# Model Pruning Enables Efficient Federated Learning on Edge Devices

Yuang Jiang, Shiqiang Wang, Víctor Valls, Bong Jun Ko, Wei-Han Lee, Kin K. Leung, Leandros Tassiulas

Abstract—Federated learning (FL) allows model training from local data collected by edge/mobile devices while preserving data privacy, which has wide applicability to image and vision applications. A challenge is that client devices in FL usually have much more limited computation and communication resources compared to servers in a datacenter. To overcome this challenge, we propose PruneFL - a novel FL approach with adaptive and distributed parameter pruning, which adapts the model size during FL to reduce both communication and computation overhead and minimize the overall training time, while maintaining a similar accuracy as the original model. PruneFL includes initial pruning at a selected client and further pruning as part of the FL process. The model size is adapted during this process, which includes maximizing the approximate empirical risk reduction divided by the time of one FL round. Our experiments with various datasets on edge devices (e.g., Raspberry Pi) show that: (i) we significantly reduce the training time compared to conventional FL and various other pruning-based methods; (ii) the pruned model with automatically determined size converges to an accuracy that is very similar to the original model, and it is also a lottery ticket of the original model.

*Index Terms*—Efficient training, federated learning, model pruning.

#### I. INTRODUCTION

T HE past decade has seen a rapid development of machine learning algorithms and applications, particularly in the area of deep neural networks (DNNs) [1]. However, a huge volume of training data is usually required to train accurate models for complex tasks such as image classification and computer vision. Due to limits in data privacy regulations and communication bandwidth, it is usually infeasible to transmit and store all training data at a central location. To address this problem, *federated learning* (FL) has emerged as a promising approach of distributed model training from decentralized data [2]–[6]. In a typical FL system, data is collected by client devices (e.g., cameras) at the network edge; the training process includes local model updates using each client's own data and the fusion of all clients' models typically through a server. In this way, the raw data remains local in clients.

Client devices in FL are usually much more resourceconstrained than server machines in a datacenter, in terms of

Y. Jiang, V. Valls, and L. Tassiulas are with Yale University, New Haven, CT, USA.

S. Wang and W.-H. Lee are with IBM T. J. Watson Research Center, Yorktown Heights, NY, USA.

B. J. Ko was with Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford, CA, USA when contributing to this work.

K. K. Leung is with Imperial College London, UK.

Contact authors: Y. Jiang (yuang.jiang@yale.edu) and S. Wang (wangshiq@us.ibm.com)

Accepted for publication in IEEE Transactions on Neural Networks and Learning Systems (TNNLS)

computation power, communication bandwidth, memory and storage size, etc. Training DNNs that can include over millions of parameters (weights) on such resource-limited edge devices can take prohibitively long and consume a large amount of energy. Therefore, a natural question is: *how can we perform FL efficiently so that a model is trained within a reasonable amount of time and energy*?

Some progress has been made towards this direction recently using model/gradient compression techniques, where instead of training the original model with full parameter vector, either a small model is extracted from the original model for training or a compressed parameter vector (or its gradient) is transmitted in the fusion stage [7]–[10]. However, the former approach may reduce the accuracy of the final model in undesirable ways, whereas the latter approach only reduces the communication overhead and does not generate a small model for efficient computation. Furthermore, how to adapt the compressed model size for the most efficient training remains a largely unexplored area, which is a challenging problem due to unpredictable training dynamics and the need of obtaining a good solution in a short time with minimal overhead.

To overcome these problems, we propose a new FL paradigm called *PruneFL*, which includes *adaptive and distributed parameter pruning* as part of the FL procedure. We make the following key contributions.

**Distributed pruning.** PruneFL includes initial pruning at a selected client followed by further distributed pruning that is intertwined with the standard FL procedure. Our experimental results show that this method outperforms alternative approaches that either prunes at a single client only or directly involves multiple FL clients for pruning, especially when the clients have heterogeneous data statistics and computational power.

Adaptive pruning. PruneFL continuously "tracks" a model that is small enough for efficient transmission and computation with low memory footprint, while maintaining useful connections and their parameters so that the model converges to a similar accuracy as the original model. The importance of model parameters evolves during training, so our method continuously updates which parameters to keep and the corresponding model size. The update follows an objective of minimizing the time of reaching intermediate and final loss values. Each FL round operates on a small pruned model, which is efficient. A small model is also obtained at any time during and after the FL process for efficient inference on edge devices, which is a lottery ticket [11] of the original model as we show experimentally. **Implementation.** We implement FL with model pruning on real edge devices, where we extend a deep learning framework to support efficient sparse matrix computation. Our code is available at: https://github.com/jiangyuang/PruneFL

## II. RELATED WORK

**Neural network pruning.** To reduce the complexity of neural network models, different ways of parameter pruning were proposed in the literature. Early work considered approximation using second-order Taylor expansion [12]. However, the computation of Hessian matrix has high complexity which is infeasible for modern DNNs. In recent years, magnitude-based pruning has become popular [13], where parameters with small enough magnitudes are removed from the network. A finding that suggests a network that is pruned by magnitude consists of an optimal substructure of the original network, known as "lottery ticket hypothesis", was presented in [11], [14]. It shows that directly training the pruned network can reach a similar accuracy as pruning a pre-trained original network.

In addition to the above approaches that train until convergence before the next pruning step, there are iterative pruning methods where the model is pruned after every few steps of training [15], [16]. There are also one-shot pruning approaches including SynFlow [17] that prunes the model at model initialization (before training), and SNIP [18] that prunes the model using first training round's gradient information. A dynamic pruning approach that allows the network to grow and shrink during training was proposed in [19]. Besides these unstructured pruning methods, structured pruning was also studied [20], which, however, often requires specific network architectures and does not conform to the lottery ticket hypothesis. The lottery ticket is useful for retraining a pruned model on a different yet similar dataset [14]. The use of pruning for efficient model training was discussed in [21], where the optimal choice of pruning rate (or final model size) remained unstudied.

These existing pruning techniques consider the centralized setting with full access to training data, which is fundamentally different from our PruneFL that works with decentralized datasets at local clients. Furthermore, the automatic adaptation of model size has not been studied before.

Efficient federated learning. The first FL method is known as federated averaging (FedAvg) [2], where each "round" of training includes multiple local gradient computation steps on each client's local data, followed by a parameter averaging step through a server. This method can be shown to converge in various settings including when the data at different clients are non-identically distributed (non-IID) [22]–[24].

To improve the communication efficiency of FL, methods for optimizing the communication frequency were studied [25]–[27]. An approach of parameter averaging using structured, sketched, and quantized updates was introduced in [7], which belongs to the broader area of gradient compression/sparsification [28]–[33]. These techniques usually consider a fixed degree of sparsity or compression that needs to be configured as a hyperparameter. An online learning approach that determines a near-optimal gradient sparsity was proposed in [9], which includes exploration steps that may slow down the training initially. This body of work does not address computation efficiency.

To reduce both communication and computation costs, efficient FL techniques using lossy compression and dropout were developed [8], [10], where the final model still has the original size and hence providing no benefit for efficient inference after the model is trained. Moreover, because the main goal of pruning is to remove less important weights from the model, it is orthogonal to other acceleration methods such as quantization [34], low-rank decomposition [35], etc., and pruning can be applied together with these other methods. In addition, since our approach considers acceleration for both training and inference, methods that accelerate inference only, e.g., knowledge distillation [36], runtime neural pruning [37], and DNN partitioning and offloading [38], do not serve our purpose. There are other distributed training methods, such as split learning [39], which are beyond our scope since we focus on FL in this paper.

Furthermore, most existing studies on FL are based on simulation. Only a few recent papers considered implementation on real embedded devices [10], [26], but they do not include parameter pruning.

**Novelty of our work.** The uniqueness of PruneFL is that we jointly address communication and computation efficiency for both training and inference phases, by extending FedAvg with minimal extra overhead. Our two-stage distributed pruning method is designed to address both data (statistical) and device (system) heterogeneity including non-IID data distribution. Our adaptive pruning method is uniquely based on gradient information, which does not require sharing clients' local data, so that existing privacy preservation and secure aggregation [40] methods for FL can be directly applied to the gradient. Thus, our approach does not introduce extra privacy concerns.

**Roadmap.** The remainder of this paper is organized as follows. Section III provides preliminaries of FL and model pruning. Section IV presents the proposed PruneFL approach and its analysis. Implementation challenges are discussed in Section V. Section VI presents the experimental setup and results. Section VII draws conclusion.

#### **III. PRELIMINARIES**

**Federated learning.** We consider an FL system with N clients. Each client  $n \in [N] := \{1, 2, ..., N\}$  has a local empirical risk  $F_n(\mathbf{w}) := \frac{1}{D_n} \sum_{i \in D_n} f_i(\mathbf{w})$  defined on its local dataset  $\mathcal{D}_n (D_n := |\mathcal{D}_n|)$  for model parameter vector  $\mathbf{w}$ , where  $f_i(\mathbf{w})$  is the loss function (e.g., cross-entropy, mean square error, etc.) that captures the difference between the model output and the desired output of data sample *i*. The system tries to find a parameter  $\mathbf{w}$  that minimizes the global empirical risk:

$$\min_{\mathbf{w}} F(\mathbf{w}) := \sum_{n \in [N]} p_n F_n(\mathbf{w})$$
(1)

where  $p_n > 0$  are weights such that  $\sum_{n \in [N]} p_n = 1$ . For example, if  $\mathcal{D}_n \cap \mathcal{D}_{n'} = \emptyset$  for  $n \neq n'$  and  $p_n = D_n/D$ 



Fig. 1. Illustration and flowchart of PruneFL.

Notation	Definition
$\odot$	element-wise product of two vectors
n,N	client index, total number of clients
k, K	iteration index, total number of iterations
Ι	number of local iterations
$p_n$	weight for client $n \ (p_n > 0, \forall n, \text{ and } \sum_n p_n = 1)$
$\mathbf{m}(k)$	weight mask in iteration $k$ (universal for all clients)
$\mathbf{w}_n(k)$	client $n$ 's parameter in iteration $k$
$\mathbf{w}(k)$	$\mathbf{w}(k) := \sum_{n=1}^{N} p_n \mathbf{w}_n(k)$
$\mathbf{w}_n'(k),\mathbf{w}'(k)$	$\mathbf{w}'_n(k) = \mathbf{w}_n(k) \odot \mathbf{m}(k),  \mathbf{w}'(k) = \mathbf{w}(k) \odot \mathbf{m}(k)$
$\mathbf{g}_n(\mathbf{w})$	client $n$ 's stochastic gradient with parameter $\mathbf{w}$
$\nabla F_n(\mathbf{w})$	client $n$ 's expected gradient with parameter $\mathbf{w}$
$ abla F(\mathbf{w})$	$\nabla F(\mathbf{w}) = \sum_{n=1}^{N} p_n F_n(\mathbf{w})$

TABLE I MAIN NOTATIONS.

with  $\mathcal{D} := \bigcup_n \mathcal{D}_n$  and  $D := |\mathcal{D}|$ , we have  $F(\mathbf{w}) = \frac{1}{D} \sum_{i \in \mathcal{D}} f_i(\mathbf{w})$ . Other ways of configuring  $p_n$  may also be used to account for fairness and other objectives [41].

In FL, each client n has a local parameter  $\mathbf{w}_n(k)$  in iteration k. The aggregation of these local parameters is defined as  $\mathbf{w}(k) := \sum_{n \in [N]} p_n \mathbf{w}_n(k)$ . The FL procedure usually involves multiple updates of  $\mathbf{w}_n(k)$  using stochastic gradient descent (SGD) on the local empirical risk  $F_n(\mathbf{w}_n(k))$ computed by every client n, followed by a parameter fusion step that involves the server collecting clients' local parameters  $\{\mathbf{w}_n(k) : \forall n \in [N]\}$  and computing the aggregated parameter  $\mathbf{w}(k)$ . After parameter fusion, the local parameters  $\{\mathbf{w}_n(k) : \forall n \in [N]\}$  are all set to be equal to the aggregated parameter  $\mathbf{w}(k)$ .

In the following, we call this procedure of multiple local SGD iterations followed by a fusion step a *round*. We use I to denote the number of local SGD iterations in each round. The main notations in this paper are listed in Table I.

It is possible that each round only involves a subset of clients, to avoid excessive delay caused by waiting for all the clients [42]. It has been shown that FedAvg converges even with random client participation, although the convergence rate is related to the degree of such randomness [23].

**Model pruning.** In the iterative training and pruning approaches for the centralized machine learning setting, the model is first trained using SGD for a given number of iterations [13], [15], [16]. Then, a certain percentage (referred to as the pruning rate) of weights that have smallest absolute values within each layer is removed (set to zero). This training and pruning process is repeated until a desired model size is reached. The benefit of this approach is that the training and pruning occurs at the same time, so that a trained model with a desired (small) size can be obtained in the end. However, existing pruning techniques require the availability of training data at a central location, which is not applicable to FL.

# IV. PRUNEFL

Our proposed PruneFL approach includes two stages: initial pruning at a selected client and further pruning involving both the server and clients during the FL process. The initial pruning can be done with biased data at a single client that has a relatively high computational capability, and the further pruning stage will "remove" the bias and refine the model. We use *adaptive pruning* in both stages. The illustration of the overall procedure is presented in Fig. 1. In the following, we introduce the two pruning stages (Section IV-A) and adaptive pruning (Section IV-B).

# A. Two-stage Distributed Pruning

**Initial pruning at a selected client.** Before FL starts, the system selects a single client to prune the model using its local data. This is important for two reasons. First, it allows us to start the FL process with a small model, which can significantly reduce the computation and communication time of each FL round. Second, when clients have heterogeneous computational capabilities, the selected client for initial pruning can be one that is powerful and trusted, so that the time required for initial pruning is short. We apply the adaptive pruning procedure that we describe in Section IV-B, where we adjust the original model iteratively while training the model on the selected client's local dataset, until the model size remains almost unchanged.

