
Efficient and Private Federated Learning with Partially Trainable Networks

Hakim Sidahmed, Zheng Xu, Ankush Garg, Yuan Cao, Mingqing Chen
Google
{hsidahmed, xuzheng, ankugarg, yuanco, mingqing}@google.com

Abstract

Federated learning is used for decentralized training of machine learning models on a large number (millions) of edge mobile devices. It is challenging because mobile devices often have limited communication bandwidth and local computation resources. Therefore, improving the efficiency of federated learning is critical for scalability and usability. In this paper, we propose to leverage partially trainable neural networks, which freeze a portion of the model parameters during the entire training process, to reduce the communication cost with little implications on model performance. Through extensive experiments, we empirically show that Federated learning of Partially Trainable neural networks (FedPT) can result in superior communication-accuracy trade-offs, with up to $46\times$ reduction in communication cost, at a small accuracy cost. Our approach also enables faster training, with a smaller memory footprint, and better utility for strong differential privacy guarantees. The proposed FedPT method can be particularly interesting for pushing the limitations of overparameterization in on-device learning.

1 Introduction

A large trove of data are being generated with the proliferation of edge devices, such as mobile phones, medical sensors, and smart home devices. While those data can be used to develop highly intelligent algorithms, they often contain private information that users are not willing to share with others. In recent years, Federated Learning (FL) [McMahan et al., 2017a] has been introduced as an alternative to centralized learning to protect user privacy when training a machine learning model. In FL, participating clients collaboratively learn a shared model under the supervision of a central server: each communication round often starts with the server broadcasting the global model to the participants; these participants then perform computations on their local data, and send their aggregated updates back to the server to update the global model [Kairouz et al., 2019]. While FL can be performed on a relatively small number of clients, many applications involve a large number of edge devices, such as mobile phones, or sensors [Ramaswamy et al., 2020; Sheller et al., 2020; Granqvist et al., 2020]. This setting is referred to as cross-device FL. Training large models on edge devices is challenging due to unreliable connections, and limited computational capabilities.

Federated averaging (FedAvg) [McMahan et al., 2017a] is one of the most popular algorithms in federated learning. It was later generalized to a framework of two stage optimization [Reddi et al., 2020]: a client optimizer is used to update local models from the client’s local data, and a server optimizer is used to update the global model from the aggregated client updates. Instead of averaging client local models to replace the global model, the client updates (the difference between the initial model received from the server, and the client local model after training on private data) are aggregated, and used as pseudo-gradients to update the global model. We build our method based on generalized FedAvg, as it can satisfy many restricting constraints of cross-device FL [Wang et al., 2021b].

By combining federated learning with differential privacy (DP), a stronger privacy protection can be provided to the clients [Kairouz et al., 2019, 2021b]. Specifically, differential privacy can prevent memorization, and protect against potential leakage of user data from a model released publicly [Ramaswamy et al., 2020]. Training large models with DP is known to be hard [Tramèr and Boneh, 2020], and reducing the dimensionality in DP is an active area of research [Zhou et al., 2020; Kairouz et al., 2021a; Wang et al., 2021b].

Deep neural networks have achieved impressive performance on various machine learning tasks, which have become popular in federated learning in recent years [Kairouz et al., 2019; Li et al., 2020; Yang et al., 2019]. The performance of deep neural networks can often be improved by increasing the model size in the overparameterized regime [Allen-Zhu et al., 2019; Neyshabur et al., 2018; Sankararaman et al., 2020; Du et al., 2018]. Even though parameters of overparameterized networks are believed to be redundant [Frankle and Carbin, 2019], and can be pruned, the size of the model is believed to help regularize the optimization landscape to facilitate training. More recently, researchers have observed that training a small fraction of the parameters of a large model, such as batch normalization layers in convolutional networks [Frankle et al., 2021], can achieve comparable performance as training all the parameters. Inspired by the recent deep learning research in centralized training, we study the effects of freezing part of the parameters of a large model in federated learning.

