



HAL
open science

FLoCoRA: FEDERATED LEARNING COMPRESSION WITH LOW-RANK ADAPTATION

Lucas Grativol, Mathieu Leonardon, Guillaume Muller, Virginie Fresse,
Matthieu Arzel

► **To cite this version:**

Lucas Grativol, Mathieu Leonardon, Guillaume Muller, Virginie Fresse, Matthieu Arzel. FLoCoRA: FEDERATED LEARNING COMPRESSION WITH LOW-RANK ADAPTATION. 32nd European Signal Processing Conference EUSIPCO, Aug 2024, Lyon, France. hal-04617632

HAL Id: hal-04617632

<https://hal.science/hal-04617632v1>

Submitted on 19 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

FLoCoRA: FEDERATED LEARNING COMPRESSION WITH LOW-RANK ADAPTATION

Lucas Grativol*, Mathieu Léonardon*, Guillaume Muller[‡], Virginie Fresse[†] and Matthieu Arzel*

*IMT Atlantique, Lab-STICC, UMR CNRS 6285, F-29238 Brest, France

[†]Hubert Curien Laboratory, Saint-Etienne, France

[‡]Mines Saint-Etienne, Institut Henri Fayol, Saint-Etienne, France

Abstract—Low-Rank Adaptation (LoRA) methods have gained popularity in efficient parameter fine-tuning of models containing hundreds of billions of parameters. In this work, instead, we demonstrate the application of LoRA methods to train small-vision models in Federated Learning (FL) from scratch. We first propose an aggregation-agnostic method to integrate LoRA within FL, named FLoCoRA, showing that the method is capable of reducing communication costs by 4.8 times, while having less than 1% accuracy degradation, for a CIFAR-10 classification task with a ResNet-8. Next, we show that the same method can be extended with an affine quantization scheme, dividing the communication cost by 18.6 times, while comparing it with the standard method, with still less than 1% of accuracy loss, tested with on a ResNet-18 model. Our formulation represents a strong baseline for message size reduction, even when compared to conventional model compression works, while also reducing the training memory requirements due to the low-rank adaptation.

Index Terms—Low-Rank Adaptation, Federated Learning, Compression

I. INTRODUCTION

The preservation of data privacy has long been a goal in the development of approaches to train machine learning models. In traditional machine learning, raw data from embedded systems are sent over a network to a powerful server for model training, raising concerns about confidentiality [1]. Federated Learning (FL) has emerged as a promising mitigation of these problems [2]. In FL, each participating client keeps its data locally, only sharing the results of its local training. These results, often in the form of model parameters or gradients, are fused in a central orchestration server. As the training process shifts from a server to a third-party device, there is a trade-off between increased privacy and computational and communication overhead. The challenge of convergence in FL [3] further supports this idea, as the diverse data distribution among devices leads to conflicting update models, which affects the convergence time.

To overcome the computation and communication problem in FL, previous works have focused on model compression techniques such as pruning [4], quantization [5] and more recently the application of low-rank adaptation (LoRA) with [6]–[8], following the original work on LoRA [9].

This work is supported by the *Futur et Ruptures* program funded by IMT and Institut Carnot TSN, and by the GdR IASIS.

Unlike previous works, that focused on large models, we demonstrate the application of LoRA for training small-vision models from scratch. Going further, we demonstrate how quantization can be applied to LoRA and FL, dividing the cost of communication by 18.6, compared to the standard method, while keeping the accuracy degradation within 1%.

We can summarize our contributions as :

- We demonstrate how LoRA adapters can be integrated to the FL framework, while still being implementable in any FL optimization method, proposing FLoCoRA. Representing a strong baseline for communication and memory reduction in FL. Our code is publicly accessible ¹.
- We study the impact of LoRA hyperparameters for a classification task with small vision models, which allows these models to be trained from scratch, contributing to reduce message sizes in FL up to 4.8 times, with 1% loss of accuracy for a ResNet-8.
- We introduce an affine quantization scheme with FLoCoRA, allowing further compression rates of 18.6 to 37.3 times with up to 1% loss of accuracy for a ResNet-18.