Algorithm 1: Adaptive pruning

1 for k = 0..., K - 1 do Initialize the set of importance measure on each 2 client:  $\mathcal{Z}_n \leftarrow \emptyset, \forall n;$ 3 for each client n, in parallel do Compute stochastic gradient 4  $\mathbf{g}_n(\mathbf{w}'_n(k)) := \mathbf{g}_n(\mathbf{w}_n(k) \odot \mathbf{m}(k));$ Update local parameters: 5  $\mathbf{w}_n(k+1) \leftarrow \mathbf{w}'_n(k) - \eta \mathbf{g}_n(\mathbf{w}'_n(k)) \odot \mathbf{m}(k);$ Add importance measure  $\mathbf{z}_n$  to  $\mathcal{Z}_n$ : 6  $\mathbf{z}_n := \mathbf{g}_n(\mathbf{w}'_n(k)) \odot \mathbf{g}_n(\mathbf{w}'_n(k));$  $\mathcal{Z}_n \leftarrow \mathcal{Z}_n \cup \mathbf{z}_n;$ if  $I \mid k+1$  then 7 Each client n sends  $\mathbf{w}_n(k+1)$  to the server; 8 Server aggregates the parameters from each 9 client:  $\mathbf{w}(k+1) \leftarrow \sum_{n=1}^{N} p_n \mathbf{w}_n(k+1);$ if k + 1 is reconfiguration iteration then 10 Each client sends the averaged importance 11 measure  $\overline{\mathbf{z}}_n := \left(\sum_{\mathbf{z}_n \in \mathcal{Z}_n} \mathbf{z}_n\right) / |\mathcal{Z}_n|$  to the Server aggregates the received importance 12 measure:  $\mathbf{z} \leftarrow \sum_{n=1}^{N} p_n \overline{\mathbf{z}}_n$ ; Reconfigure using Algorithm 2: 13  $\mathbf{w}'(k+1), \mathbf{m}(k+1) \leftarrow$ reconfigure( $\mathbf{w}(k+1), \mathbf{z}$ ); Reset:  $\mathcal{Z}_n \leftarrow \emptyset, \forall n;$ 14 else 15 No reconfiguration:  $\mathbf{w}'(k+1) \leftarrow \mathbf{w}(k+1)$ ; 16 Server sends new parameters to each client: 17  $\mathbf{w}'_n(k+1) \leftarrow \mathbf{w}'(k+1), \forall n;$ 

**Further pruning during FL process.** The model produced by initial pruning may not be optimal, because it is obtained based on data at a single client. However, it is a good starting point for the FL process involving all clients. During FL, we perform further adaptive pruning together with the standard FedAvg procedure, where the model can either grow or shrink depending on which way makes the training most efficient. In this stage, data from all participating clients are involved.

## B. Adaptive Pruning

For our adaptive pruning method, the notion of pruning broadly includes both removing and adding back parameters. Hence, we also refer to such pruning operations as *reconfiguration*. We reconfigure the model at a given interval of multiple iterations. For initial pruning, the reconfiguration interval can be any number of local iterations at the selected client. For further pruning, reconfiguration is done at the server after receiving parameter updates from clients (i.e., at the boundary between two rounds), and the reconfiguration interval in this case is always an integer multiple of the number of iterations (i.e., I) in each round.

**Definitions.** Let k denote the iteration index,  $\mathbf{g}_n(\mathbf{w}(k))$  denote the stochastic gradient of  $F_n(\mathbf{w}(k))$  evaluated at  $\mathbf{w}(k)$ 

and computed on the full parameter space on client n. Also, let  $\mathbf{m}_{\mathbf{w}}(k)$  denote a mask vector that is zero if the corresponding component in  $\mathbf{w}(k)$  is pruned and one if not pruned, and  $\odot$  denote the element-wise product.

In each reconfiguration step, adaptive pruning finds an optimal set of remaining (i.e., not pruned) model parameters. Then, parameters are pruned or added back accordingly, and the resulting model and mask are used for training until the next reconfiguration step. This procedure is illustrated in Algorithm 1, where  $a \mid b$  denotes that a divides b, i.e., b is an integer multiple of a. More details are given in Appendix C.

Our goal is to find the subnetwork that learns the "fastest". We do so by estimating the empirical risk reduction divided by the time required for completing an FL round, for any given subset of parameters chosen to be pruned. Note that at the beginning of each round (after averaging the clients' parameters), all clients start with the same parameter vector  $\mathbf{w}$ , i.e.,  $\mathbf{w}_n(k) = \mathbf{w}(k)$  for all n if a new round starts at iteration k (see Section III). For approximation purpose, we first consider the change of empirical risk after *one SGD iteration* starting with a *common* parameter  $\mathbf{w}(k)$ , for both initial and further pruning stages. The full FL procedure will be considered later in Theorem 2.

When the model is reconfigured at the end of iteration k, parameter update in the next iteration will be done on the reconfigured parameter  $\mathbf{w}'(k)$ , so we have an SGD update step that follows:

$$\mathbf{w}(k+1) = \mathbf{w}'(k) - \eta \sum_{n=1}^{N} p_n \mathbf{g}_n(\mathbf{w}'(k)) \odot \mathbf{m}_{\mathbf{w}'}(k)$$
$$= \mathbf{w}'(k) - \eta \mathbf{g}(\mathbf{w}'(k)) \odot \mathbf{m}(k)$$
(2)

where  $\eta$  is the learning rate. For simplicity, we define  $\mathbf{g}(\mathbf{w}'(k)) := \sum_{n=1}^{N} p_n \mathbf{g}_n(\mathbf{w}'(k))$ , and we omit the subscript  $\mathbf{w}'$  of  $\mathbf{m}$  in the following when it is clear from the context.

Let  $\mathcal{M}$  denote the index set of components that are not pruned, which corresponds to the indices of all non-zero values of the mask  $\mathbf{m}(k)$ .

**Empirical risk reduction.** To analyze the empirical risk reduction, we use a first-order approximation, which is a common practice in the literature [18], [43], [44]. We have

$$F(\mathbf{w}(k+1))$$

$$\approx F(\mathbf{w}'(k)) + \langle \nabla F(\mathbf{w}'(k)), \mathbf{w}(k+1) - \mathbf{w}'(k) \rangle$$
(3)

$$= F(\mathbf{w}'(k)) - \eta \langle \nabla F(\mathbf{w}'(k)), \mathbf{g}(\mathbf{w}'(k)) \odot \mathbf{m}(k) \rangle \quad (4)$$

$$\approx F'(\mathbf{w}(k)) - \eta \|\mathbf{g}(\mathbf{w}'(k)) \odot \mathbf{m}(k)\|^2$$
(5)

where  $\langle \cdot, \cdot \rangle$  is the inner product, (3) is from Taylor expansion, (4) is because of (2), and (5) is obtained by using the stochastic gradient to approximate the actual gradient, i.e.,  $\mathbf{g}(\mathbf{w}'(k)) \approx$  $\nabla F(\mathbf{w}'(k))$ . Then, the approximate decrease of empirical risk after the SGD step (2) is:

$$F(\mathbf{w}'(k)) - F(\mathbf{w}(k+1)) \approx \eta \|\mathbf{g}(\mathbf{w}'(k)) \odot \mathbf{m}(k)\|^2$$
$$\propto \|\mathbf{g}(\mathbf{w}'(k)) \odot \mathbf{m}(k)\|^2$$
$$= \sum_{j \in \mathcal{M}} g_j^2 =: \Delta(\mathcal{M})$$
(6)

where  $g_j$  is the *j*-th component of  $\mathbf{g}(\mathbf{w}'(k))$  and we define the set function  $\Delta(\mathcal{M})$  in the last line. The learning rate  $\eta$ is omitted since it is independent to the relative importance between parameter components, no matter if it is constant or varying. We use  $\Delta(\mathcal{M})$  as the *approximate risk reduction*, where we ignore the proportionality coefficient because our optimization problem is independent of the coefficient. As  $\Delta(\mathcal{M})$  is defined as the sum of  $g_j^2$  in (6), we use  $g_j^2$  as the *importance measure* for the *j*-th component of the parameter vector.

*Remark.* During the further pruning stage, the stochastic gradient  $\mathbf{g}(\mathbf{w}'(k))$  is the aggregated stochastic gradient from clients in FL. Since clients cannot compute  $\mathbf{g}(\mathbf{w}'(k))$  before receiving  $\mathbf{w}'(k)$  from the server, they compute  $\mathbf{g}(\mathbf{w}(k))$  and we use  $\mathbf{g}(\mathbf{w}'(k)) \approx \mathbf{g}(\mathbf{w}(k))$ , both of which are denoted by  $\mathbf{g}(\mathbf{w}'(k))$  with components  $\{g_j\}$  in the following. The additional overhead for clients to compute and transmit gradients on the full parameter space in a reconfiguration is small because pruning is done once in many FL rounds (the interval between two reconfigurations is 50 rounds in our experiments). Further details are given in Appendix C.

**Time of one FL round.** We define the (approximate) time of one FL round when the model has remaining parameters  $\mathcal{M}$  as a set function  $T(\mathcal{M}) := c + \sum_{j \in \mathcal{M}} t_j$ , where  $c \ge 0$  is a fixed constant and  $t_j > 0$  is the time corresponding to the *j*-th parameter component. Note that this is a linear function which is sufficient according to our empirical observations (see Appendix D1). In particular, the quantity  $t_j$  has a value that can be dependent on the neural network layer, and *c* captures a constant system overhead. From our experiments, we observed that  $t_j$  remains the same for all *j* that belong to the same neural network layer. Therefore, we can estimate the quantities  $\{t_j\}$ and *c* by measuring the time of one FL round for a small subset of different model sizes, before the overall pruning and FL procedure starts. An extension to the general case with non-linear  $T(\mathcal{M})$  is also discussed in Appendix A.

**Optimization of reconfiguration.** We would like to find the set of remaining parameters  $\mathcal{M}$  that maximizes the empirical risk reduction per unit training time. However,  $\Delta(\mathcal{M})$  only captures the risk reduction in the *next* SGD step when starting from the reconfigured parameter vector  $\mathbf{w}'(k)$ , as defined in (6). It does not capture the change in empirical risk when using  $\mathbf{w}'(k)$  instead of the original parameter vector  $\mathbf{w}(k)$  before reconfiguration. In other words, in addition to maximizing  $\Gamma(\mathcal{M}) := \frac{\Delta(\mathcal{M})}{T(\mathcal{M})}$ , we also need to ensure that  $F(\mathbf{w}'(k)) \approx F(\mathbf{w}(k))$ .

To ensure  $F(\mathbf{w}'(k)) \approx F(\mathbf{w}(k))$  after reconfiguration, we define an index set  $\overline{\mathcal{P}}$  to denote the parameters that *are* not allowed to be pruned. Usually,  $\overline{\mathcal{P}}$  includes parameters whose magnitudes are larger than a certain threshold, because pruning them can cause  $F(\mathbf{w}'(k))$  to become much larger than  $F(\mathbf{w}(k))$ . Among the remaining parameters that *can* be pruned (or added back if they are already pruned before), denoted by  $\mathcal{P}$ , we find which of them to prune to maximize  $\Gamma(\mathcal{M})$ . This yields the following optimization problem:

 Algorithm 2: Solving (7)

 Input : importance measure  $g_j^2$  and time coefficient  $t_j$ , for each parameter index j 

 Output: the optimal subset of parameters  $\mathcal{A}$  

 1  $\mathcal{A} \leftarrow \emptyset$ ;

 2  $\mathcal{S} \leftarrow \arg \operatorname{sort}_{j \in \mathcal{P}} \frac{g_j^2}{t_j}$ ;

 3 for  $j \in \mathcal{S}$  do

 4

 if  $\frac{g_j^2}{t_j} \ge \Gamma \left(\mathcal{A} \cup \overline{\mathcal{P}}\right)$  then

 5
  $\lfloor \mathcal{A} \leftarrow \mathcal{A} \cup \{j\}$ ;

 6
 else

 7
  $\lfloor$  break;

 8
 return  $\mathcal{A}$ ;

where  $\mathcal{A}$  is the set of parameters in  $\mathcal{P}$  that remain (i.e., are not pruned). The final set of remaining parameters is then  $\mathcal{M} = \mathcal{A} \cup \overline{\mathcal{P}}$ . Note that  $\mathcal{P} \cup \overline{\mathcal{P}}$  is the set of all parameters in the original model.

The algorithm for solving (7) is given in Algorithm 2, where sorting is in non-increasing order and S is an ordered set that includes the sorted indices. In essence, this algorithm sorts the ratios of components in the sums of  $\Delta(\mathcal{M})$  and  $T(\mathcal{M})$ . When the individual ratio  $g_j^2/t_j$  is larger than the current overall ratio  $\Gamma$ , then adding j to  $\mathcal{A}$  increases  $\Gamma$ . The bottleneck of this algorithm is the sorting operation. Hence, the overall time complexity of this algorithm is  $O(|\mathcal{P}| \log |\mathcal{P}|)$ .

**Theorem 1.** We have  $\Gamma(\mathcal{A} \cup \overline{\mathcal{P}}) \geq \Gamma(\mathcal{A}' \cup \overline{\mathcal{P}})$ , where  $\mathcal{A}$  from Algorithm 2 and  $\mathcal{A}'$  is any subset of  $\mathcal{P}$  with  $\mathcal{A}' \neq \mathcal{A}$ .

Theorem 1 shows that the result obtained from our Algorithm 2 is a global optimal solution to (7).

**Convergence of adaptive pruning.** As adaptive pruning can both increase and decrease the model size over time, a natural question is whether the model parameter vector will converge to a fixed value in a regular FL procedure with  $I \ge 1$  local SGD iterations in each round. We study this problem in the following.

We first make the following minimal set of assumptions that are common in the literature [22], [23].