In this paper, we propose the use of partially trainable networks (PTNs) to tackle the communication and computation burdens of training large models, which can make FL more accessible to various applications. Based on the generalized FedAvg algorithm, we only communicate the trainable parameters (which can represent as little as 2% of the total network parameter count in our study), and a random seed from server to clients. Clients reconstruct the full model by regenerating the frozen parameters from the random seed, perform local training on private data, and send back the updates on the trainable parameters to the server. The communication can be significantly saved by the size of frozen parameters. Clients can also save local computations, and memory on gradient calculations for the frozen parameters. The proposed FedPT algorithm is used to train various network architectures, including convolutional networks for computer vision, and Transformers for language tasks. We empirically show the advantages of FedPT by running experiments on three benchmark datasets: EMNIST, CIFAR-10, and Stack Overflow. Reducing the communication cost by $1.1\times$ to $46\times$ results in a loss in accuracy between 0.1% and 4%. The simulation training times are reduced by 25%, and the memory footprint is reduced by 10% in some settings. Moreover, differentially private FedPT achieves better utility (accuracy) than training the full model when the privacy protection is strong (small ϵ).

2 Related Work

Communication Efficient Federated Learning McMahan et al. [2017a] exploit local updates to save the number of communication rounds in FL. Konečný et al. [2016] introduce compression techniques to reduce the uplink communication costs for the clients. They show a reduction of up to two orders of magnitude in these communication costs. Caldas et al. [2019b] apply the lossy compression techniques used in the clients in [Konečný et al., 2016] to the model on the server. The authors also introduce federated dropout, in which each client receives a submodel from the server by randomly removing some parameters. Closely related to federated dropout, ordered dropout [Horvath et al., 2021] extracts submodels from a main model, and adapts the computation and communication costs to the edge device capabilities. Singhal et al. [2021] propose a method to partially reconstruct a model on the edge devices, which results in a reduced communication cost. The reconstructed parameters are local to the clients, and never sent to the server. Wang et al. [2021a] proposed training low-rank, pre-factorized deep networks to reduce communication in distributed learning. Other methods, like compression, and knowledge distillation have been used in FL to reduce the communication costs. We refer the interested reader to [Wang et al., 2021b] for more details. The proposed FedPT is complementary to these methods, and can be combined with them to further improve the communication efficiency in federated learning.

Partially Trainable Neural Networks Partially trainable neural networks have a long history in machine learning. Extreme learning machines can be considered single layer networks with frozen hidden nodes [Huang et al., 2015]. Echo state networks [Jaeger, 2002], and liquid state machines

[Maass et al., 2002] disentangle randomly projected inputs with a (generally linear) readout layer. More recently, Zhang et al. [2019] studied the properties of layers in overparameterized deep models. They distinguish *ambient* layers from *critical* ones, and show that the ambient layers are robust to re-randomization and re-initialization in the post-training validation. Rosenfeld and Tsotsos [2018] show competitive results on vision tasks where they freeze most of the weights of the model. Frankle et al. [2021] only train the parameters of batch normalization layers of vision models, and freeze the other weights. Randomized networks have also been studied for language tasks, such as sentence classification [Wieting and Kiela, 2019], machine translation [Garg et al., 2020; Shen et al., 2021], and speech recognition [Shrivastava et al., 2021], where learning only linear mappings on top of random features achieves surprisingly high accuracy. We are not aware of any previous work applying PTNs to federated learning.

Efficient Transformers A lot of attention has been put lately on designing efficient architectures for Transformers [Vaswani et al., 2017; Fournier et al., 2021]. FNETs [Lee-Thorp et al., 2021] replace the self-attention layer by a Fourier transform, reducing its computational complexity from quadratic in the input sequence size to linear. Shen et al. [2021] use random layers as a computationally efficient way to increase network depth. They intersperse reservoir frozen Transformer layers with regular, trainable ones. They also prove that backward pass can be skipped altogether for these layers. Choromanski et al. [2021] approximate the attention using a kernel with randomized mappings. Wang et al. [2020] approximate the attention with a low-rank factorization. Child et al. [2019] introduce the Sparse Transformer, in which the attention matrix is approximated with a sparse factorization. We apply FedPT to efficiently train Transformers in federated learning.