II. BACKGROUND

A. Federated Learning

As the most used aggregation method in FL, we use FedAvg (Federated Averaging) [2] as the showcase for a typical FL framework, while also presenting the baseline used in this work. It operates by sampling a subset K , of a pool of clients C , in each round to train a model with parameters w for a specified number of local epochs. Each client has its own dataset of size n_i , and the total size of the datasets of the participating clients in a round is denoted n . Each client $k \in K$, seeks to find the parameters w that minimize its local loss $f_k(w) = \mathbb{E}[l(X, y, w)]$, for a set of examples X and labels y , and the loss function $l(\cdot)$. The general objective of FedAvg is then to find a global w that minimizes Equation 1, where $F_k(w) = \mathbb{E}[f_k]$. When this process is iterating for a certain number of rounds R , the final result is expected to increase the individual performance of each client [10], without the need for data sharing.

¹https://anonymous.4open.science/r/FLoCoRA_eusipco24-A6F1/

$$\min_w f(w) = \sum_{k=1}^K \frac{n_i}{n} F_k(w) \quad (1)$$

During training, model parameters are downloaded from server to client and uploaded from client to server. Equation 2 represents the total communication cost (TCC) for a certain number of rounds R . Here, we consider Q_p as the number of bits of each element in w and $|w|$ as the total number of elements/parameters. To decrease the size of a message, one can reduce the number of parameters or reduce the value of Q_p . Equation 2 can be utilized to calculate the communication cost for a single client in FL. Therefore, we employ this equation throughout this work to determine the TCC.

$$TCC(R) = 2RQ_p|w| \quad (2)$$

B. Model Compression

This is the point where traditional model compression techniques help the FL framework. Model compression is a widely adopted solution [11] to reduce computational and memory requirements in centralised machine learning. As such, the ideas of pruning [12] and quantization [5], two widely adopted techniques in model compression, have been studied in the FL scenario [3].

Within the FL framework, pruning can be applied to induce more sparsity in a model, leading to more compressible parameters [4], or to eliminate certain architectural elements, such as kernel or filters, to transmit fewer parameters [13]. By introducing sparsity or removing certain parts of a model, pruning can also reduce the computation needs of a model during training [14].

Another method, sometimes complementary to pruning [15], is quantization, where the objective is to reduce the binary representation of the parameters. As usually models are represented using 32- or 16-bits floating point (FP) numbers, the idea is to reduce it to 8- or sub-8-bit FP, or even change the data to fixed-point or integer formats [16].

Pruning has been used successfully in FL, either post-training to introduce sparsity and reduce message size, or during training to reduce computations and message size [12]. However, such methods require additional steps to induce sparsity, resulting in overhead.

C. Low-Rank Adaptation

In the domain of transfer learning, low-rank adaptation methods have gained popularity for task adaptation of model containing hundreds of billions of parameters [9], [17]. LoRA comes from the need to adapt a large-scale already trained model to a new task, in order to save computation time and reduce the memory footprint. For example, the parameters, $W_l \in \mathbb{R}^{d_1 \times d_2}$, of pretrained linear layer, l , are reformulated with a new adapter, a parallel layer, obtaining $W_l^* = W_l + \frac{\alpha}{r}BA$, where $A \in \mathbb{R}^{d_1 \times r}$ and $B \in \mathbb{R}^{r \times d_2}$ are two matrices of maximum rank r . The term $\frac{\alpha}{r}$ is a scaling factor, based on a hyperparameter α and rank r . The idea is to keep W_l

frozen and to train BA to represent a lower-rank update of W_l . An advantage of this approach is that the matrix BA can be incorporated back into the original pretrained weights, W_l^* , without any additional latency. LoRA has shown that for $r \ll \min(d_1, d_2)$, W_l can be adapted to a new task, largely reducing the need to fine-tune the entire model.