# Assumption 1.

(a) Smoothness:

$$\left\|\nabla F_{n}(\mathbf{w}_{1}) - \nabla F_{n}(\mathbf{w}_{2})\right\| \leq \beta \left\|\mathbf{w}_{1} - \mathbf{w}_{2}\right\|, \forall n, \mathbf{w}_{1}, \mathbf{w}_{2},$$

where  $\beta$  is a positive constant.

(b) Lipschitzness:

$$\left\|F(\mathbf{w}_{1}) - F(\mathbf{w}_{2})\right\| \leq L \left\|\mathbf{w}_{1} - \mathbf{w}_{2}\right\|, \forall \mathbf{w}_{1}, \mathbf{w}_{2},$$

where L is a positive constant.

(c) Unbiasedness:

$$\mathbb{E}\left[\mathbf{g}_{n}(\mathbf{w})\right] = \nabla F_{n}(\mathbf{w}), \forall n, \mathbf{w}$$

(d) Bounded variance:

$$\mathbb{E} \|\mathbf{g}_n(\mathbf{w}) - \nabla F_n(\mathbf{w})\|^2 \le \sigma^2, \forall n, \mathbf{w} .$$

## (e) Bounded divergence:

$$\left\|\nabla F(\mathbf{w}) - \nabla F_n(\mathbf{w})\right\|^2 \le \epsilon^2, \forall n, \mathbf{w}$$

where  $F(\mathbf{w}) := \sum_{n=1}^{N} p_n F_n(\mathbf{w})$  as defined in (1). (f) **Time independence in SGD**: The stochastic gradients

- f) Time independence in SGD: The stochastic gradients obtained in different iterations are independent from each other.
- (g) *Client independence:* The stochastic gradients obtained from different clients are always independent from each other, even in the same iteration.

**Theorem 2.** When Assumption 1 holds, and  $\eta \leq \frac{1}{2\sqrt{6}I\beta}$ , we have

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla F(\mathbf{w}'(k)) \odot \mathbf{m}_{\mathbf{w}'}(k) \right\|^{2} \\
\leq \frac{2(F_{0} - F^{*})}{\eta K} + \alpha \eta \beta \sigma^{2} + 4\beta^{2} \left( (1 - \alpha) I \sigma^{2} + 3I^{2} \epsilon^{2} \right) \eta^{2} \\
+ \frac{2L}{\eta K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \mathbf{w}(k) - \mathbf{w}'(k) \right\| ,$$
(8)

where  $\alpha := \sum_{n=1}^{N} p_n^2$ ,  $F_0 := F(\mathbf{w}(0))$ ,  $F^* := \min_{\mathbf{w}} F(\mathbf{w})$ , and I is the number of iterations per round.

Under some conditions, this convergence can be bounded asymptotically by  $\mathcal{O}\left(\frac{1}{\sqrt{NK}}\right) + \mathcal{O}\left(\frac{1}{K}\right)$ , achieving linear speedup<sup>1</sup> with the number of clients N [31], [45] for sufficiently large K. See Appendix B2 for more details.

If we do not reconfigure in an iteration k, we have  $\mathbf{w}(k) =$  $\mathbf{w}'(k)$ . In the right-hand side (RHS) of (8), the first two terms go to zero as  $K \to \infty$ . The last term is related to how well  $\mathbf{w}'$  approximates  $\mathbf{w}$  after pruning. To ensure that the sum in the last term grows slower than  $\sqrt{K}$ , the number of *non-zero* prunable parameters (which belong to  $\mathcal{P}$ ) should decrease over time. Note that we consider all zero parameters to be prunable and they also belong to  $\mathcal{P}$ , thus the size of  $\mathcal{P}$  itself may not decrease over time. This convergence result shows that the gradient components corresponding to the remaining (i.e., not pruned) parameters vanishes over time, which suggests that we will get a "stable" parameter vector in the end, because when the gradient norm is small, the change of parameters in each iteration is also small. In addition to gradient convergence on the subspace after pruning as suggested in Theorem 2, our experiments show that our pruned model also converges to an accuracy close to that of the full-sized model.

**Tracking a small model.** By choosing the size of  $\overline{P}$  properly over time, our adaptive pruning algorithm can keep reducing the model size as long as such reduction does not adversely impact further training. Intuitively, the model that we obtain from this process is one that has a small size while maintaining full "trainablity" in future iterations. Parameter components for which the corresponding gradient components remain zero (or close to zero) will be pruned.

In cases where a target maximum model size should be reached at convergence (e.g., for efficient inference later), we can also enforce a maximum size constraint in each reconfiguration that starts with the full size and gradually decreases to the target size as training progresses, which allows the model to train quickly in initial rounds while converging to the target size in the end.

## V. IMPLEMENTATION

## A. Using Sparse Matrices

Although the benefit of model pruning in terms of computation is constantly mentioned in the literature from a theoretical point of view [13], most existing implementations substitute sparse parameters by applying binary masks to dense parameters. Applying masks increases the overhead of computation, instead of reducing it. We implement sparse matrices for model pruning, and we show its efficacy in our experiments. We use dense matrices for full-sized models, and sparse matrices for weights in both convolutional and fully-connected layers in pruned models.

#### B. Complexity Analysis

Storage, memory, and communication. We implement two types of storage for sparse matrices: bitmap and valueindex tuple. Bitmap uses one extra bit to indicate whether the specific value is zero. For 32-bit floating point parameter components, bitmap incurs 1/32 extra storage and communication overhead. Value-index tuple stores the values and both row and column indices of all non-zero entries. In our implementation, we use 16-bit integers to store row and column indices and 32-bit floating point numbers to store parameter values. Since each parameter component is associated with a row index and a column index, the storage and communication overhead doubles compared to storing the values only. We dynamically choose between the two ways of storage, and thus, the ratio of the sparse parameter size to the dense parameter size is  $\min \{2 \times d, \frac{1}{32} + d\}$ , where d is the model's *density* (percentage of non-zero parameters). This ratio is further optimized when the matrix sparsity pattern is fixed (in most FL rounds, see Appendix C). In this case, there is no extra cost since only values of the non-zero entries need to be exchanged.

**Computation.** Because dense matrix multiplication is extremely optimized, sparse matrices will show advantage in computation time only when the matrix is below a certain density, where this density threshold depends on specific hardware and software implementations. In our implementation, we choose either dense or sparse representation depending on which one is more efficient. The complexity (computation time) of the matrix multiplication between a sparse matrix S and a dense matrix D is linear to the number of non-zero entries in S (assuming D is fixed).

#### C. Implementation Challenges

As of today, well-known machine learning frameworks have limited support of sparse matrix computation. For instance, in PyTorch version 1.6.0, the persistent storage of a matrix in

<sup>&</sup>lt;sup>1</sup>The dominant term is  $\mathcal{O}\left(\frac{1}{\sqrt{NK}}\right)$  when K is sufficiently large. The notion of linear speedup means that, to reach the same error bound, the total number of rounds K can proportionally decrease as N increases.

sparse form takes  $5\times$  space compared to its dense form; the computations on sparse matrices are slow; and sparse matrices are not supported for the kernels in convolutional layers, etc. To benefit from using sparse matrices in real systems, we extend the PyTorch library by implementing a more efficient sparse storage, and the support for sparse convolutional kernels. We only partially improve backward passes due to implementation limitations (more details in Appendix C2). This problem, however, can be improved in the future by implementing and further optimizing efficient sparse matrix multiplication on low-level software, as well as developing specific hardware for this purpose. Nevertheless, the novelty in our implementation is that we use sparse matrices in both fully-connected and convolutional layers in the pruned model.

## VI. EXPERIMENTS

In this section, we present the experimental setup and results.

**Datasets.** We evaluate PruneFL on four image classification tasks:

- (a) Conv-2 model on FEMNIST [46],
- (b) VGG-11 model [47] on CIFAR-10 [48],
- (c) ResNet-18 model [49] on ImageNet-100 [50],
- (d) MobileNetV3-Small model [51] on CelebA [46],

all of which represent typical FL tasks. Due to practical considerations of edge devices' training time and storage capacity, we select data corresponding to 193 writers for FEMNIST, and the first 100 classes of the ImageNet dataset (referred to as ImageNet-100). We adapt some layers in VGG-11, ResNet-18 and MobileNetV3-Small to match with the number of output labels in our datasets.

When using full client participation, because we only have 10 clients in total, for FEMNIST, we partition all the 193 writers' images into 10 clients (the first 9 clients each has 19 writers' images and the last client has 22 writers' images). For CelebA, we partition all the 9,343 persons' images into 10 clients (the first 9 clients each has 934 persons' images and the last client has 937 persons' images). Note that such partitioning is still non-IID.

**Model architectures.** The architecture details are presented in Table C.1 in the appendix. VGG-11 ResNet-18, and MobileNetV3-Small are well-known architectures, and we directly acquire Conv-2 from its original work [46].

**Platform.** To study the performance of our proposed approach, we conduct experiments in (i) a real edge computing prototype, where a personal computer serves as both the server and a client, and the other clients are Raspberry Pi devices, and (ii) a simulated setting with multiple clients and a server, where computation and communication times are obtained from measurements involving either Raspberry Pi devices or Android phones.

Unless otherwise specified, the prototype system includes nine Raspberry Pi (version 4, with 2 GB RAM, 32 GB SD card) devices as clients and a personal computer without GPU as both a client and the server (totaling 10 clients). Three of the Raspberry Pis use wireless connections and the remaining six use wired connections. The communication speed is stable and is approximately 1.4 MB/s. The simulated system uses the same setting as in the prototype. We use time measurements from Raspberry Pis, except for the ImageNet-100 dataset where we replace the computation time by measurements from Android virtual machine (VM).

We consider FL with full client participation in the main paper and present results with random client selection [42] in Appendix D2. The results are similar. FEMNIST and CelebA data are partitioned into clients in a non-IID manner according to writer/person identity, and CIFAR-10 and ImageNet-100 are partitioned into clients in an IID manner.

**Baselines.** We compare the test accuracy vs. *time* curve of PruneFL with five baselines: (i) conventional FL [2], (ii) iterative pruning [13], (iii) online learning [9], (iv) SNIP [18], and (v) SynFlow [17]. Because iterative pruning and SNIP cannot automatically determine the model size, we consider an enhanced version of these baselines that obtain the same model size as PruneFL at convergence. Additional baselines are also considered in Section VI-B.

Since our experiments try to minimize the training time using pruning, there is no direct way of comparing with the baselines that either are not specifically designed for pruning (the online learning baseline) or do not adapt the pruned model size (all other baselines). We compare with the baselines as follows. In every round, the online learning approach produces a model size for the next round, and we adjust the model accordingly while keeping each layer's density the same. To compare with SNIP, after the first round, we let SNIP prune the original model in a one-shot manner to the same density as the final model found by our adaptive pruning method, and keep the architecture afterwards. Similarly, to compare with SynFlow, we let SynFlow prune the model (before training) to the same density as the final model found by our adaptive pruning method, and keep the architecture afterwards. To compare with iterative pruning, we let the model be pruned with a fixed rate for 20 times (at an equal interval) in the first half of the total number of rounds, such that the remaining number of parameter components equals that of the model found by our adaptive pruning method, and the pruning rate is equal across layers. See Fig. 9 for the illustration of the baseline settings.

**Pruning configurations.** The initial pruning stage is done on the personal computer client. We end the initial pruning stage either when the model size is "stable", or when it exceeds certain maximum number of iterations. We consider the model size as "stable" when its relative change is below 10% for 5 consecutive reconfigurations.

For adaptive pruning, to ensure convergence of the last term on the RHS of (8) in Theorem 2, we exponentially decrease the number of *non-zero* prunable parameters in  $\mathcal{P}$  over rounds. We note that  $\mathcal{P}$  includes both zero and non-zero parameters, hence the size of  $\mathcal{P}$  itself may not decrease. For a given size of  $\mathcal{P}$ , the  $|\mathcal{P}|$  parameters with the smallest magnitude belong to  $\mathcal{P}$  that can be pruned (or added back), and the rest belong to  $\overline{\mathcal{P}}$  that cannot be pruned.

Biases (if any) in the DNNs are not pruned. In ResNet-18, BatchNorm layers and downsampling layers are not pruned since the number of parameters in such layers is negligible

 TABLE II

 EVALUATION CONFIGURATIONS (C.S. STANDS FOR CLIENT SELECTION; LR STANDS FOR LEARNING RATE).

Dataset	FEMNIST	CIFAR-10	ImageNet-100	CelebA
SGD params in round r	LR = 0.25	$LR = 0.1 \cdot 0.5 \frac{r}{10000}$	$LR = 0.05 \cdot 0.5^{\left\lfloor \frac{r}{1000} \right\rfloor \cdot 0.1}$	LR = 0.2
Fraction of non-zero prunable parameters in round r	$0.3\cdot 0.5rac{r}{10000}$	$0.3\cdot 0.5rac{r}{10000}$	$0.3\cdot 0.5rac{r}{10000}$	$0.3\cdot 0.5rac{r}{10000}$
Number of data samples used in initial pruning	200	200	500	500
Number of clients (Non-C.S., C.S.)	10, 193	10, 100	10, 100	10, 934
Mini-batch size, local iterations in each round	20, 5	20, 5	20, 5	20, 5
Reconfiguration	every 50 rounds	every 50 rounds	every 50 rounds	every 50 rounds
Total number of FL rounds	10,000	10,000	20,000	1,000
Evaluation	prototype (Pi 4), simulation (Pi 4)	simulation (Pi 4)	simulation (Android VM)	simulation (Pi 4)



Fig. 2. Training time on Raspberry Pi 4 (FEMNIST).



Fig. 3. Inference time on Raspberry Pi 4 (FEMNIST).

compared to the size of convolutional and fully-connected layers.

Lottery ticket analysis. To verify whether the final model from adaptive pruning is a lottery ticket [11], [14], we reinitialize this converged model using the original random seed, and compare its accuracy vs. *round* curve with (i) conventional FL, (ii) random reinitialization (same architecture as the lottery ticket but initialized with a different random seed), (iii) SNIP, and (iv) SynFlow.