Model Compression Pruning [LeCun et al., 1990] reduces the model size by deleting the weights whose removal have a minimal effect on the model performance. It was observed in Frankle and Carbin [2019] that sub-networks can be found in a dense neural network that can match the test accuracy of the original network. Knowledge distillation [Ba and Caruana, 2014; Hinton et al., 2015] transfers the knowledge of a large teacher into a smaller student, which can then be used at inference time. While the methods mentioned can be leveraged to make the inference of federated trained models lighter, they do not alleviate the bottlenecks present during on-device training. Our proposed FedPT, on the other hand, aims at making the training process more efficient.

3 Partially Trainable Networks for Federated Learning

3.1 FedPT Method

Algorithm 1: FedPT: federated learning of partially trainable neural networks

Input: Initial model $\mathbf{x}^{(0)}$; ClientOpt, ServerOpt with learning rates η, α

- 1 Split $\mathbf{x}^{(0)}$ into trainable part $\mathbf{y}^{(0)}$, and non-trainable part generated by random seed z
- 2 **for** $t \in \{0, 1, \dots, T - 1\}$ **do**
- 3 Send $(\mathbf{y}^{(t)}, z)$ to a subset $\mathcal{S}^{(t)}$ of clients
- 4 **for client** $i \in \mathcal{S}^{(t)}$ **in parallel do**
- 5 Initialize local model $\mathbf{x}_i^{(t,0)} = \text{Reconstruct}(\mathbf{y}^{(t)}, z)$
- 6 **for** $k = 0, \dots, \tau_i - 1$ **do**
- 7 Compute local stochastic gradient $g_i(\mathbf{y}_i^{(t,k)})$ by backprop through $\mathbf{x}_i^{(t,k)}$
- 8 Perform local update $\mathbf{y}_i^{(t,k+1)} = \text{ClientOpt}(\mathbf{y}_i^{(t,k)}, g_i(\mathbf{y}_i^{(t,k)}), \eta, t)$
- 9 **end**
- 10 Compute and send back local model changes $\Delta_i^{(t)} = \mathbf{y}_i^{(t,\tau_i)} - \mathbf{y}_i^{(t,0)}$
- 11 **end**
- 12 Aggregate local changes $\Delta^{(t)} = \sum_{i \in \mathcal{S}^{(t)}} p_i \Delta_i^{(t)} / \sum_{i \in \mathcal{S}^{(t)}} p_i$
- 13 Update global model $\mathbf{y}^{(t+1)} = \text{ServerOpt}(\mathbf{y}^{(t)}, -\Delta^{(t)}, \alpha, t)$
- 14 **end**

We use PTNs in federated learning tasks in the proposed FedPT algorithm. Inspired by previous partially trainable models in the centralized setting [Huang et al., 2015; Giryas et al., 2016], we freeze a set of parameters after initialization. This allows us to summarize the frozen parameters into a single random seed, provided that the server and clients share the same random number generator. This seed is sent to the clients (edge devices), and used to reconstruct the frozen parameters. Clients do not need to send updates for the frozen parameters back. We summarized the method in algorithm 1, which is based on the generalized FedAvg algorithm with two stage optimization by ServerOpt and ClientOpt. In cross-device FL, only a small subset of the clients $\mathcal{S}^{(t)}$ (compared to the large population) can be accessed at each communication round t . We use the number of local samples on client i as weight p_i to aggregate the local updates following [McMahan et al., 2017a].

The design of PTNs highly depends on the network architecture. Since freezing a small number of parameters does not impact the communication efficiency notably, we focus on freezing layers that contain a large proportion of the parameters in our experiments, and choose different layers for different architectures. To maximize the communication and computation efficiencies of the PTN, we suggest the following principles for experimental design:

1. Start from freezing the largest parameter block of a network.
2. Add more blocks if it did not degrade the model performance on utility (accuracy).
3. Switch to a smaller block if it did degrade the model performance by a large margin.
4. Repeat these steps to find the best PTN.

Once an optimal PTN is found for a specific network architecture, the same PTN can be used for various application tasks that use the same network architecture.

We study several popular network architectures: ResNet [He et al., 2016] with group normalization for image tasks, small convolutional neural networks as feature extractors with a few fully connected layers for classification, and Transformer [Vaswani et al., 2017] for language tasks. We experiment with freezing the following layers, and test on FL benchmark datasets:

- Convolutional layers in resnet-18 trained on CIFAR-10.
- Encoder dense layers in the Transformer architecture used in Stack Overflow.
- Dense layer following the convolutional layers in the convolutional model used in EMNIST.