When LoRA is applied to FL, the original model is kept frozen and we only need to train the adapters, B and A . The adapters have fewer parameters than the original model, resulting in a lighter model for training because they require less memory for the gradients. Since the original model remains frozen, it is not necessary to communicate it between the server and client in every round, thus the communication is also lighter.

Previous studies like SLoRA [7] have investigated the potential impact of LoRA techniques on FL. They suggested merging standard FL training with matrix decomposition to achieve favorable initialization for both the original model and matrices A and B . HLoRA [8] allowed clients to select different LoRA ranks, according to their individual constraints, to adapt a pretrained base model. FedPara [6], redefined the LoRA adapter as $W_l^* = W_l + (X_1 \cdot Y_1^T) \odot (X_2 \cdot Y_2^T)$, where \odot is the Hadamard product, in order to obtain an update of higher rank than conventional LoRA.

However, SLoRA and HLoRA have not investigated the application of LoRA methods for the training of small CNNs from scratch, having focused on how LoRA was applied to bigger models like large-language models (LLM) and foundation models. Closer to our work, FedPara [6] has tested on small CNNs, but its low-rank method and quantization is limited to FP-16 and the adaptation of FedPq [5] on a rather simple test scenario.

III. OUR METHOD - FLoCoRA

In this work, we propose never to update the parameters of the randomly initialized neural network parameters shared between clients at the start of training. Only adapters parameters will be trained, exchanged, and updated. Since the original parameters remain unchanged, the exchange of LoRA parameters is sufficient to represent the updates from each client. The server continues to receive updated parameters from clients, which means that this method can also be integrated with other FL techniques, without further changes.

Figure 1 illustrates our approach, **Federated Learning Compression with Low-Rank Adaptation (FLoCoRA)**, for a single communication round. All clients start with identical weights, denoted as $W_{initial}$, which remain fixed throughout the training process. In the first step (1), the server transmits the global LoRA adapters parameters, $\bar{\Delta}_t L$, to the selected subset K of clients. Subsequently, in step (2), each client independently trains its LoRA adapter locally and uploads the result, $\Delta_{t+1}^k L$, to the server in step (3). Finally, in step (4), the server uses the same weighted averaging mechanism as FedAvg to obtain the updated global LoRA adapter parameters for the next round, denoted $\bar{\Delta}_{t+1} L$.

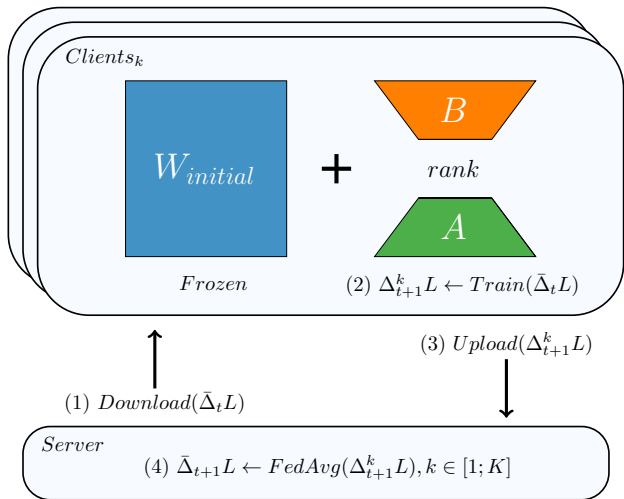


Fig. 1: FLoCoRA training loop.

TABLE I: Number of parameters for different sizes of r . For each value, we have the total number of parameters to be trained/sent and the total number of the parameters with the original model plus the LoRA adapter.