**Hyperparameters.** The hyperparameters above are chosen empirically only with coarse tuning by experience. We observe that our and other methods are insensitive to these hyperparameters. Hence, we do not perform fine tuning on any parameter.

The detailed evaluation configurations are given in Table II.

#### A. Time Measurement

We present the time measurements of one FL round on the prototype system to show the effectiveness of model pruning on edge devices. We implement the full-sized Conv-2 model



Fig. 4. Comparing conventional FL and PruneFL with both prototype and simulation results (FEMNIST).

in dense form as well as the pruned models in sparse form at different densities, and measure the average elapsed time of FL on these pruned models involving both the server and clients over 10 rounds.

Fig. 2 shows the average total time, computation time, and communication time in one round as we vary the model density. Note that the model is in dense form at 100% on the *x*-axis and sparse form elsewhere. We also plot the actual size of the parameters that are exchanged between server and clients in this figure for one FL round.

**Computation time.** We see from Fig. 2 that as the model density decreases, the computation time (for five local iterations) decreases from 11.24 seconds per round to 6.34 seconds per round. This reduction in computation time is moderate since our implementation of sparse computation only partially improves backward passes (see Section V). Additionally, we plot in Fig. 3 the total inference time and the number of floating-point operations (FLOPs) for 200 data samples (see Appendix C3 for details of FLOPs computation). The inference time result shows a similar trends as in Fig. 2, and the number of FLOPs keeps decreasing as we reduce the model size.

**Communication time.** Our implementation of sparse matrices reduces the storage requirement significantly (see Section V). Compared to computation time, the decrease in the communication time is more noticeable. It drops from 35.88 seconds per round to 1.04 seconds per round.

**Enabling FL on low-power edge device.** In addition, we ran experiments with the LeNet-300-100 [52] architecture, and we observed that when training the MNIST [52] dataset on the



Fig. 5. Test accuracy vs. time results of four datasets.



Fig. 6. Test accuracy vs. accumulated FLOPs per client (FEMNIST).

full-sized, dense-form LeNet-300-100 model on Raspberry Pi version 3 (with 1 GB RAM, 32 GB SD card), the system dies during the first mini-batch due to resource exhaustion, while the models in sparse form can be trained. Thus, our approach of using sparse models enables model training on low-power edge devices, which is otherwise impossible on Raspberry Pi 3 in this experiment due to the device's resource limitation.

#### B. Training Cost Reduction

In the following, we study PruneFL's cost reduction in terms of both time and FLOPs for training.

**Comparing conventional FL and PruneFL.** Fig. 4 shows the test accuracy vs. time results on both the prototype and simulated systems, for Conv-2 on FEMNIST. The time for initial pruning of PruneFL is included in this figure, which is negligible (it takes less than 500 seconds) compared to the further pruning stage. We see that PruneFL outperforms conventional FL by a significant margin. Since the prototype and simulation results match closely, we present the simulation results in subsequent experiments due to their excessive training time on the prototype system.

**Training time reduction.** In Fig. 5, we compare the test accuracy vs. time results for all datasets, models, and baselines. Clearly, PruneFL demonstrates a consistent advantage on training speed over baselines. Moreover, PruneFL always converges to similar accuracy achieved by conventional FL (see Appendix D3). Other methods may have suboptimal performance, e.g., SNIP does not converge to conventional FL's accuracy with CIFAR-10. We also observe that some approaches such as online learning and SynFlow in Fig. 5(b) always stay at the random guess accuracy. The reason could be that such approaches are unstable and can get stuck in local optimal points at the beginning of training.



(c) ResNet-18 on ImageNet-100

(d) MobileNetV3-Small on CelebA

TABLE III TIME AND ACCUMULATED FLOPS PER CLIENT TO REACH TARGET ACCURACY (FEMNIST).

Approach	Time (FLOPs) to reach	Time (FLOPs) to reach		
	70% accuracy	80% accuracy		
Conventional FL	17,929 s (3.5 TFLOPs)	52,153 s (10.5 TFLOPs)		
PruneFL (ours)	3,187 s (1.6 TFLOPs)	15,009 s (6.8 TFLOPs)		
SNIP	6,801 s (3.3 TFLOPs)	22,467 s (11.7 TFLOPs)		
SynFlow	7,132 s (3.6 TFLOPs)	22,327 s (12.3 TFLOPs)		
Online	18,042 s (3.5 TFLOPs)	54,593 s (10.7 TFLOPs)		
Iterative	17,495 s (3.5 TFLOPs)	46,521 s (10.1 TFLOPs)		

**Training FLOPs reduction.** Although we observe that the training time on Raspberry Pis is relatively consistent across different models and tasks, there are still factors that can affect the training time (e.g., environment temperature). To further validate our approach's advantage on accelerating training, we present the results on test accuracy vs. accumulated FLOPs per client for FEMNIST in Fig. 6. We find that this result shares the same characteristics with Fig. 5(a) in terms of acceleration (we present only one set of results here due to the similarity).

Time and FLOPs to reach target accuracy. Table III lists the time and accumulated FLOPs per client that an algorithm first reaches a certain accuracy with FEMNIST. PruneFL takes less than 1/3 of time compared to conventional FL to reach 80% accuracy, and it also saves more than 33% of time (more than 2 hours) compared to SNIP and SynFlow. The savings of FLOPs are similar.

**Comparing with additional baselines.** To avoid bottlenecks, our algorithm and implementation ensures that all components in PruneFL, including communication, computation, and reconfiguration, are orchestrated and inexpensive. For this reason, some approaches in the literature that are not specifically designed for the edge computing environment with low-power devices may perform poorly if applied to our system setup, as we illustrate next.

Considering computation time, PruneTrain [21] applies regularization on every input and output channel in every layer. When the same regularization is applied to our system, we find that the computation time (using FEMNIST and Conv-2) takes 17.65 seconds per round, which is a 57% increase compared to PruneFL.

Considering communication time, dynamic pruning with feedback (DPF) [19] maintains a full-sized model, and clients have to upload full-sized gradients to the server (but only download a subset of model parameters) in every round.



(a) Conv-2 on FEMNIST(b) VGG-11 on CIFAR-10Fig. 7. Lottery ticket results of four datasets.



Fig. 8. PruneFL with only one pruning stage (FEMNIST).

Thus, assuming unit model size and model density d, the communication cost per round, including both uploading and downloading, is 1 + d. In comparison, clients in PruneFL only upload the full-sized model to the server at a reconfiguration round (every 50 rounds in our experiments), and always exchange pruned models otherwise. This gives an average cost of  $\frac{(1+d)+2\times49d}{50} = 0.02 + 1.96d$  including both uploading and downloading. For instance, when the model density is 10%, DPF incurs  $5.1 \times$  communication cost compared to PruneFL. Finally, our reconfiguration algorithm (Algorithm 2) runs in quasi-linear time, making it possible to be implemented on edge devices.

#### C. Finding a Lottery Ticket

Unlike some existing pruning techniques such as SNIP [18], dynamic pruning [19], and SynFlow [17], PruneFL finds a lottery ticket (although not necessarily the smallest). In Fig. 7, FL with the reinitialized pruned model obtained from PruneFL learns comparably fast as FL with the original model, in terms of test accuracy vs. *round*, confirming that they are lottery tickets. In Fig. 7(b), the Random Reinit curve stays at the random guess accuracy. This is not surprising since the lottery ticket, i.e., the final pruned model found by PruneFL, needs to be reinitialized to their original values to learn comparably fast as the full-sized model [11]. When reinitialized with different values, the training of the "lottery ticket" can be suboptimal. In this experiment, it is stuck at the beginning of training.

Fig. 8 compares the test accuracy vs. round curves of PruneFL with alternative methods that either only includes initial pruning (at a single client) or only includes further pruning (during FL). It shows that the model obtained from the initial pruning stage does not converge to the optimal accuracy, and only performing further pruning without initial pruning results in a slower learning speed. PruneFL with both stages



(c) ResNet-18 on ImageNet-100

(d) MobileNetV3-Small on CelebA

avoids drawbacks from methods that includes only one stage. Furthermore, the model obtained from initial pruning is not a lottery ticket of the original model, while PruneFL with only further pruning or both pruning stages find a lottery ticket.

Therefore, one can view PruneFL as a two-stage procedure to find a lottery ticket of the given model, which is in line with our claims in Section IV. The ability that we can find lottery tickets is useful when we need to retrain a pruned model on slightly different but similar datasets [14].

## D. Model Size Adaptation

An illustration of the change in model size is shown in Fig. 9. The small negative part of the x-axis shows the initial pruning stage which is unique to PruneFL. Since there is no notion of "round" in the initial pruning stage, we consider five local iterations in this stage as one round, which is consistent with our FL setting. Conventional FL always keeps the full model size. The model sizes provided by online learning is unstable. It fluctuates in initial rounds due to its exploration. SNIP and SynFlow prune the initial model to the target size in a one-shot manner at the beginning of training. Iterative pruning gradually reduces the model size until reaching the target. We notice that PruneFL also discovers the degree of overparameterization. Empirically, Conv-2, an overparameterized model for FEMNIST, converges to a small density (13.4%), while ResNet-18, an underparameterized model for ImageNet-100, converges to a density of around 67.7%.

It is worth mentioning that finding a proper target density for pruning is non-trivial. Usually, foresight pruning methods such as [18] and [43] prune the model to the (manually selected) density before training. Fig. 10 shows two cases where we use SNIP to prune Conv-2 (with FEMNIST) to 30% and 1%, the training speed becomes slower, and if the density is too small (1%), the sparse model cannot converge to the same accuracy as the original model. In comparison, PruneFL automatically determines a proper density.

# E. Training with Limited/Targeted Model Sizes

There are cases where a hard limit on the maximum model size or a targeted final model size (or both) is desired. For example, if some of the client devices have limited memory or storage so that only a partial model can be loaded, then the model size must be constrained after initial pruning. Targeted model size may be needed in the case where the goal of the FL system is to obtain a model with a certain small size at the end of training.



Fig. 9. Number of parameters vs. round for four datasets.



Fig. 10. SNIP with different densities (FEMNIST).



Fig. 11. Training with limited and targeted size (FEMNIST).

Next, we present an extended PruneFL with limited and targeted model sizes, and show that with reasonable constraints, PruneFL still achieves good results. We use a heuristic way to limit the model size: we stop Algorithm 2 early when the number of remaining parameters reaches the maximum allowed size, and we schedule the maximum size of the model to decrease linearly. Assuming  $d_l$  is the density limit,  $d_t$  is the target model density at the end of the further pruning stage  $(d_t \leq d_l)$ , and PruneFL is run for  $r_{\text{max}}$  rounds after the initial pruning stage, then the maximum density at round r is  $d_{\max}(r) = (r \cdot d_t + (r_{\max} - r) \cdot d_l)/r_{\max}$ . The results of selecting  $d_l = 10\%$  and  $d_t = 5\%$  for Conv-2 on FEMNIST are given in Fig. 11. We see that if we do not impose these model size constraints, PruneFL exceeds the density limits  $d_l = 10\%$  and  $d_t = 5\%$  defined in this example, and obtains a model that is much larger than the target density  $d_t = 5\%$ at the end of training. We see that PruneFL with effective size limit and target still achieves fast convergence and similar convergence accuracy, and the model size is always limited below the threshold  $d_l = 10\%$  and reaches the target density  $d_t = 5\%$  at the end of training.



(c) ResNet-18 on ImageNet-100

(d) MobileNetV3-Small on CelebA

TABLE IV Density of each layer at convergence (no C.S ).

Experiment	VGG-11 on CIFAR-10			
Convolutional	1.0, 0.80, 0.84, 0.84, 0.50, 0.07, 0.01, 0.03			
Fully-connected	0.14,  0.15,  0.7			

#### F. Relative Importance Between Layers

Similar to SynFlow [17] and SNIP [18], our algorithm discovers the relative importance between layers automatically. Taking CIFAR-10 on VGG-11 as an example, the densities of convolutional layers and fully-connected layers at convergence are listed in Table IV in a sequential manner. We observe that the input and output layers are not pruned to a low density, indicating that they are relatively important in the neural network architectures. This also agrees with the pruning scheme in [11], where the authors empirically set the pruning rate of the output layer to a small percentage or even to zero. Some large convolutional layers, such as the last two convolutional layers in VGG-11, are identified as redundant, and thus have small densities at convergence.

#### VII. CONCLUSION

We have proposed PruneFL for FL in edge/mobile computing environments, where the goal is to effectively reduce the size of neural networks so that resource-limited clients can train them within a short time. Our PruneFL method includes initial and further pruning stages, which improves the performance compared to only having a single stage. PruneFL also includes a low-complexity adaptive pruning method for efficient FL, which finds a desired model size that can achieve a similar prediction accuracy as the original model but with much less time. Our experiments on Raspberry Pi devices confirm that we improve the cost-efficiency of FL while obtaining a lottery ticket. Our method can be applied together with other compression techniques, such as quantization, to further reduce the communication overhead.

#### ACKNOWLEDGMENT

This work was partially supported by the U.S. Office of Naval Research under Grant N00014-19-1-2566, the U.S. National Science Foundation AI Institute Athena under Grant CNS-2112562, and the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of U.S. Office of Naval Research, the U.S. National Science Foundation, the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. V. Valls has also received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 795244.