The networks and benchmarks are introduced in [McMahan et al., 2017a; Hsu et al., 2019; Reddi et al., 2020; Wang et al., 2021b], and are widely used for cross-device FL. We notice that freezing the normalization layers in FedPT notably hurts the performance of ResNet, as observed in [Frankle et al., 2021] in the centralized setting. Since the normalization layers usually have a small number of parameters, we always train them in FedPT. The frozen parameters are set to their initial values, which are generated from Gaussian initializers.

3.2 Advantages of FedPT

Communication Efficiency Communication is one of the main bottlenecks in cross-device federated learning. Model transmission from server to devices can be a serious constraint for the server, particularly when stragglers that have limited network connections exist. Sending the model updates back to the server can be even more challenging, as uplink is typically much slower than downlink. The download and upload bandwidths in a real cross-device FL system are estimated at 0.75MB/s and 0.25MB/s respectively by Wang et al. [2021b]. FedPT can mitigate this issue: the frozen parameters are compressed into a random seed sent from server to clients, and participating clients do not need to send any updates back for these frozen parameters.

Differential Privacy (DP) Federated learning is designed for privacy protection, as the clients do not share their private data. Combining FL with DP can provide even stronger privacy defenses [Kairouz et al., 2019; Wang et al., 2021b]. The most popular method to achieve user-level DP in FL is a variant of generalized FedAvg, called DP-FedAvg, which clips the model updates $\Delta_i^{(t)}$, and adds noise to the aggregated $\Delta^{(t)}$ [McMahan et al., 2017b]. It is nontrivial to get good utility-privacy trade-offs when applying DP in FL, as it is challenging for practical cross-device FL to perform

sampling for privacy amplification [Wang et al., 2021b]. Recently, Kairouz et al. [2021b] proposed to use the Differentially Private Follow the Regularized Leader (DP-FTRL) method on the server in the generalized FedAvg framework to get formal (ϵ, δ) -DP guarantees. We combine DP-FTRL with FedPT to get (ϵ, δ) privacy guarantees. In FedPT, we only need to make the trainable parameters private, and this dimension reduction effect can be beneficial for differential privacy. We only clip the client updates $\Delta_i^{(t)}$, and add noise to the aggregated $\Delta^{(t)}$ of trainable parameters, which can potentially tolerate smaller clip norms, and larger noises. In our experiments, partially trained models are more resilient to high noise levels than fully trained ones.

Training time Reducing the client training time allows more devices to complete their local computations in the allotted time in a round, which is often desirable in practical FL systems [Bonawitz et al., 2019; Wang et al., 2021b]. In addition to that, reducing the training time makes it possible to train larger models in production settings where the federated learning tasks have a limited amount of time to run on edge devices. FedPT can reduce the training time, as it does not need to calculate gradients for the frozen parameters. We observe, in some of our simulations, that FedPT has a reduction in the client training time as the number of frozen parameters increases. Our study shows that FedPT has a significant decrease in runtime for deep convolutional models, for which we freeze the convolutional layers.

Memory footprint FedPT can also improve the memory footprint of model training in FL, as we do not need to save intermediate activations of frozen layers after backprop, and we do not need to calculate or save the gradients for the frozen layers. Besides, when computing the model updates on clients, two copies of the trainable parameters (new value, and old one) are needed to generate the client model update Δ_i , but none is needed for the frozen ones. We compare the memory footprints of fully/partially trainable models, and see a reduction in the settings where we freeze the convolutional layers.

4 Experiments

Setup We use three datasets: EMNIST¹ [Cohen et al., 2017], CIFAR-10 [Krizhevsky et al., 2009], and Stack Overflow². We follow [Hsu et al., 2019] to use a Dirichlet distribution of parameter 1 to create a heterogeneous dataset for CIFAR-10, and [Wang et al., 2021b] to preprocess this dataset. We follow [Reddi et al., 2020; Wang et al., 2021b] to preprocess the EMNIST, and Stack Overflow datasets. EMNIST and CIFAR-10 are image datasets, for which we train convolutional networks. Stack Overflow is a language dataset, for which we train a Transformer-based next word prediction model. We measure runtime and memory footprint on desktop CPUs to simulate on-device training. More details on models and hyperparameters used are provided in the Appendix.