Method	Total Params	Trained Params	% of Trained Params
FedAvg	1.23M	1.23M	100
FLoCoRA ($r = 8$)	1.30M	69.45K	5.35
FLoCoRA ($r = 16$)	1.36M	131.92K	9.70
FLoCoRA ($r = 32$)	1.48M	256.84K	17.30
FLoCoRA ($r = 64$)	1.73M	506.70K	29.22
FLoCoRA ($r = 128$)	2.23M	1.00 M	45.05

Table I presents the size of the trainable parameters for a ResNet-8 [18], for different values of r . It should be noted that we adapt the convolution layers using LoRA, but the normalization and fully connected (FC) layers of the original model are trained; more details of this are discussed in Section IV.

For convolution layers, we follow the decomposition proposed in [19]. Let $P_l \in \mathbb{R}^{O \times I \times K_r \times K_r}$ be a convolution layer; then we define its LoRA adapter matrices as $B \in \mathbb{R}^{r \times I \times K_r \times K_r}$ and $A \in \mathbb{R}^{O \times r \times 1 \times 1}$. Output channels are denoted as O , the input channels as I , and the kernel size as K_r .

To verify the impacts of using LoRA with FL, we first investigate which layers of the original model need to be trained in a normal way and which can receive the LoRA adapters. Then, we compare our work with recent literature on compression methods for FL, showing that LoRA methods are a strong baseline for future works in the domain. Finally, we show that quantizing LoRA parameters can lead to additional message compression, even when dealing with highly diverse data distributions across clients.

IV. EXPERIMENTS AND DISCUSSION

Our basic setup for FL consists of 100 clients, where at each round 10 % of clients are sampled. We train for a total of 100 rounds. For simplicity, based on [6], we fix the batch size, the

TABLE II: The effect of training different layers with or without LoRA adapters.

Method	Nb. of Params. to update	Accuracy
FedAvg	1.23 M	76.14 ± 0.74
FLoCoRA Vanilla	0.26 M	22.14 ± 3.99
+ Norm. layers	0.26 M	39.80 ± 12.05
+ Final FC	0.26 M	75.51 ± 1.34

learning rate, the number of local epochs and momentum at 32, 0.01, 5 and 0.9, respectively, for all clients. We use SGD with momentum as the local optimizer for clients. We use FedAvg as our aggregation algorithm. We ran all the experiments three times with different seeds.

Initially, to verify how FLoCoRA training impacts convergence in FL, we trained a ResNet-8 for an image classification task with the CIFAR-10 dataset. Following an LDA distribution with parameter 0.5 [2], we also replace the batch normalization layer with the group normalization layer as suggested by [20]. Table II presents an ablation of the layers to be trained with LoRA hyperparameters $r = 32$ and $\alpha = 512$. We chose to use $r = 32$ because it was the minimum rank of which the accuracy had less than 1% of degradation. Starting with a randomly initialized ResNet-8 model, we freeze the entire model, introducing LoRA adapters for all convolutions and to the final FC layer, corresponding to "FLoCoRA Vanilla" in Table II. Subsequently, we unfreeze the normalization layers, "+ Norm. Layers", then we remove the LoRA adapter from the final FC layer and unfreeze it, "+ Final FC".

Normalization layers need to be trained to capture running statistics, as it cannot be adapted with LoRA methods, so we need to train it. For the FC layer, we hypothesize that the lower rank is not enough for the last classification layer, as it is a layer that tends to be highly specialized and sensible [21]. For the remainder of this paper, we keep this configuration for FLoCoRA experiments.

Expanding on the importance of choosing the appropriate




Fig. 2: The relationship between the r hyperparameter in FLoCoRA and the scaling factor α , in $\frac{\alpha}{r}$. Two scenarios are evaluated, $\alpha = 2r$ and $\alpha = 16r$, against FedAvg.

layers, we examine the trade-off between rank r and the scaling parameter α . We compare a FedAvg baseline with two different setups, where α is set to $2r$ and $16r$, the results are shown in Figure 2. In LoRA’s original paper [9], α is scaled to twice the rank for LLMs. However, when training small CNNs with FL, we found that increasing this factor further led to improvements in accuracy of up to 4.4%. Increasing the scaling factor of the LoRA adapters layers is a way to adjust their learning rate to higher values while keeping more sensible layers, the normalization and the FC layers, in a lower learning rate, improving training stability and performance.