#### REFERENCES

- I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.
- [2] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *AISTATS*, 2017.
- [3] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *arXiv preprint arXiv*:1908.07873, 2019.
- [4] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, "Wireless network intelligence at the edge," *Proceedings of the IEEE*, vol. 107, no. 11, pp. 2204–2239, 2019.
- [5] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Transactions on Intelligent Systems and Technology (TIST), vol. 10, no. 2, p. 12, 2019.
- [6] P. Kairouz, H. B. McMahan et al., "Advances and open problems in federated learning," arXiv preprint arXiv:1912.04977, 2019.
- [7] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," arXiv preprint arXiv:1610.05492, 2016.
- [8] S. Caldas, J. Konečny, H. B. McMahan, and A. Talwalkar, "Expanding the reach of federated learning by reducing client resource requirements," arXiv preprint arXiv:1812.07210, 2018.
- [9] P. Han, S. Wang, and K. K. Leung, "Adaptive gradient sparsification for efficient federated learning: An online learning approach," in *IEEE ICDCS*, 2020.
- [10] Z. Xu, Z. Yang, J. Xiong, J. Yang, and X. Chen, "Elfish: Resourceaware federated learning on heterogeneous edge devices," arXiv preprint arXiv:1912.01684, 2019.
- [11] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *ICLR*, 2019.
- [12] Y. LeCun, J. S. Denker, and S. A. Solla, "Optimal brain damage," in Advances in neural information processing systems, 1990, pp. 598–605.
- [13] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in Advances in neural information processing systems, 2015, pp. 1135–1143.
- [14] A. Morcos, H. Yu, M. Paganini, and Y. Tian, "One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers," in Advances in Neural Information Processing Systems, 2019, pp. 4933– 4943.
- [15] S. Narang, E. Elsen, G. Diamos, and S. Sengupta, "Exploring sparsity in recurrent neural networks," in *International Conference on Learning Representations*, 2017.
- [16] M. Zhu and S. Gupta, "To prune, or not to prune: exploring the efficacy of pruning for model compression," *arXiv preprint arXiv:1710.01878*, 2017.
- [17] H. Tanaka, D. Kunin, D. L. Yamins, and S. Ganguli, "Pruning neural networks without any data by iteratively conserving synaptic flow," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [18] N. Lee, T. Ajanthan, and P. Torr, "Snip: Single-shot network pruning based on connection sensitivity," in *International Conference* on Learning Representations, 2019. [Online]. Available: https: //openreview.net/forum?id=B1VZqjAcYX
- [19] T. Lin, S. U. Stich, L. Barba, D. Dmitriev, and M. Jaggi, "Dynamic model pruning with feedback," in *International Conference* on Learning Representations, 2020. [Online]. Available: https: //openreview.net/forum?id=SJem8ISFwB
- [20] S. Anwar, K. Hwang, and W. Sung, "Structured pruning of deep convolutional neural networks," ACM Journal on Emerging Technologies in Computing Systems (JETC), vol. 13, no. 3, pp. 1–18, 2017.

- [21] S. Lym, E. Choukse, S. Zangeneh, W. Wen, S. Sanghavi, and M. Erez, "Prunetrain: fast neural network training by dynamic sparse model reconfiguration," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–13.
- [22] H. Yu, S. Yang, and S. Zhu, "Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 5693–5700.
- [23] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of fedavg on non-iid data," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/forum?id=HJxNAnVtDS
- [24] J. Wang, S. Wang, R.-R. Chen, and M. Ji, "Local averaging helps: Hierarchical federated learning and convergence analysis," *arXiv preprint arXiv:2010.12998*, 2020.
- [25] J. Wang and G. Joshi, "Adaptive communication strategies to achieve the best error-runtime trade-off in local-update sgd," in *Machine Learning* and Systems (MLSys), 2019.
- [26] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1205–1221, June 2019.
- [27] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," arXiv preprint arXiv:1910.06378, 2019.
- [28] S. P. Karimireddy, Q. Rebjock, S. Stich, and M. Jaggi, "Error feedback fixes SignSGD and other gradient compression schemes," in *International Conference on Machine Learning*, vol. 97, Jun. 2019, pp. 3252– 3261.
- [29] D. Alistarh, T. Hoefler, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli, "The convergence of sparsified gradient methods," in *Advances in Neural Information Processing Systems*, 2018, pp. 5973– 5983.
- [30] S. Shi, K. Zhao, Q. Wang, Z. Tang, and X. Chu, "A convergence analysis of distributed sgd with communication-efficient gradient sparsification," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 2019, pp. 3411–3417.
- [31] P. Jiang and G. Agrawal, "A linear speedup analysis of distributed deep learning with sparse and quantized communication," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., 2018, pp. 2525–2536.
- [32] W. Du, X. Zeng, M. Yan, and M. Zhang, "Efficient federated learning via variational dropout," 2018.
- [33] A. Li, J. Sun, B. Wang, L. Duan, S. Li, Y. Chen, and H. Li, "Lotteryfl: Personalized and communication-efficient federated learning with lottery ticket hypothesis on non-iid datasets," arXiv preprint arXiv:2008.03371, 2020.
- [34] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep learning with limited numerical precision," in *International Conference* on *Machine Learning*, 2015, pp. 1737–1746.
- [35] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," in *Proceedings of the British Machine Vision Conference. BMVA Press*, 2014.
- [36] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," arXiv preprint arXiv:1503.02531, 2015.
- [37] J. Lin, Y. Rao, J. Lu, and J. Zhou, "Runtime neural pruning," in Proceedings of the 31st International Conference on Neural Information Processing Systems, 2017, pp. 2178–2188.
- [38] C. Hu, W. Bao, D. Wang, and F. Liu, "Dynamic adaptive dnn surgery for inference acceleration on the edge," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 1423– 1431.
- [39] P. Vepakomma, O. Gupta, T. Swedish, and R. Raskar, "Split learning for health: Distributed deep learning without sharing raw patient data," arXiv preprint arXiv:1812.00564, 2018.
- [40] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *proceedings of the 2017* ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1175–1191.
- [41] T. Li, M. Sanjabi, A. Beirami, and V. Smith, "Fair resource allocation in federated learning," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/ forum?id=ByexEISYDr

- [42] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konecny, S. Mazzocchi, H. B. McMahan *et al.*, "Towards federated learning at scale: System design," *arXiv preprint arXiv:1902.01046*, 2019.
- [43] C. Wang, G. Zhang, and R. Grosse, "Picking winning tickets before training by preserving gradient flow," in *International Conference on Learning Representations*, 2019.
- [44] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, "Pruning convolutional neural networks for resource efficient inference," arXiv preprint arXiv:1611.06440, 2016.
- [45] H. Yu, R. Jin, and S. Yang, "On the linear speedup analysis of communication efficient momentum sgd for distributed non-convex optimization," in *International Conference on Machine Learning*. PMLR, 2019, pp. 7184–7193.
- [46] S. Caldas, P. Wu, T. Li, J. Konecný, H. B. McMahan, V. Smith, and A. Talwalkar, "LEAF: A benchmark for federated settings," *CoRR*, vol. abs/1812.01097, 2018. [Online]. Available: http://arxiv.org/abs/1812. 01097
- [47] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.
- [48] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.
- [49] 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.
- [50] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in 2009 IEEE conference on computer vision and pattern recognition. Ieee, 2009, pp. 248–255.
- [51] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan *et al.*, "Searching for mobilenetv3," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1314–1324.
- [52] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner *et al.*, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [53] R. Tang, A. Adhikari, and J. Lin, "Flops as a direct optimization objective for learning sparse neural networks," arXiv preprint arXiv:1811.03060, 2018.

## APPENDIX

# A. Extension to Non-linear $T(\mathcal{M})$

For the case where  $T(\mathcal{M})$  is non-linear, but a general monotone and positive set function instead, we can still find a local optimal solution to (7) using Algorithm A.1. We can see that the complexity of Algorithm A.1 is  $O(|\mathcal{P}|^2)$ .

Algorithm A.1: Solving (7), general $T(\cdot)$	
1 $\mathcal{A} \leftarrow \emptyset;$	
2 $j^* \leftarrow$ None;	
3 repeat	
4   if $j^*$ is not None then	
5 $ \mid \mathcal{A} \leftarrow \mathcal{A} \cup \{j^*\}; $	
$6  j^* \leftarrow \arg\max_{j \in \mathcal{P} \setminus \mathcal{A}} \ \frac{g_j^2}{t_j(\mathcal{A} \cup \overline{\mathcal{P}})};$	
7 <b>until</b> $\frac{g_j^2}{t_j(\mathcal{A}\cup\overline{\mathcal{P}})} < \Gamma\left(\mathcal{A}\cup\overline{\mathcal{P}}\right);$	
s return $\mathcal{A}$ ;	// final result

**Theorem 3.** For general  $T(\mathcal{M})$ , we have  $\Gamma\left(\mathcal{A}\cup\overline{\mathcal{P}}\right) \geq \Gamma\left(\mathcal{A}'\cup\overline{\mathcal{P}}\right)$ , where  $\mathcal{A}$  is given by Algorithm A.1 and  $\mathcal{A}' = \mathcal{A}\cup\{j\}$  for any  $j \in \mathcal{P} \setminus \mathcal{A}$ .

Theorem 3 shows that for general  $T(\mathcal{M})$ , adding another component to  $\mathcal{A}$  cannot improve the solution to (7).

*Remark.* Theorem 3 gives a weaker result for general  $T(\cdot)$  compared to the global optimality result in Theorem 1 for linear  $T(\cdot)$ , because when  $\mathcal{A}$  and  $\mathcal{A}'$  differ by more than one element, it is non-straightforward to express the change in cost for general  $T(\cdot)$ . Furthermore, there may exist multiple local optimal solutions for general  $T(\cdot)$ .

# B. Proofs

1) Proof of Theorems 1 and 3: Recall that  $\Gamma(\mathcal{M}) := \frac{\Delta(\mathcal{M})}{T(\mathcal{M})}$ , where  $\Delta(\mathcal{M})$  and  $T(\mathcal{M})$  are both monotone and positive functions, i.e., for any  $\mathcal{M} \subseteq \mathcal{M}'$ , we have  $0 \leq \Delta(\mathcal{M}) \leq \Delta(\mathcal{M}')$  and  $0 \leq T(\mathcal{M}) \leq T(\mathcal{M}')$ .

**Lemma 1.** For any  $\mathcal{M}$  and  $\mathcal{M}'$ , let  $\delta_{\Delta}(\mathcal{M}, \mathcal{M}') := \Delta(\mathcal{M}') - \Delta(\mathcal{M})$  and  $\delta_T(\mathcal{M}, \mathcal{M}') := T(\mathcal{M}') - T(\mathcal{M})$ . We have  $\Gamma(\mathcal{M}') \leq \Gamma(\mathcal{M})$  if and only if  $\delta_{\Delta}(\mathcal{M}, \mathcal{M}') \leq \Gamma(\mathcal{M}) \cdot \delta_T(\mathcal{M}, \mathcal{M}')$ .

Proof.

$$\begin{split} \Gamma(\mathcal{M}') &:= \frac{\Delta(\mathcal{M}')}{T(\mathcal{M}')} \leq \Gamma(\mathcal{M}) \\ \iff \Delta(\mathcal{M}') \leq \Gamma(\mathcal{M}) \cdot T(\mathcal{M}') \\ \iff \Delta(\mathcal{M}) + \delta_{\Delta}(\mathcal{M}, \mathcal{M}') \leq \Gamma(\mathcal{M}) \cdot T(\mathcal{M}) + \Gamma(\mathcal{M}) \cdot \delta_{T}(\mathcal{M}, \mathcal{M}') \\ \iff \Delta(\mathcal{M}) + \delta_{\Delta}(\mathcal{M}, \mathcal{M}') \leq \Delta(\mathcal{M}) + \Gamma(\mathcal{M}) \cdot \delta_{T}(\mathcal{M}, \mathcal{M}') \\ \iff \delta_{\Delta}(\mathcal{M}, \mathcal{M}') \leq \Gamma(\mathcal{M}) \cdot \delta_{T}(\mathcal{M}, \mathcal{M}') . \end{split}$$
 (by definition of  $\delta_{\Delta}(\cdot, \cdot)$  and  $\delta_{T}(\cdot, \cdot)$ )  
$$\iff \delta_{\Delta}(\mathcal{M}, \mathcal{M}') \leq \Gamma(\mathcal{M}) \cdot \delta_{T}(\mathcal{M}, \mathcal{M}') . \end{split}$$

We are now ready to prove Theorems 1 and 3.

Proof of Theorem 1. By definition, we have  $\Delta(\mathcal{M}) := \sum_{j \in \mathcal{M}} g_j^2$  and  $T(\mathcal{M}) := c + \sum_{j \in \mathcal{M}} t_j$  for any  $\mathcal{M}$ . In the following, we let  $\mathcal{M} := \mathcal{A} \cup \overline{\mathcal{P}}$  and  $\mathcal{M}' := \mathcal{A}' \cup \overline{\mathcal{P}}$ . We have

 $\delta_T(.$ 

$$\delta_{\Delta}(\mathcal{M}, \mathcal{M}') = \Delta(\mathcal{M}') - \Delta(\mathcal{M})$$
  
=  $\sum_{j \in \mathcal{M}' \setminus \mathcal{M}} g_j^2 - \sum_{j \in \mathcal{M} \setminus \mathcal{M}'} g_j^2$  (B.1)

$$\mathcal{M}, \mathcal{M}') = T(\mathcal{M}') - T(\mathcal{M})$$
  
=  $\sum_{j \in \mathcal{M}' \setminus \mathcal{M}} t_j - \sum_{j \in \mathcal{M} \setminus \mathcal{M}'} t_j$ . (B.2)

For  $\mathcal{A}$  obtained from Algorithm 2, we can easily see that  $\frac{g_j^2}{t_j} < \Gamma(\mathcal{M})$  for any  $j \in \mathcal{M}' \setminus \mathcal{M}$  and  $\frac{g_{j'}}{t_{j'}} \ge \Gamma(\mathcal{M})$  for any  $j' \in \mathcal{M} \setminus \mathcal{M}'$ . Hence,

$$\sum_{j \in \mathcal{M}' \setminus \mathcal{M}} g_j^2 < \Gamma(\mathcal{M}) \cdot \sum_{j \in \mathcal{M}' \setminus \mathcal{M}} t_j$$
(B.3)

$$\sum_{\in \mathcal{M} \setminus \mathcal{M}'} g_j^2 \ge \Gamma(\mathcal{M}) \cdot \sum_{j \in \mathcal{M} \setminus \mathcal{M}'} t_j .$$
(B.4)

Combining with (B.1) and (B.2), we have

$$\delta_{\Delta}(\mathcal{M}, \mathcal{M}') \leq \Gamma(\mathcal{M}) \cdot \delta_T(\mathcal{M}, \mathcal{M}')$$
.