4.1 Efficiency of FedPT

Table 1: Percentage of trainable parameters, reduction in communication cost, average accuracies, and training times for EMNIST.

Trainable Parameters (%)	Reduction in Communication	Final Accuracy	Runtime (s)
4.97	20×	83.75 ± 0.16	2.90 ± 0.42
100	1×	85.47 ± 0.21	2.72 ± 0.37

EMNIST with CNN Table 1 shows the results for EMNIST. We freeze the dense layer following the convolutional ones for this model presented in table 6. The reduction in communication cost is 20 fold, for an absolute loss in accuracy of less than 2%, and similar runtimes.

¹https://www.tensorflow.org/federated/api_docs/python/tff/simulation/datasets/emnist

²https://www.tensorflow.org/federated/api_docs/python/tff/simulation/datasets/stackoverflow

Table 2: Percentage of trainable parameters, reduction in communication cost, average accuracies, and training times for CIFAR-10.

Trainable Parameters (%)	Reduction in Communication	Final Accuracy	Runtime (s)
2.16	46.3×	79.49 ± 1.1	16.22 ± 0.076
3.47	28.8×	80.76 ± 0.72	16.58 ± 0.053
8.07	12.4×	82.82 ± 0.18	16.96 ± 0.021
26.25	3.8×	83.57 ± 0.41	17.89 ± 0.057
100	1.0×	83.65 ± 0.26	20.51 ± 0.080

CIFAR-10 with ResNet Table 2 summarizes the results for the CIFAR-10 task, trained on resnet-18. The training time per round monotonically decreases as we freeze more parameters. While the accuracy decreases as well, a reduction of 4× in the communication costs results in almost the same accuracy.

Table 3: Percentage of trainable parameters, reduction in communication cost, average accuracies, and training times for Stack Overflow.

Trainable Parameters (%)	Reduction in Communication	Final Accuracy	Runtime (s)
73.8	1.4×	23.23 ± 0.22	9.37 ± 0.34
82.6	1.2×	23.71 ± 0.16	9.23 ± 0.46
91.3	1.1×	24.09 ± 0.16	9.42 ± 0.55
100	1.0×	24.50 ± 0.19	9.74 ± 0.46

Stack Overflow with Transformer Table 3 shows the results for Stack Overflow Next Word Prediction. In this task, the training times are not sensibly changed, and we can trade-off 10% of communication costs for less than 0.5% accuracy.

Memory footprint In order to verify our intuition that the memory footprint can be positively impacted by FedPT, we profile the memory in the CIFAR-10 task in simulation experiments. Our results, summarized in table 4, show a reduction in the peak memory usage as the proportion of frozen parameters increases.

Table 4: Peak memory at different percentages of trainable parameters for CIFAR-10.

Trainable Parameters (%)	Peak Memory (MiB)
2.16	505.94
3.47	506.37
8.07	506.37
26.25	520.43
100	565.43

4.2 Differential Privacy

We experiment with DP-FTRL [Kairouz et al., 2021b] on the Stack Overflow next word prediction task. We use the same noise multipliers, number of clients per round (100), and total number of rounds (1600) as in [Kairouz et al., 2021b], and hence achieve the same (ϵ, δ) privacy guarantees. The results are presented in Table 5. We observe that FedPT is more resilient to high levels of noise, making it a better fit to offer strong privacy guarantees. An intuitive explanation is that, by updating less parameters, FedPT is less affected by a given noise level (the noise is added to less parameters).

5 Discussion and Future Work

Our experiments show that PTNs can be used to improve the efficiency and privacy in FL. By leveraging the ability to summarize the frozen parameters into a random seed when broadcasting

Table 5: Training accuracies (%) for Fully Trainable (FT), and Partially Trainable (PT) models with DP.

Noise multiplier	0.0	1.13	2.33	4.03	6.21	8.83
Privacy ϵ	∞	18.71	7.83	4.19	2.60	1.77
FT	23.22	19.67	18.37	17.28	16.24	14.60
PT	22.44	19.58	18.38	17.03	16.28	15.01

them from the server to the clients, the communication cost can be drastically reduced. Reducing the number of trainable parameters also has implications for the privacy guarantees. We suspect that the accuracy gap we observe in our experiments between fully trainable, and partially trainable models can be reduced in multiple ways, such as running a grid search to find the optimal parameters for the partially trainable parameters, or adding layers of frozen parameters.