Looking at the results in Figure 2, FLoCoRA is able to achieve an accuracy drop of less than 1%, while sharing only 0.26M parameters, representing a reduction of $4.8\times$ compared to sharing the entire model. For a rank of 128, the number of parameters to send is closer to the original model, as seen in Table I, but we see an improvement in accuracy of 2%. Indeed, for the ResNet-8 design used, the convolution layers throughout the model have output channel numbers of 64, 128, or 256. Consequently, opting for a rank of 128 results in an augmentation of rank for certain layers while maintaining the same rank for others. This modification results in a greater overall rank update for the shallower convolution layers compared to the original model, and the low-rank update is actually applied to the deeper convolution layers. Finally, the model has slightly less parameters because the larger convolution layers, with 256 output channels, are adapted with a lower rank of 128, effectively reducing the number of parameters.

So far FLoCoRA has reduced the number of trainable parameters, which reduces the necessary memory to train the model, and consequently reducing the number of parameters that must be communicated in each round. Next, we investigate the effects of applying an affine quantization scheme [22] for both the client and the server message. We calculate the scaling factor and zero point values per channel for the convolution layers and per column for the FC layer. Normalization layers are not quantized. We use 2/4/8-bit formats to quantify trainable layers. For a model such as ResNet-8 this translates to reductions of 56.3, 32.6 and 17.7 times in the TCC, respectively, for 2/4/8-bits. We included the overhead to transmit the scaling factors and zero points in FP format. Table III resumes the TTC, as expressed by Equation 2, to train a model with FLoCoRA and quantization.

We show in Figure 3 the evolution of the accuracy of the FP FedAvg and FLoCoRA, and also the quantized version of

TABLE III: Total communication cost (TTC) for different quantization levels with FLoCoRA, for a rank of 32 and alpha of 512, during 100 rounds of FL.

Method	Quantization	TCC	Accuracy
FedAvg	FP	982.07 MB ($\div 1$)	76.14 \pm 0.74
FLoCoRA	FP	205.47 MB ($\div 4.8$)	75.51 \pm 1.34
	int8	55.56 MB ($\div 17.7$)	74.21 \pm 1.05
	int4	30.15 MB ($\div 32.6$)	73.15 \pm 0.18
	int2	17.44 MB ($\div 56.3$)	55.03 \pm 1.90




Fig. 3: Convergence behavior between FedAvg, FLoCoRA with rank of 32 and its quantized versions of 2/4/8-bits.

FLoCoRA. The convergence time for FLoCoRA in FP and its quantized version with 8-bits is not affected, demonstrating that it is an effective method for communication savings. As per Table III the quantized versions suffer further degradation, with 2% for the 8-bit case. Even with close convergence times, we see that the quantized versions added more instability to the training, this could be a further interesting research direction to improve the quantization scheme with FLoCoRA. We would expect that revisiting the literature on model compression [23], [24] and combining it with FLoCoRA would find a better quantization scheme. Here, we demonstrate that even a simple technique, based on a round-to-nearest method, is capable of achieving competitive results.

FedPara [6] method has been applied to a VGG-16 model, which is compared to FedAvg and a low-rank tensor parameterization based on Tucker decomposition, resulting in communication reductions of 2.8-10.1 times. Base FLoCoRA achieves comparable compression ratios, showing that combining the low-rank adapter for convolution proposed in [19] is capable of training a small CNN model from scratch. The quantized version of FLoCoRA further reduces communication cost, resulting in a better trade-off of communication and accuracy.

TABLE IV: Comparing LoRA and quantization to ZeroFL and Magnitude Pruning methods.