Then, the result follows from Lemma 1.

Proof of Theorem 3. Let  $\mathcal{M} := \mathcal{A} \cup \overline{\mathcal{P}}$  and  $\mathcal{M}' := \mathcal{A}' \cup \overline{\mathcal{P}}$  in this proof. As  $\mathcal{A}' := \mathcal{A} \cup \{j\}$  for some  $j \notin \mathcal{A}$  by definition in this theorem, we note that  $\delta_{\Delta}(\mathcal{M}, \mathcal{M}') = g_j^2$  and  $\delta_T(\mathcal{M}, \mathcal{M}') = t_j(\mathcal{M})$ .

For  $\mathcal{A}$  obtained from Algorithm A.1, it is easy to see that  $\frac{g_j^2}{t_j(\mathcal{M})} < \Gamma(\mathcal{M})$  for  $j \notin \mathcal{A}$ . Hence,  $\delta_{\Delta}(\mathcal{M}, \mathcal{M}') < \Gamma(\mathcal{M}) \cdot \delta_T(\mathcal{M}, \mathcal{M}')$  and the result follows from Lemma 1.

2) Proof of Theorem 2: The analysis of this section is an extension of Theorem 1 in [24] (proof given in Section A.1). Note that Assumption 1(a)-(e) still hold if we apply the same mask to all gradients or function values in the LHS, since applying masks is equivalent to replacing a subset of the entries in the parameter with zeros. For the same reason, Assumption 1(f) and 1(g) also hold if the gradients are masked. For convenience, we define three expedient notations in addition to the notations in Table I for pruned values:  $\mathbf{g}'_n(\mathbf{w}) := \mathbf{g}_n(\mathbf{w}) \odot \mathbf{m}(k), \nabla' F_n(\mathbf{w}) := \nabla F_n(\mathbf{w}) \odot \mathbf{m}(k), and \nabla' F(\mathbf{w}) := \nabla F(\mathbf{w}) \odot \mathbf{m}(k).$ 

We first present a special form of Jensen's inequality, which has the original form below:

j

$$\phi\left(\frac{\sum_{i} a_{i} x_{i}}{\sum_{i} a_{i}}\right) \leq \frac{\sum_{i} a_{i} \phi(x_{i})}{\sum_{i} a_{i}}$$

where  $\phi(\cdot)$  is a convex function,  $a_i$ 's are positive weights.

**Lemma 2.** Assume  $\phi(\cdot)$  is a convex function,  $p_n$ 's are positive weights that sum to 1, we have

$$\left\|\sum_{n} p_n \mathbf{x}_n\right\|^2 \le \sum_{n} p_n \|\mathbf{x}_n\|^2$$

Proof.

$$\left\|\sum_{n} p_{n} \mathbf{x}_{n}\right\|^{2} = \left\|\begin{bmatrix}\sum_{n} p_{n} x_{n,1}\\ \sum_{n} p_{n} x_{n,2}\\ \vdots \end{bmatrix}\right\|^{2} = \sum_{i} \left(\sum_{n} p_{n} x_{n,i}\right)^{2} \le \sum_{i} \sum_{n} p_{n} x_{n,i}^{2} = \sum_{n} p_{n} \sum_{i} x_{n,i}^{2} = \sum_{n} p_{n} \|\mathbf{x}_{n}\|^{2},$$

where  $x_{n,i}$  denotes the *i*-th component of  $\mathbf{x}_n$ .

The local updating rule is

$$\mathbf{w}'_n(k+1) = \mathbf{w}_n(k) \odot \mathbf{m}(k) - \eta \mathbf{g}_n \big( \mathbf{w}_n(k) \odot \mathbf{m}(k) \big) \odot \mathbf{m}(k) = \mathbf{w}'_n(k) - \eta \mathbf{g}'_n(\mathbf{w}'_n(k)) ,$$

Consequently, the updating rule for the averaged parameters is

$$\mathbf{w}(k+1) = \sum_{n=1}^{N} p_n \mathbf{w}'_n(k+1) = \mathbf{w}'(k) - \eta \sum_{n=1}^{N} p_n \mathbf{g}'_n(\mathbf{w}'_n(k))$$

Note that  $\mathbf{w}(k)$ ,  $\mathbf{w}'(k)$  are observable only in iterations where the clients send their local parameters to the server for aggregation (see Algorithm 1), and  $\mathbf{g}'(\mathbf{w}'_n(k))$  is dependent on  $\mathbf{m}(k)$ . We know that

$$\begin{split} & \mathbb{E}\left[F(\mathbf{w}(k+1))|\{\mathbf{w}_{n}(k)\},\mathbf{m}(k)\right] \\ &= \mathbb{E}\left[F\left(\mathbf{w}'(k) - \eta \sum_{n=1}^{N} p_{n}\mathbf{g}'_{n}(\mathbf{w}'_{n}(k))\right)\Big|\{\mathbf{w}_{n}(k)\},\mathbf{m}(k)\right] \\ &\stackrel{(a)}{\leq} F(\mathbf{w}'(k)) - \eta \mathbb{E}\left[\left\langle \nabla F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_{n}\mathbf{g}'_{n}(\mathbf{w}'_{n}(k))\right\rangle\Big|\{\mathbf{w}_{n}(k)\},\mathbf{m}(k)\right] + \frac{\eta^{2}\beta}{2} \mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n}\mathbf{g}'_{n}(\mathbf{w}'_{n}(k))\right\|^{2}\Big|\{\mathbf{w}_{n}(k)\},\mathbf{m}(k)\right] \\ &\stackrel{(b)}{\leq} F(\mathbf{w}(k)) + L \|\mathbf{w}(k) - \mathbf{w}'(k)\| - \eta \mathbb{E}\left[\left\langle \nabla F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_{n}\mathbf{g}'_{n}(\mathbf{w}'_{n}(k))\right\rangle\Big|\{\mathbf{w}_{n}(k)\}\right\rangle\Big|\{\mathbf{w}_{n}(k)\},\mathbf{m}(k)\right] \end{split}$$

$$+ \frac{\eta^{2}\beta}{2} \mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \mathbf{g}_{n}'(\mathbf{w}_{n}'(k))\right\|^{2} \left|\{\mathbf{w}_{n}(k)\}, \mathbf{m}(k)\right]\right]$$
$$= F(\mathbf{w}(k)) + L \left\|\mathbf{w}(k) - \mathbf{w}'(k)\right\| - \eta \mathbb{E}\left[\left\langle\nabla' F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_{n} \mathbf{g}_{n}'(\mathbf{w}_{n}'(k))\right\rangle \left|\{\mathbf{w}_{n}(k)\}, \mathbf{m}(k)\right]\right]$$
$$+ \frac{\eta^{2}\beta}{2} \mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \mathbf{g}_{n}'(\mathbf{w}_{n}'(k))\right\|^{2} \left|\{\mathbf{w}_{n}(k)\}, \mathbf{m}(k)\right]\right]$$
(B.5)

where (a) is due to Assumption 1(a) ( $\beta$ -smoothness), (b) is due to Assumption 1(b) (L-Lipschitzness), and (B.5) is because

$$\begin{split} \left\langle \nabla F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_n \mathbf{g}'_n(\mathbf{w}'_n(k)) \right\rangle &= \left\langle \nabla F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_n \mathbf{g}_n(\mathbf{w}'_n(k)) \odot \mathbf{m}(k) \right\rangle \\ &= \left\langle \nabla F(\mathbf{w}'(k)) \odot \mathbf{m}(k), \sum_{n=1}^{N} p_n \mathbf{g}_n(\mathbf{w}'_n(k)) \odot \mathbf{m}(k) \right\rangle \\ &= \left\langle \nabla' F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_n \mathbf{g}'_n(\mathbf{w}'_n(k)) \right\rangle \,. \end{split}$$

The third term in (B.5) can be rewritten as

$$-\eta \mathbb{E}\left[\left\langle \nabla' F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_n \mathbf{g}'_n(\mathbf{w}'_n(k)) \right\rangle \Big| \{\mathbf{w}_n(k)\}, \mathbf{m}(k) \right]$$

$$\stackrel{(a)}{=} -\eta \left\langle \nabla' F(\mathbf{w}'(k)), \sum_{n=1}^{N} p_n \nabla' F_n(\mathbf{w}'_n(k)) \right\rangle$$

$$= \frac{\eta}{2} \left( \left\| \nabla' F(\mathbf{w}'(k)) - \sum_{n=1}^{N} p_n \nabla' F_n(\mathbf{w}'_n(k)) \right\|^2 - \left\| \nabla' F(\mathbf{w}'(k)) \right\|^2 - \left\| \sum_{n=1}^{N} p_n \nabla' F_n(\mathbf{w}'_n(k)) \right\|^2 \right), \quad (B.6)$$

where (a) uses Assumption 1(c) (unbiasedness). Our fourth term in (B.5) is bounded by

$$\mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \mathbf{g}_{n}'(\mathbf{w}_{n}'(k))\right\|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right]\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right] + \left(\mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \nabla' F_{n}(\mathbf{w}_{n}'(k))\right\|\right| \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right]\right)^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right] + \left(\mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \nabla' F_{n}(\mathbf{w}_{n}'(k))\right\|\right| \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right]\right)^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right] + \left(\mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \nabla' F_{n}(\mathbf{w}_{n}'(k))\right\|\right| \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right]\right)^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right] + \left(\mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \nabla' F_{n}(\mathbf{w}_{n}'(k))\right\|\right| \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right]\right)^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right] + \left(\mathbb{E}\left[\left\|\sum_{n=1}^{N} p_{n} \nabla' F_{n}(\mathbf{w}_{n}'(k))\right\|\right| \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right]\right)^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2}\right]^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2}\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right\}\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left\{\mathbf{w}_{n}(k)\right\}, \mathbf{m}(k)\right|^{2} \left|\left$$

where (a) uses the definition of variance, i.e.,  $\mathbb{E}[\|\mathbf{x}\|^2] = \mathbb{E}[\|\mathbf{x} - \mathbb{E}[\mathbf{x}]\|^2] + [\mathbb{E}\|\mathbf{x}\|]^2$ ; (b) uses Assumption 1(g) (client independence), expands the first term and removes the zero-valued cross-product terms; (c) uses Assumption 1(d) (bounded variance). Substituting (B.6) and (B.7) into (B.5), (B.5) becomes

$$\mathbb{E}\left[F(\mathbf{w}(k+1))|\{\mathbf{w}_{n}(k)\},\mathbf{m}(k)\right] \\
\leq F(\mathbf{w}(k)) + L\|\mathbf{w}(k) - \mathbf{w}'(k)\| + \frac{\eta^{2}\beta\sum_{n=1}^{N}p_{n}^{2}}{2}\sigma^{2} - \frac{\eta}{2}\|\nabla'F(\mathbf{w}'(k))\|^{2} + \frac{\eta}{2}\|\nabla'F(\mathbf{w}'(k)) - \sum_{n=1}^{N}p_{n}\nabla'F_{n}(\mathbf{w}_{n}'(k))\|^{2} \\
- \left(\frac{\eta}{2} - \frac{\eta^{2}\beta}{2}\right)\|\sum_{n=1}^{N}p_{n}\nabla'F_{n}(\mathbf{w}_{n}'(k))\|^{2}.$$
(B.8)

Assuming  $\eta \leq \frac{1}{\beta}$ , we have  $\frac{\eta}{2} - \frac{\eta^2 \beta}{2} \geq 0$ , and the last term in (B.8) can be removed. Then we have the following:

 $\mathbb{E}\left[F(\mathbf{w}(k+1))|\{\mathbf{w}_n(k)\},\mathbf{m}(k)\right]$ 

$$\leq F(\mathbf{w}(k)) + L \|\mathbf{w}(k) - \mathbf{w}'(k)\| + \frac{\eta^2 \beta \sum_{n=1}^{N} p_n^2}{2} \sigma^2 - \frac{\eta}{2} \|\nabla' F(\mathbf{w}'(k))\|^2 + \frac{\eta}{2} \|\nabla' F(\mathbf{w}'(k)) - \sum_{n=1}^{N} p_n \nabla' F_n(\mathbf{w}'_n(k))\|^2 \\ \leq F(\mathbf{w}(k)) + L \|\mathbf{w}(k) - \mathbf{w}'(k)\| + \frac{\eta^2 \beta \sum_{n=1}^{N} p_n^2}{2} \sigma^2 - \frac{\eta}{2} \|\nabla' F(\mathbf{w}'(k))\|^2 + \frac{\eta}{2} \sum_{n=1}^{N} p_n \|\nabla' F_n(\mathbf{w}'(k)) - \nabla' F_n(\mathbf{w}'_n(k))\|^2 \\ \leq F(\mathbf{w}(k)) + L \|\mathbf{w}(k) - \mathbf{w}'(k)\| + \frac{\eta^2 \beta \sum_{n=1}^{N} p_n^2}{2} \sigma^2 - \frac{\eta}{2} \|\nabla' F(\mathbf{w}'(k))\|^2 + \frac{\eta\beta^2}{2} \sum_{n=1}^{N} p_n \|\mathbf{w}'(k) - \mathbf{w}'_n(k)\|^2 ,$$
(B.9)

where (a) uses Jensen's inequality, and (b) uses Assumption 1(a) (smoothness). Taking expectation on both sides of (B.9), we get

$$\mathbb{E}\left[F(\mathbf{w}(k+1))\right] \le \mathbb{E}\left[F(\mathbf{w}(k))\right] + L\mathbb{E}\left\|\mathbf{w}(k) - \mathbf{w}'(k)\right\| + \frac{\eta^2 \beta \sum_{n=1}^N p_n^2}{2} \sigma^2 - \frac{\eta}{2} \mathbb{E}\left\|\nabla' F(\mathbf{w}'(k))\right\|^2 + \frac{\eta \beta^2}{2} \sum_{n=1}^N p_n \mathbb{E}\left\|\mathbf{w}'(k) - \mathbf{w}'_n(k)\right\|^2.$$
(B.10)

