNet and FedRepOpt-VGG obtain 2.57% and 1.21% better performance compared to Fed-RepGhostNet-Inf and Fed-RepVGG-Inf respectively while achieving similar training speed. Interestingly, Figure 3 shows that the FedRepOpt-based models follow the learning curve of their respective Fed-CSLA-based models, and the Fed-Rep-Inf models follow the learning curve of Fed-Rep-Tr models.

5.10 Effects of the Client Participation in Cross-Device Setting

In practice, since only a fraction of the participants can be connected to the central server at a given time, the participation ratio of the clients is an important attribute in the cross-device setting. As shown in Table 6, FedRepOpt-RepGhost




Fig. 3: Learning behavior of CSLA-based models and their rep-optimized versions on cross-device Non-IID setting

Table 6: Tiny ImageNet top-1 accuracy on the cross-device Non-IID setting with a different number of aggregation clients. The search dataset in this case was CIFAR100.

Model	50%	10%	5%
Fed-RepGhost-Tr 0.5x	20.76 ± 0.60	18.54 ± 0.44	15.33 ± 0.81
Fed-RepGhost-Inf 0.5x	21.62 ± 0.08	18.34 ± 0.44	14.19 ± 1.16
Fed-CSLA-Ghost 0.5x	24.29 ± 1.25	21.07 ± 0.41	18.01 ± 0.22
FedRepOpt-RepGhost 0.5x	24.16 ± 0.15	21.11 ± 0.09	18.35 ± 0.15

suffers less degradation in accuracy compared to the conventional Fed-RpGhost-Inf method (i.e., 5.43% vs 7.43%) when the participation ratio is decreased from 50% to 5%.

5.11 Effects of the Search Dataset

As mentioned in [6], the modified optimizers (called RepOptimizers) for RepOpt-based architectures require scalar multipliers that are multiplied with the gradients before updating the parameters of the model. These scalars are obtained by training the CSLA versions of the RepOpt-based models on a small *Hyperparameter Search* (HS) dataset. The authors in [6] state that RepOptimizers are architecture-specific but agnostic to the type of HS datasets. We went one step further and evaluated the effects of the HS dataset being IID and Non-IID as shown in Table 7 and the effect of different search datasets as shown in Table 8. For the first purpose, we chose the full (50,000 training samples with 100 classes), three different IID subsets (4900 training samples with 100 classes for each), and three different Non-IID subsets (5500 samples with 45 classes, 6000 samples with 46 classes and 5000 samples with 59 classes) versions of CIFAR-100 as the HS dataset. For the second objective, we chose the complete Fashion-MNIST [33], iNaturalist [29] and CIFAR100 as the HS datasets for comparison.

As shown in Table 7, one can observe that in all cases the FedRepOpt results in similar performance with the full, IID, and Non-IID versions of the HS datasets, which further validates the claim of [6] even in FL. Furthermore, as shown in Table 8, the hyper-parameters searched on Fashion-MNIST and iNaturalist demonstrated comparable accuracy to those searched on CIFAR100. Such results bear obvious advantages, for instance, in cases where it is hard to obtain

Table 7: Ablation study on the effects of search dataset. We chose three configurations of CIFAR-100 as a search dataset, i.e., Full (complete dataset), iid, and non-iid subsets. The accuracy is reported on Tiny-ImageNet.

Model	Param (M)	Full set	IID subset	NIID subset	No. of classes/client	HS Samples/client	Top-1 Acc.		
							Non-IID	IID	
FedRepOpt-GhostNet 0.5x	1.284	✓	✗	✗	100	50K	15.29 / 16.19		
		✗	✓	✗	100 / 100 / 100	4.9K / 4.9K / 4.9K	14.81 / 15.12 / 15.16	15.74 / 15.37 / 15.51	
		✗	✗	✓	45 / 46 / 59	5.5K / 6.0K / 5.0K	15.28 / 15.70 / 15.37	16.00 / 15.18 / 15.06	
FedRepOpt-VGG	50.2	✓	✗	✗	100	50K	37.27 / 42.15		
		✗	✓	✗	100 / 100 / 100	4.9K / 4.9K / 4.9K	37.12 / 37.41 / 36.65	41.67 / 41.73 / 42.11	
		✗	✗	✓	45 / 46 / 59	5.5K / 6.0K / 5.0K	36.56 / 37.04 / 36.80	41.27 / 41.97 / 40.97	

Table 8: Tiny ImageNet accuracy on Non-IID and IID setting with different hyperparameter search datasets including CIFAR100, Fashion-MNIST, and iNaturalist.