We present extensive simulation experiments to suggest PTNs can have a positive impact on real world FL applications. Note that our profiling experiments ran in a data center for simulation, while real-world training happens on edge devices for cross-device FL. We are interested in studying the effects on memory and runtime in real-world on-device training, which may depend on the network architectures, the parameters which are frozen, and implementation details. Though the on-device performance can be reasonably inferred from data center simulation [Wang et al., 2021b], the data center and on-device settings can use different computation optimizations in implementation and hardware.

We are also interested in combining FedPT with other methods for improving communication and computation efficiency, such as federated dropout [Caldas et al., 2019b], to further ease the communication bottleneck inherent to FL. Furthermore, we presume that FedPT can benefit from better generalization bounds, similar to the ones shown in [Zhang et al., 2019] based on parameter counting [Anthony and Bartlett, 2009]. The application of FedPT to edge devices can result in more clients contributing to the training of the global model, potentially leading to better fairness.

Another possible extension to our work is to selectively freeze more parameters for devices with smaller bandwidth and/or computational capacity, while training more parameters on devices that do not suffer such limitations. Finally, a systematic study of the relationship between the final test accuracy, and the architecture-specific frozen layers can provide more insights and guidelines for improving FedPT’s performance.

6 Conclusion

This paper introduces FedPT for training partially trainable neural networks in the context of federated learning. Our empirical results suggest that it is possible to achieve competitive accuracy using federated models when a significant portion of the parameters are untrained, which improves efficiency in terms of communication cost, and computational complexity. We also show that partially trainable models can be more resilient to high additive noise, enabling better utility for stronger privacy guarantees than fully trained models. Based on these findings, FedPT is a promising method to improve efficiency and privacy for practical federated learning applications.

Acknowledgements

The authors would like to thank Zachary Garrett, Jakub Konečný, Keith Rush, Changwan Ryu, Om Thakkar, and Petr Zadrazil for helpful discussions.

References

Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. *Advances in neural information processing systems*, 2019.

- Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, USA, 1st edition, 2009. ISBN 052111862X.
- Lei Jimmy Ba and Rich Caruana. Do deep nets really need to be deep?, 2014.
- Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, H Brendan McMahan, et al. Towards federated learning at scale: System design. *SysML*, 2019.
- Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. Leaf: A benchmark for federated settings, 2019a.
- Sebastian Caldas, Jakub Konečný, H. Brendan McMahan, and Ameet Talwalkar. Expanding the reach of federated learning by reducing client resource requirements, 2019b.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers, 2019.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. Rethinking attention with performers, 2021.
- Gregory Cohen, Saeed Afshar, Jonathan Tanson, and André van Schaik. Emnist: an extension of mnist to handwritten letters, 2017.
- Simon S Du, Xiyu Zhai, Barnabas Póczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- Quentin Fournier, Gaétan Marceau Caron, and Daniel Aloi. A practical survey on faster and lighter transformers, 2021.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks, 2019.
- Jonathan Frankle, David J. Schwab, and Ari S. Morcos. Training batchnorm and only batchnorm: On the expressive power of random features in cnns, 2021.
- Ankush Garg, Yuan Cao, and Qi Ge. Echo state neural machine translation, 2020.
- Raja Giryes, Guillermo Sapiro, and Alex M. Bronstein. Deep neural networks with random gaussian weights: A universal classification strategy? *IEEE Transactions on Signal Processing*, 64 (13):3444–3457, Jul 2016. ISSN 1941-0476. doi: 10.1109/tsp.2016.2546221. URL <http://dx.doi.org/10.1109/TSP.2016.2546221>.
- Filip Granqvist, Matt Seigel, Rogier van Dalen, Áine Cahill, Stephen Shum, and Matthias Paulik. Improving on-device speaker verification using federated learning with privacy, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
- Samuel Horvath, Stefanos Laskaridis, Mario Almeida, Ilias Leontiadis, Stylianos I Venieris, and Nicholas D Lane. Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout. *arXiv preprint arXiv:2102.13451*, 2021.
- Kevin Hsieh, Amar Phanishayee, Onur Mutlu, and Phillip B. Gibbons. The non-iid data quagmire of decentralized machine learning, 2020.
- Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification, 2019.
- Gao Huang, Guang-Bin Huang, Shiji Song, and Keyou You. Trends in extreme learning machines: A review. *Neural Networks*, 61:32–48, 2015.