Taking average over time on (B.10) and rearranging, we get

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla' F(\mathbf{w}'(k)) \right\|^{2} \\
\leq \frac{2}{\eta K} [F(\mathbf{w}(0)) - F^{*}] + \frac{2L}{\eta K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \mathbf{w}(k) - \mathbf{w}'(k) \right\| + \eta \beta \Big( \sum_{n=1}^{N} p_{n}^{2} \Big) \sigma^{2} + \frac{\beta^{2}}{K} \sum_{k=0}^{K-1} \sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \mathbf{w}'(k) - \mathbf{w}'_{n}(k) \right\|^{2}. \quad (B.11)$$

Now we bound the last term of (B.11).

$$\begin{split} &\sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \mathbf{w}'(k) - \mathbf{w}'_{n}(k) \right\|^{2} \\ &= \sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \left( \mathbf{w}'(k-1) - \eta \sum_{i=1}^{N} p_{i} \mathbf{g}'_{i}(\mathbf{w}'_{i}(k-1)) \right) - \left( \mathbf{w}'_{n}(k-1) - \eta \mathbf{g}'_{n}(\mathbf{w}'_{n}(k-1)) \right) \right\|^{2} \\ &\stackrel{(a)}{=} \eta^{2} \sum_{n=1}^{N} p_{n} \mathbb{E} \right\| \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \left( \mathbf{g}'_{n}(\mathbf{w}'_{n}(\tau)) - \sum_{i=1}^{N} p_{i} \mathbf{g}'_{i}(\mathbf{w}'_{i}(\tau)) \right) \right\|^{2} \\ &= \eta^{2} \sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \left( \left( \mathbf{g}'_{n}(\mathbf{w}'_{n}(\tau)) - \nabla' F_{n}(\mathbf{w}'_{n}(\tau)) + \sum_{i=1}^{N} p_{i} \nabla' F_{i}(\mathbf{w}'_{i}(\tau)) - \sum_{i=1}^{N} p_{i} \mathbf{g}'_{i}(\mathbf{w}'_{i}(\tau)) \right) \right\|^{2} \\ &+ \left( \nabla' F_{n}(\mathbf{w}'_{n}(\tau)) - \sum_{i=1}^{N} p_{i} \nabla' F_{i}(\mathbf{w}'_{i}(\tau)) \right) \right\|^{2} \\ &\leq 2\eta^{2} \sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \left( \mathbf{g}'_{n}(\mathbf{w}'_{n}(\tau)) - \nabla' F_{n}(\mathbf{w}'_{n}(\tau)) + \sum_{i=1}^{N} p_{i} \nabla' F_{i}(\mathbf{w}'_{i}(\tau)) - \sum_{i=1}^{N} p_{i} \mathbf{g}'_{i}(\mathbf{w}'_{i}(\tau)) \right) \right\|^{2} \end{aligned} \tag{B.12} \\ &+ 2\eta^{2} \sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \left( \nabla' F_{n}(\mathbf{w}'_{n}(\tau)) - \sum_{i=1}^{N} p_{i} \nabla' F_{i}(\mathbf{w}'_{i}(\tau)) \right) \right\|^{2} \end{aligned}$$

$$-2\eta^{2}\sum_{n=1}^{N}p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{n-1}\left(\nabla'F_{n}(\mathbf{w}_{n}'(\tau))-\sum_{i=1}^{N}p_{i}\nabla'F_{n}(\mathbf{w}_{i}'(\tau))\right)\right\|,\qquad(B.13)$$

where in (a), we trace back to the nearest iteration where all local parameters are synchronized. For (B.12),

$$2\eta^{2}\sum_{n=1}^{N}p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\left(\mathbf{g}_{n}'(\mathbf{w}_{n}'(\tau))-\nabla'F_{n}(\mathbf{w}_{n}'(\tau))+\sum_{i=1}^{N}p_{i}\nabla'F_{i}(\mathbf{w}_{i}'(\tau))-\sum_{i=1}^{N}p_{i}\mathbf{g}_{i}'(\mathbf{w}_{i}'(\tau))\right)\right\|^{2}$$

$$\stackrel{(a)}{=}2\eta^{2}\sum_{n=1}^{N}p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\left(\mathbf{g}_{n}'(\mathbf{w}_{n}'(\tau))-\nabla'F_{n}(\mathbf{w}_{n}'(\tau))\right)\right\|^{2}-2\eta^{2}\sum_{n=1}^{N}p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\sum_{i=1}^{N}p_{i}\left(\mathbf{g}_{i}'(\mathbf{w}_{i}'(\tau))-\nabla'F_{i}(\mathbf{w}_{i}'(\tau))\right)\right\|^{2}$$

$$\stackrel{(b)}{=}2\eta^{2}\sum_{n=1}^{N}p_{n}\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\mathbb{E}\left\|\mathbf{g}_{n}'(\mathbf{w}_{n}'(\tau))-\nabla'F_{n}(\mathbf{w}_{n}'(\tau))\right\|^{2}-2\eta^{2}\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\mathbb{E}\left\|\sum_{i=1}^{N}p_{i}\left(\mathbf{g}_{i}'(\mathbf{w}_{i}'(\tau))-\nabla'F_{i}(\mathbf{w}_{i}'(\tau))\right)\right\|^{2}$$

$$\stackrel{(c)}{=} 2\eta^{2} \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \sum_{n=1}^{N} p_{n} \mathbb{E} \left\| \mathbf{g}_{n}'(\mathbf{w}_{n}'(\tau)) - \nabla' F_{n}(\mathbf{w}_{n}'(\tau)) \right\|^{2} - 2\eta^{2} \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \sum_{n=1}^{N} \mathbb{E} \left\| p_{n} \left( \mathbf{g}_{n}'(\mathbf{w}_{n}'(\tau)) - \nabla' F_{n}(\mathbf{w}_{n}'(\tau)) \right) \right\|^{2}$$

$$\leq 2\eta^{2} \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{k-1} \sum_{n=1}^{N} (p_{n} - p_{n}^{2}) \mathbb{E} \left\| \mathbf{g}_{n}'(\mathbf{w}_{n}'(\tau)) - \nabla' F_{n}(\mathbf{w}_{n}'(\tau)) \right\|^{2}$$

$$\leq 2\left(1 - \sum_{n=1}^{N} p_{n}^{2}\right) I \eta^{2} \sigma^{2} ,$$

$$(B.14)$$

where (a) is due to the definition of variance; (b) is due to Assumption 1(f) (time independence) and  $\sum_n p_n = 1$ ; and (c) is due to Assumption 1(g) (client independence). For (B.13),

$$\begin{split} &2\eta^{2}\sum_{n=1}^{N}p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot[k/I]}^{k-1}\left(\nabla'F_{n}(\mathbf{w}_{n}'(\tau))-\sum_{i=1}^{N}p_{i}\nabla'F_{i}(\mathbf{w}_{i}'(\tau))\right)\right\|^{2} \\ &=2\eta^{2}\sum_{n=1}^{N}p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot[k/I]}^{k-1}\left(\left(\nabla'F_{n}(\mathbf{w}_{n}'(\tau))-\nabla'F_{n}(\mathbf{w}'(\tau))\right)+\left(\nabla'F_{n}(\mathbf{w}'(\tau))-\sum_{i=1}^{N}p_{i}\nabla'F_{i}(\mathbf{w}'(\tau))\right)\right) \\ &+\left(\sum_{i=1}^{N}p_{i}\nabla'F_{i}(\mathbf{w}'(\tau))-\sum_{i=1}^{N}p_{i}\nabla'F_{i}(\mathbf{w}_{i}(\tau))\right)\right)\right\|^{2} \\ &\leq 6\eta^{2}\sum_{n=1}^{N}\left(p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot[k/I]}^{k-1}\left(\nabla'F_{n}(\mathbf{w}_{n}'(\tau))-\nabla'F_{n}(\mathbf{w}'(\tau))\right)\right\|^{2} +p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot[k/I]}^{k-1}\sum_{i=1}^{N}p_{i}\left(\nabla'F_{i}(\mathbf{w}'(\tau))-\nabla'F_{i}(\mathbf{w}'(\tau))\right)\right\|^{2} \\ &+p_{n}\mathbb{E}\left\|\sum_{\tau=I\cdot[k/I]}^{k-1}\left(\nabla'F_{n}(\mathbf{w}'(\tau))-\nabla'F_{n}(\mathbf{w}'(\tau))\right)\right\|^{2} +\mathbb{E}\left\|\sum_{i=1}^{N}p_{i}\left(\nabla'F_{i}(\mathbf{w}'(\tau))-\nabla'F_{i}(\mathbf{w}_{i}(\tau))\right)\right\|^{2} \right) \\ &+6\eta^{2}I\mathbb{E}\left[\sum_{n=1}^{k}p_{n}\left(\mathbb{E}\left\|\nabla'F_{n}(\mathbf{w}_{n}(\tau))-\nabla'F_{n}(\mathbf{w}'(\tau))\right\right\|^{2} +\mathbb{E}\left\|\sum_{i=1}^{N}p_{i}\left(\nabla'F_{i}(\mathbf{w}'(\tau))-\nabla'F_{i}(\mathbf{w}_{i}(\tau))\right)\right\|^{2} \right) +6I^{2}\eta^{2}\epsilon^{2} \\ &\leq (12\eta^{2}I\sum_{\tau=I\cdot[k/I]}^{k-1}\sum_{n=1}^{k-1}\mathbb{E}\left\|\nabla'F_{n}(\mathbf{w}_{n}(\tau))-\nabla'F_{n}(\mathbf{w}'(\tau))\right\|^{2} +6I^{2}\eta^{2}\epsilon^{2} \\ &\leq (12\eta^{2}\beta^{2}\sum_{n=1}^{k-1}\sum_{\tau=I\cdot[k/I]}^{k-1}\mathbb{E}\left\|\mathbf{w}_{n}'(\tau)-\mathbf{w}'(\tau)\right\|^{2} +6I^{2}\eta^{2}\epsilon^{2} , \end{aligned}$$
(B.15)

where in (a), we use the fact that,  $\forall \mathbf{x}_{\tau}$ ,

$$\left\|\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\mathbf{x}_{\tau}\right\|^{2} \leq \left((k-1)-I\cdot\lfloor k/I\rfloor\right)\cdot\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\|\mathbf{x}_{\tau}\|^{2} \leq I\cdot\sum_{\tau=I\cdot\lfloor k/I\rfloor}^{k-1}\|\mathbf{x}_{\tau}\|^{2}.$$

In (b), we use Assumption 1(e) (bounded divergence), and in (c), we use Jensen's inequality. Substituting (B.14) and (B.15) into (B.12) and (B.13), respectively, we get

$$\frac{1}{K} \sum_{k=0}^{K-1} \sum_{n=1}^{N} p_n \mathbb{E} \| \mathbf{w}'(k) - \mathbf{w}'_n(k) \|^2 
\leq 2 \Big( 1 - \sum_{n=1}^{N} p_n^2 \Big) I \eta^2 \sigma^2 + 6 \eta^2 I^2 \epsilon^2 + \frac{12 I \eta^2 \beta^2}{K} \sum_{n=1}^{N} \sum_{k=0}^{K-1} \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{K-1} \mathbb{E} \| \mathbf{w}'_n(\tau) - \mathbf{w}'(\tau) \|^2 
\stackrel{(a)}{\leq} 2 \Big( 1 - \sum_{n=1}^{N} p_n^2 \Big) \eta^2 I \sigma^2 + 6 \eta^2 I^2 \epsilon^2 + \frac{12 I \eta^2 \beta^2}{K} \sum_{n=1}^{N} \sum_{k=0}^{K-1} \sum_{\tau=I \cdot \lfloor k/I \rfloor}^{K-1} \mathbb{E} \| \mathbf{w}'_n(\tau) - \mathbf{w}'(\tau) \|^2$$

$$\overset{(b)}{\leq} 2\Big(1 - \sum_{n=1}^{N} p_n^2\Big)\eta^2 I\sigma^2 + 6\eta^2 I^2 \epsilon^2 + \frac{12I^2 \eta^2 \beta^2}{K} \sum_{k=0}^{K-1} \sum_{n=1}^{N} \mathbb{E}\Big\|\mathbf{w}_n'(\tau) - \mathbf{w}'(\tau)\Big\|^2 ,$$
(B.16)

where (a) holds because  $k - 1 \le \min\{I \cdot \lceil k/I \rceil, K - 1\}$  is always true, and (b) holds because in the innermost summation in its last term, we always have

$$\min\{I \cdot \lceil k/I \rceil, K-1\} - I \cdot \lfloor k/I \rfloor \le I .$$

Rearranging yields

$$\frac{1}{K} \sum_{k=0}^{K-1} \sum_{n=1}^{N} p_n \mathbb{E} \left\| \mathbf{w}'(k) - \mathbf{w}'_n(k) \right\|^2 \le \frac{2\left(1 - \sum_{n=1}^{N} p_n^2\right) \eta^2 I \sigma^2 + 6\eta^2 I^2 \epsilon^2}{1 - 12I^2 \eta^2 \beta^2} . \tag{B.17}$$

Applying (B.17) into (B.11)'s last term, (B.11) becomes

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla' F(\mathbf{w}'(k)) \right\|^2 \le \frac{2}{\eta K} [F(\mathbf{w}(0)) - F^*] + \eta \beta \Big( \sum_{n=1}^N p_n^2 \Big) \sigma^2 + \frac{2 \Big( 1 - \sum_{n=1}^N p_n^2 \Big) I \sigma^2 + 6I^2 \epsilon^2}{1 - 12I^2 \eta^2 \beta^2} \eta^2 \beta^2 \\ + \frac{2L}{\eta K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \mathbf{w}(k) - \mathbf{w}'(k) \right\| .$$
(B.18)

Assume  $\eta \leq \frac{1}{2\sqrt{6}I\beta}$ , we have  $1 - 12\eta^2 I^2 \eta^2 \beta^2 \geq \frac{1}{2}$ , and

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla' F(\mathbf{w}'(k)) \right\|^2 \le \frac{2(F_0 - F^*)}{\eta K} + \alpha \eta \beta \sigma^2 + 4\beta^2 \left( (1 - \alpha) I \sigma^2 + 3I^2 \epsilon^2 \right) \eta^2 + \frac{2L}{\eta K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \mathbf{w}(k) - \mathbf{w}'(k) \right\|, \quad (B.19)$$

where  $\alpha := \sum_{n=1}^{N} p_n^2$ ,  $F_0 := F(\mathbf{w}(0))$ ,  $F^* := \min_{\mathbf{w}} F(\mathbf{w})$ . Note that under the assumption that  $\eta \le \frac{1}{2\sqrt{6}I\beta}$ , the previous assumption that  $\eta \le \frac{1}{\beta}$  used for (B.8) is automatically satisfied.