Model	Non-IID			IID		
	CIFAR100	Fashion-MNIST	iNaturalist	CIFAR100	Fashion-MNIST	iNaturalist
FedRepOpt-VGG-B1	37.27 ± 0.34	35.81 ± 0.06	37.78 ± 0.10	42.15 ± 0.03	41.54 ± 0.58	43.37 ± 0.34
FedRepOpt-GhostNet	15.29 ± 0.85	15.53 ± 0.91	14.93 ± 0.35	16.19 ± 0.47	14.67 ± 0.69	15.04 ± 0.53

the HS dataset due to data privacy concerns. Since the HS dataset only serves to find the *model-specific* hyperparameters at the initial round of FL, one can also obtain the *model-specific* hyperparameters by first asking one of the clients with large computational resources to train the CSLA-model locally first. The *model-specific* hyperparameter obtained from the local client can then be utilized to perform the FL training on the rest of the clients. Instead of conducting the hyperparameters search on the client side, the central server can utilize the public datasets on the Internet to obtain the parameters in the initial stage and distribute them to the local clients during the FL training.

6 Conclusion

In this work, we presented FedRepOpt, a novel federated learning approach that addresses the computational constraints of client devices in the training process. By leveraging gradient re-parameterization and RepOptimizers, FedRepOpt enables the deployment of complex multi-branch models on edge devices without compromising performance. Our experiments with VGG and GhostNet architectures have demonstrated that FedRepOpt achieves a significant boost in performance of 16.7% and 11.4% compared to the RepGhost-style and RepVGG-style networks. We also obtained faster convergence time of 11.7% and 57.4% compared to their complex structure. Moreover, FedRepOpt allows for collaborative on-device training, ensuring the privacy of user data while harnessing the full potential of distributed learning. The proposed approach opens up new possibilities for deploying advanced deep-learning models in IoT applications.

References

1. Acar, D.A.E., Zhao, Y., Navarro, R.M., Mattina, M., Whatmough, P.N., Saligrama, V.: Federated learning based on dynamic regularization. arXiv preprint arXiv:2111.04263 (2021)
2. Alistarh, D., Grubic, D., Li, J., Tomioka, R., Vojnovic, M.: Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in neural information processing systems* **30** (2017)
3. Beutel, D.J., Topal, T., Mathur, A., Qiu, X., Parcollet, T., Lane, N.D.: Flower: A friendly federated learning research framework. arXiv preprint arXiv:2007.14390 (2020)
4. Chen, C., Xu, H., Wang, W., Li, B., Li, B., Chen, L., Zhang, G.: Communication-efficient federated learning with adaptive parameter freezing. In: 2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS). pp. 1–11. IEEE (2021)
5. Chen, C., Guo, Z., Zeng, H., Xiong, P., Dong, J.: Repghost: A hardware-efficient ghost module via re-parameterization. arXiv e-prints pp. arXiv-2211 (2022)
6. Ding, X., Chen, H., Zhang, X., Huang, K., Han, J., Ding, G.: Re-parameterizing your optimizers rather than architectures. In: The Eleventh International Conference on Learning Representations (2023), https://openreview.net/forum?id=B92TMCg_7rp
7. Ding, X., Zhang, X., Han, J., Ding, G.: Diverse branch block: Building a convolution as an inception-like unit. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10886–10895 (2021)
8. Ding, X., Zhang, X., Han, J., Ding, G.: Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11963–11975 (2022)
9. Ding, X., Zhang, X., Ma, N., Han, J., Ding, G., Sun, J.: Repvgg: Making vgg-style convnets great again. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13733–13742 (2021)
10. Dryden, N., Moon, T., Jacobs, S.A., Van Essen, B.: Communication quantization for data-parallel training of deep neural networks. In: 2016 2nd Workshop on Machine Learning in HPC Environments (MLHPC). pp. 1–8. IEEE (2016)
11. Gao, Y., Parcollet, T., Zaiem, S., Fernandez-Marques, J., de Gusmao, P.P., Beutel, D.J., Lane, N.D.: End-to-end speech recognition from federated acoustic models. In: ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 7227–7231. IEEE (2022)
12. Hsieh, K., Harlap, A., Vijaykumar, N., Konomis, D., Ganger, G.R., Gibbons, P.B., Mutlu, O.: Gaia: {Geo-Distributed} machine learning approaching {LAN} speeds. In: 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17). pp. 629–647 (2017)
13. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S., Stich, S., Suresh, A.T.: Scaffold: Stochastic controlled averaging for federated learning. In: International conference on machine learning. pp. 5132–5143. PMLR (2020)
14. Konečný, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T., Bacon, D.: Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492 (2016)
15. Li, X., Jiang, M., Zhang, X., Kamp, M., Dou, Q.: Fedbn: Federated learning on non-iid features via local batch normalization. arXiv preprint arXiv:2102.07623 (2021)