Method	Config.	Message Size (MB)	TCC (GB)	Accuracy
FedAvg	Full Model	44.7($\div 1.0$)	62.6	84.43 \pm 0.36
ZeroFL [12]	90% SP+ 0.2 MR	27.3($\div 1.6$)	38.2	81.04 \pm 0.28
	90% SP+ 0.0 MR	10.1($\div 4.4$)	14.1	73.87 \pm 0.50
Magnitude Pruning [4]	40% prune	27.1 ($\div 1.6$)	38.0	85.20 \pm 0.20
	80% prune	9.8 ($\div 4.6$)	13.7	80.70 \pm 0.24
FLoCoRA	r=64	9.2 ($\div 4.9$)	12.9	85.17 \pm 0.44
	r=32	4.6 ($\div 9.7$)	6.5	83.90 \pm 0.20
	r=16	2.4 ($\div 18.6$)	3.3	82.33 \pm 0.35
	r=64, Q=8	2.4 ($\div 18.6$)	3.3	85.24 \pm 0.23
	r=32, Q=8	1.2 ($\div 37.3$)	1.7	83.95 \pm 0.32
	r=16, Q=8	0.7 ($\div 63.9$)	1.0	81.89 \pm 1.01

Finally, combining FLoCoRA with the suggested quantization approach, we compare it with the techniques introduced in ZeroFL [12] and with Magnitude Pruning [4]. Table IV illustrates that, although used as a fine-tuning method, LoRA serves as a strong baseline for communication savings in FL. For this experiment, we reproduced the setup proposed by [4], [12], where we have 100 clients, training for 1 local epoch a ResNet-18 model, with an LDA parameter of 1.0, for 700 communication rounds.

ZeroFL reports their baseline accuracy as 80.62%, in our work we use the baseline as in Table III. The difference comes from the batch size used for the clients in the two works. As the client batch size is not reported by ZeroFL, we used the values indicated by Magnitude Pruning. FLoCoRA achieves less accuracy degradation for smaller messages, compared to conventional methods. The same is observed when FLoCoRA is combined with quantization, a reduction between 18.6-63.9 times with 4% accuracy loss for a case where we spend 0.7 MB per client to train a model with 44.7 MB. Interestingly, compared to our previous experience in Table III, the quantized versions for ResNet-18 show less degradation than before. We see this to be possible because we have a model that is $9\times$ larger, being trained $7\times$ longer, with an easy training scenario. The first results use an LDA of 0.5 and for ResNet-18 an LDA 1.0, where the higher the LDA parameter, the more identical and close the data distribution between the clients becomes.

V. CONCLUSION

In this work, we showed that Low Rank Adaption (LoRA), commonly used for fine-tuning, can be used to train small vision models from scratch in Federated Learning. It allows to decrease the training complexity of previous methods, reducing the amount of communication between server and clients, as well as the training memory requirements and computational power. Moreover, we showed that our method, FLoCoRA, can be deployed with quantization, allowing to further reduce the communication cost, while maintaining strong accuracies. In our setting, using ResNet18 networks on the CIFAR10 dataset, we showed that communication cost is divided by 18.6, with less than 1% accuracy loss. FLoCoRA paves the way for the deployment of Federated Learning on edge devices, where communication is a bottleneck, and computational power is limited. It also raises some questions to be explored in future work. Indeed, our method is able to train small models without updating the original one. This opens the door to the exploration of model architecture heterogeneity between clients, by using low-rank methods. It would also be possible to explore quantization and rank heterogeneity to further reduce the communication cost.