- Alex Ingerman and Krzys Ostrowski. Tensorflow federated. <https://medium.com/tensorflow/introducing-tensorflow-federated-a4147aa20041>, 2019.
- Herbert Jaeger. Adaptive nonlinear system identification with echo state networks. *Advances in neural information processing systems*, 15:609–616, 2002.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- Peter Kairouz, Monica Ribero Diaz, Keith Rush, and Abhradeep Thakurta. (nearly) dimension independent private erm with adagrad rates via publicly estimated subspaces. In *Conference on Learning Theory*, pages 2717–2746. PMLR, 2021a.
- Peter Kairouz, Brendan McMahan, Shuang Song, Om Thakkar, Abhradeep Thakurta, and Zheng Xu. Practical and private (deep) learning without sampling or shuffling, 2021b.
- Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pages 598–605, 1990.
- James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, and Santiago Ontanon. Fnet: Mixing tokens with fourier transforms, 2021.
- Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560, 2002.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agueria y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017a.
- H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*, 2017b.
- Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018.
- Swaroop Ramaswamy, Om Thakkar, Rajiv Mathews, Galen Andrew, H Brendan McMahan, and Françoise Beaufays. Training production language models without memorizing user data. *arXiv preprint arXiv:2009.10031*, 2020.
- Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*, 2020.
- Amir Rosenfeld and John K. Tsotsos. Intriguing properties of randomly weighted networks: Generalizing while learning next to nothing, 2018.
- Karthik Abinav Sankararaman, Soham De, Zheng Xu, W Ronny Huang, and Tom Goldstein. The impact of neural network overparameterization on gradient confusion and stochastic gradient descent. In *International Conference on Machine Learning*, pages 8469–8479. PMLR, 2020.

- Micah J Sheller, Brandon Edwards, G Anthony Reina, Jason Martin, Sarthak Pati, Aikaterini Kotrotsou, Mikhail Milchenko, Weilin Xu, Daniel Marcus, Rivka R Colen, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports*, 10(1):1–12, 2020.
- Sheng Shen, Alexei Baevski, Ari S. Morcos, Kurt Keutzer, Michael Auli, and Douwe Kiela. Reservoir transformers, 2021.
- Harsh Shrivastava, Ankush Garg, Yuan Cao, Yu Zhang, and Tara Sainath. Echo state speech recognition. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5669–5673, 2021. doi: 10.1109/ICASSP39728.2021.9414495.
- Karan Singhal, Hakim Sidahmed, Zachary Garrett, Shanshan Wu, Keith Rush, and Sushant Prakash. Federated reconstruction: Partially local federated learning. *arXiv preprint arXiv:2102.03448*, 2021.
- Florian Tramèr and Dan Boneh. Differentially private learning needs better features (or much more data). *arXiv preprint arXiv:2011.11660*, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.
- Hongyi Wang, Saurabh Agarwal, and Dimitris Papailiopoulos. Pufferfish: Communication-efficient models at no extra cost. *arXiv preprint arXiv:2103.03936*, 2021a.
- Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H Brendan McMahan, Blaise Aguerre y Arcas, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, et al. A field guide to federated optimization. *arXiv:2107.06917*, 2021b.
- Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity, 2020.
- John Wieting and Douwe Kiela. No training required: Exploring random encoders for sentence classification. *CoRR*, abs/1901.10444, 2019. URL <http://arxiv.org/abs/1901.10444>.
- Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- Chiyuan Zhang, Samy Bengio, and Yoram Singer. Are all layers created equal?, 2019.
- Yingxue Zhou, Zhiwei Steven Wu, and Arindam Banerjee. Bypassing the ambient dimension: Private sgd with gradient subspace identification. *arXiv preprint arXiv:2007.03813*, 2020.