16. Li, Y., Tao, X., Zhang, X., Liu, J., Xu, J.: Privacy-preserved federated learning for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems* **23**(7), 8423–8434 (2021)
17. Liu, W., Chen, L., Chen, Y., Zhang, W.: Accelerating federated learning via momentum gradient descent. *IEEE Transactions on Parallel and Distributed Systems* **31**(8), 1754–1766 (2020)
18. Liu, Y., Nie, J., Li, X., Ahmed, S.H., Lim, W.Y.B., Miao, C.: Federated learning in the sky: Aerial-ground air quality sensing framework with uav swarms. *IEEE Internet of Things Journal* **8**(12), 9827–9837 (2020)
19. Long, G., Tan, Y., Jiang, J., Zhang, C.: Federated learning for open banking. In: *Federated Learning: Privacy and Incentive*, pp. 240–254. Springer (2020)
20. Luping, W., Wei, W., Bo, L.: Cmf1: Mitigating communication overhead for federated learning. In: *2019 IEEE 39th international conference on distributed computing systems (ICDCS)*. pp. 954–964. IEEE (2019)
21. 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. PMLR (2017)
22. Nguyen, D.C., Ding, M., Pathirana, P.N., Seneviratne, A., Li, J., Poor, H.V.: Federated learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials* **23**(3), 1622–1658 (2021)
23. Qayyum, A., Ahmad, K., Ahsan, M.A., Al-Fuqaha, A., Qadir, J.: Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge. *IEEE Open Journal of the Computer Society* **3**, 172–184 (2022)
24. Reddi, S.J., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., McMahan, H.B.: Adaptive federated optimization. In: *International Conference on Learning Representations (2021)*, <https://openreview.net/forum?id=LkFG31B13U5>
25. Rehman, Y.A.U., Gao, Y., de Gusmao, P.P.B., Alibeigi, M., Shen, J., Lane, N.D.: L-dawa: Layer-wise divergence aware weight aggregation in federated self-supervised visual representation learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 16464–16473 (2023)
26. Seide, F., Fu, H., Droppo, J., Li, G., Yu, D.: 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In: *Fifteenth annual conference of the international speech communication association (2014)*
27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
28. Ström, N.: Scalable distributed dnn training using commodity gpu cloud computing (2015)
29. 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*. pp. 8769–8778 (2018)
30. Wang, H., Yurochkin, M., Sun, Y., Papailiopoulos, D., Khazaeni, Y.: Federated learning with matched averaging. *arXiv preprint arXiv:2002.06440* (2020)
31. Wang, Y., Shi, Q., Chang, T.H.: Why batch normalization damage federated learning on non-iid data? *IEEE Transactions on Neural Networks and Learning Systems* (2023)
32. Xianjia, Y., Queralta, J.P., Heikkonen, J., Westerlund, T.: Federated learning in robotic and autonomous systems. *Procedia Computer Science* **191**, 135–142 (2021)
33. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017)

34. Zhao, Y., Li, M., Lai, L., Suda, N., Civan, D., Chandra, V.: Federated learning with non-iid data. arXiv preprint arXiv:1806.00582 (2018)
35. Zhuang, W., Gan, X., Wen, Y., Zhang, S., Yi, S.: Collaborative unsupervised visual representation learning from decentralized data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4912–4921 (2021)