REFERENCES

- [1] Y. Sun, J. Zhang, Y. Xiong, and G. Zhu, "Data security and privacy in cloud computing," *International Journal of Distributed Sensor Networks*, vol. 10, no. 7, p. 190903, 2014.
- [2] B. McMahan and et al, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [3] P. Kairouz and et al, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [4] L. Grativol, M. Léonardon, G. Muller, V. Fresse, and M. Arzel, "Federated learning compression designed for lightweight communications," in *2023 30th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*. IEEE, 2023, pp. 1–4.
- [5] A. Reiszadeh, A. Mokhtari, H. Hassani, A. Jadbabaie, and R. Pedarsani, "Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2021–2031.
- [6] N. Hyeon-Woo, M. Ye-Bin, and T.-H. Oh, "Fedpara: Low-rank hadamard product for communication-efficient federated learning," in *International Conference on Learning Representations*, 2021.
- [7] S. Babakniya, A. R. Elkordy, Y. H. Ezzeldin, Q. Liu, K.-B. Song, M. El-Khamy, and S. Avestimehr, "Slora: Federated parameter efficient fine-tuning of language models," *arXiv preprint arXiv:2308.06522*, 2023.
- [8] Y. J. Cho, L. Liu, Z. Xu, A. Fahrezi, M. Barnes, and G. Joshi, "Heterogeneous lora for federated fine-tuning of on-device foundation models," in *International Workshop on Federated Learning in the Age of Foundation Models in Conjunction with NeurIPS 2023*, 2023.
- [9] E. J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen et al., "Lora: Low-rank adaptation of large language models," in *International Conference on Learning Representations*, 2021.
- [10] A. Z. Tan, H. Yu, L. Cui, and Q. Yang, "Towards personalized federated learning," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [11] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "A survey of model compression and acceleration for deep neural networks," *arXiv:1710.09282*, 2017.
- [12] X. Qiu, J. Fernandez-Marques, P. P. Gusmao, Y. Gao, T. Parcollet, and N. D. Lane, "Zerofl: Efficient on-device training for federated learning with local sparsity," *arXiv preprint arXiv:2208.02507*, 2022.
- [13] Y. Jiang, S. Wang, V. Valls, B. J. Ko, W.-H. Lee, K. K. Leung, and L. Tassiulas, "Model pruning enables efficient federated learning on edge devices," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [14] H. Tessier, V. Gripon, M. Léonardon, M. Arzel, T. Hannagan, and D. Bertrand, "Rethinking weight decay for efficient neural network pruning," *Journal of Imaging*, vol. 8, no. 3, p. 64, 2022.
- [15] S. Bai, J. Chen, X. Shen, Y. Qian, and Y. Liu, "Unified data-free compression: Pruning and quantization without fine-tuning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 5876–5885.
- [16] S. Wu, G. Li, F. Chen, and L. Shi, "Training and inference with integers in deep neural networks," *arXiv preprint arXiv:1802.04680*, 2018.
- [17] R. Bensaïd, V. Gripon, F. Leduc-Primeau, L. Mauch, G. B. Hacene, and F. Cardinaux, "A novel benchmark for few-shot semantic segmentation in the era of foundation models," *arXiv preprint arXiv:2401.11311*, 2024.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [19] M. Huh, H. Mobahi, R. Zhang, B. Cheung, P. Agrawal, and P. Isola, "The low-rank simplicity bias in deep networks," *Transactions on Machine Learning Research*, 2022.
- [20] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," *arXiv:1909.06335*, 2019.
- [21] B. Neyshabur, H. Sedghi, and C. Zhang, "What is being transferred in transfer learning?" *Advances in neural information processing systems*, vol. 33, pp. 512–523, 2020.
- [22] M. Nagel, M. Fournarakis, R. A. Amjad, Y. Bondarenko, M. Van Baalen, and T. Blankevoort, "A white paper on neural network quantization," *arXiv preprint arXiv:2106.08295*, 2021.
- [23] M. Nagel, M. Fournarakis, Y. Bondarenko, and T. Blankevoort, "Overcoming oscillations in quantization-aware training," in *International Conference on Machine Learning*. PMLR, 2022, pp. 16 318–16 330.
- [24] S. K. Esser, J. L. McKinstry, D. Bablani, R. Appuswamy, and D. S. Modha, "Learned step size quantization," in *International Conference on Learning Representations*, 2019.