**Discussion.** Consider the situation where all clients have equal weight, i.e.  $p_n = \frac{1}{N}$ ,  $\forall n$ , we have  $\alpha = \frac{1}{N}$ . Letting  $\eta = \frac{1}{\sqrt{\alpha K}} = \sqrt{\frac{N}{K}}$ , (B.19) becomes

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla' F(\mathbf{w}'(k)) \right\|^2 \le \frac{2(F_0 - F^*) + \beta \sigma^2}{\sqrt{NK}} + \frac{4\beta^2 \left( (N-1)I\sigma^2 + 3NI^2\epsilon^2 \right)}{K} + \frac{2L}{\sqrt{NK}} \sum_{k=0}^{K-1} \mathbb{E} \left\| \mathbf{w}(k) - \mathbf{w}'(k) \right\| .$$
(B.20)

For the last term in (B.20), when k is not a reconfiguration iteration,  $\mathbf{w}(k) = \mathbf{w}'(k)$ , and when k is a reconfiguration iteration, the difference between  $\mathbf{w}(k)$  and  $\mathbf{w}'(k)$  is the subset of parameters that get pruned from  $\mathbf{w}(k)$ . Assuming the norm of  $\mathbf{w}(k)$  is bounded by B, i.e.,  $\|\mathbf{w}(k)\| \le B$ ,  $\forall k$ , the initial fraction of non-zero prunable parameters is  $r_0$  and this fraction halves every h iterations (as we do in our experiments, see Table II), then the last term in (B.20) is bounded by

$$\frac{2L}{\sqrt{NK}} \sum_{k=0}^{K-1} \mathbb{E} \left\| \mathbf{w}(k) - \mathbf{w}'(k) \right\| \stackrel{(a)}{\leq} \frac{2L}{\sqrt{NK}} \sum_{k=0}^{K-1} B \cdot r_0 \cdot 2^{-k/h} \leq \frac{2LBr_0}{\sqrt{NK}} \sum_{k=0}^{\infty} 2^{-k/h} \leq \frac{2^{\frac{1}{h}+1}}{2^{\frac{1}{h}}-1} \cdot \frac{LBr_0}{\sqrt{NK}} . \tag{B.21}$$

In (B.21), (a) holds for the following reason: the reconfiguration on the parameter vector  $\mathbf{w}(k)$  includes both adding back parameters and removing parameters, resulting in a new parameter vector  $\mathbf{w}'(k)$ . However, parameters that are added back in  $\mathbf{w}'(k)$  in reconfigurations are assigned zero values<sup>2</sup>, which are equal to their corresponding values in  $\mathbf{w}(k)$ . Thus, the difference between  $\mathbf{w}(k)$  and  $\mathbf{w}'(k)$  is the part that is *removed* from  $\mathbf{w}(k)$ , whose maximum fraction is bounded by  $r_0 \cdot 2^{-k/h}$ . In consequence,  $\|\mathbf{w}(k) - \mathbf{w}'(k)\|$  is bounded by  $B \cdot r_0 \cdot 2^{-k/h}$ . Plugging (B.21) into (B.20), we get

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla' F(\mathbf{w}'(k)) \right\|^2 \le \left( 2(F_0 - F^*) + \beta \sigma^2 + \frac{2^{\frac{1}{h} + 1}}{2^{\frac{1}{h}} - 1} LBr_0 \right) \frac{1}{\sqrt{NK}} + \frac{4\beta^2 \left( (N-1)I\sigma^2 + 3NI^2\epsilon^2 \right)}{K} = \mathcal{O} \left( \frac{1}{\sqrt{NK}} \right) + \mathcal{O} \left( \frac{1}{K} \right) \ . \tag{B.22}$$

Thus, when the additional assumptions (i)  $K \ge 24NI^2\beta^2$ , (ii)  $\eta = \sqrt{\frac{N}{K}}$ , (iii)  $p_n = \frac{1}{N}$ ,  $\forall n$ , (iv)  $\|\mathbf{w}(k)\| \le B$ ,  $\forall k$ , and (v) the fraction of non-zero prunable parameters decreases exponentially all hold, we have a convergence bound provided in (B.22) that is dominated by  $\mathcal{O}\left(\frac{1}{\sqrt{NK}}\right)$ . This means using more clients can accelerate the convergence (by a factor of  $\frac{1}{\sqrt{N}}$ ).

Architecture	Conv-2	VGG-11	ResNet-18	MobileNetV3-Small
Convolutional Fully-connected	22 <b>p</b> ool	64, pool, 128, pool,	64, pool, $2 \times [64, 64],$	16, 16, 8, 16, 16, 72, 72, 24, 88, 88, 24, 96, 96, 24, 96, 40, 240, 240, 64, 240, 40, 240, 240, 64, 240, 40,
	64, pool	$2 \times 256$ , pool, $2 \times 512$ , pool, $2 \times 512$ pool	$2 \times [128, 128],$ $2 \times [256, 256],$ $2 \times [512, 512]$	120, 120, 32, 120, 48, 144, 144, 40, 144, 48, 288, 288, 72, 288, 96, 576, 576, 144, 576, 96, 576, 576, 144, 576, 96, 576, 144, 144, 144, 144, 144, 144, 144, 14
Fully-connected	2048, 62	512, 512, 10	avgpool, 100	avgpool, 1280, dropout (0.2), 2
Conv/FC/all params	52.1K/6.6M/6.6M	9.2M/530.4K/9.8M	11.2M/102.6K/11.3M	927.0K/592.9K/1.5M

TABLE C.1 MODEL ARCHITECTURES.

## C. PruneFL details

1) Model architecture details: The details of the model architectures are listed in Table C.1.

2) Gradient Computation: The forward pass in neural networks with sparse matrices is straightforward. Taking an FC layer as an example (convolutional layers are more complex but similar in principle), the input data is multiplied by a sparse weight, and produces a dense output to be passed to the next layer. Let  $\mathbf{u} \in \mathbb{R}^{n_{in} \times n_{out}}$  be its (sparse) weight,  $\mathbf{x} \in \mathbb{R}^{N \times n_{in}}$  be the (dense) input and  $\mathbf{y} \in \mathbb{R}^{N \times n_{out}}$  be the (dense) output, where  $n_{in}$ ,  $n_{out}$ , and N are the number of input neurons, output neurons of the FC layer, and the mini-batch size for SGD, respectively. Assuming there is no bias, the forward pass is given by  $\mathbf{y} = \mathbf{x} \cdot \mathbf{u}$ . The backward pass is slightly more complex. By simple calculations, the gradient of  $\mathbf{u}$  is given by  $\mathbf{g}_{\mathbf{u}} = \mathbf{x}^{T} \mathbf{g}_{\mathbf{y}}$ , and the gradient of  $\mathbf{x}$  (when required) is given by  $\mathbf{g}_{\mathbf{x}} = \mathbf{g}_{\mathbf{y}} \mathbf{u}^{T}$ . Here,  $\mathbf{g}_{\mathbf{y}}$  is the (dense) gradient in backpropagation fed by the next layer.

For the computation of  $g_x$ , since the weight u is sparse, we can accelerate the computation using our sparse matrix implementation.

The computation of  $g_u$  is however different. Note that both x and  $g_y$  are dense, and thus current implementations (e.g., PyTorch) first compute the dense gradient with u's dense form that has all zero values included, and then select values from the dense gradient according to u's sparse pattern. There is currently no better way to accelerate this process as far as we know. Therefore, this implementation does not improve the backward pass's speed of weights' gradient computations. For the above reason, in our implementation we collect the gradients of zero-valued components of u at the same time with no extra overhead (although those zero-valued components themselves are not updated). This characteristic is useful in our adaptive pruning procedure.

3) FLOPs Computation: Following the discussion in Section C2, we now explain the computation of FLOPs in both forward and backward passes in our models. Using convention in the literature, we consider that one addition and one multiplication each counts as a FLOP [53]. Taking the same notations and assumptions from Section C2, the FLOPs for the forward pass is  $2Nn_{in}n_{out} \times d$ , where d is the density of this FC layer. In the backward pass, the FLOPs for the gradient computation of weight u is  $2Nn_{in}n_{out}$  since the computation does not involve sparse matrices, while for the gradient of input x, the FLOPs is  $2Nn_{in}n_{out} \times d$ . Therefore, the total FLOPs for the backward pass is  $2Nn_{in}n_{out} \times (1 + d)$ , and the FLOPs for both forward and backward passes is  $2Nn_{in}n_{out} \times (1 + 2d)$ . FLOPs computation for convolutional layers is similar, so we skip this discussion.

4) Starting Reconfiguration Round: We start the reconfiguration in the initial pruning stage when the *training accuracy* on *local sample data* of the selected client exceeds 1.5 times the random guess accuracy. There are two advantages: (i) if the task is easy, reconfiguration starts early, which results in early model size decrease that saves training time; and (ii) it also avoids pruning the model too early when the prediction is still close to random guess, meaning the parameter values are still in the random initialization stage. In the further pruning stage, reconfiguration happens periodically with a fixed interval.

5) Further Details of Adaptive Pruning: In the following, we provide detailed information for the adaptive pruning procedure described in Section IV-B. In both reconfiguration and non-reconfiguration rounds, the importance measure is summed locally after every local update, until the sum is sent to the server in the next reconfiguration round. Note that in non-reconfiguration rounds, only the remaining model parameters are used in computation and exchanged between server and clients. Since the parameter set is fixed in non-reconfiguration rounds, only the *values* of the parameters need to be exchanged between the server and clients, which incurs no extra communication cost. An illustration of the two types of rounds is shown in Fig. C.1.

6) PyTorch on Raspberry Pi devices: To install PyTorch on Raspberry Pi devices, we follow the instructions described at https://bit.ly/3e6I7tG, where acceleration packages such as MLKDNN and NNPACK are disabled due to possible compatibility issues and their lack of support of sparse computation. We compare our implementation with the *plain* implementation by PyTorch without accelerations. We expect that similar results can be obtained if acceleration packages could support sparse computation. This is an active area of research on its own where methods for efficient sparse computation on both CPU and GPU have been developed in recent years. Integrating such methods into our experiments is left for future work.

 $<sup>^{2}</sup>$ In practice, we use a small perturbation instead of zero values for parameters that are added back, but we use zero values in our analysis for ease of exposition.





Fig. D.1. Linearity of average computation time vs. number of parameters.

#### D. Additional Experimental Results

1) Validation of Assumptions: Agreeing with our assumptions and analysis in Sections IV-B and V-B, we observe that the training time in each layer is generally independent of the other layers. Within each layer, the time is approximately linear with the number of parameters in the layer with sparse implementation. In Fig. D.1, we fix the parameters in other layers and increase the number of parameters in Conv-2's largest (first) FC layer and VGG-11's largest convolutional layer (last convolutional layer with 512 channels), respectively, and measure for 50 times. The  $R^2$  values of linear regression for Conv-2 and VGG-11 are 0.997 and 0.994, respectively.

2) Client Selection Results for Section VI: We present results under same settings as in Section VI, but with random client selection. We partition the IID CIFAR-10 and ImageNet-100 datasets uniformly into 100 equal-sized, non-overlapping clients. FEMNIST and CelebA are intrinsically non-IID datasets. For the FEMNIST dataset, we use its original 193-user partition; and for the CelebA dataset, since some of its users in the original partition have too few images (e.g. 4 images), we merge the 9343 persons' images into 934 clients (first 933 clients have 10 persons' images and the last client has 13 persons' images). In each round, we sample 10 clients randomly from the aforementioned partitions when using client selection. Fig. D.2 (corresponding to Fig. 5) shows the training time reduction; Fig. D.3 (corresponding to Fig. 7) shows the lottery ticket result; and Fig. D.4 (corresponding to Fig. 9) shows the model size adaptation. We observe similar behaviors as with full client participation that is described in the main paper, thus we omit further discussion in this section.

3) Convergence Accuracy Results for All Experiments: The convergence accuracies with/without client selection are shown in Table D.1. The results are taken from the test accuracies in the last five evaluations of each simulation. We see that the convergence accuracy of PruneFL is similar to that of conventional FL. Conventional FL sometimes shows a slight advantage because all methods run for the same number of *rounds*, and full-sized models in conventional FL learn faster when the accuracy is measured in *rounds* instead of time.



Fig. D.2. Test accuracy vs. time results of four datasets (client selection).





Fig. D.3. Lottery ticket results of four datasets (client selection).



Fig. D.4. Number of parameters vs. round for four datasets (client selection).

	FEMNIST		CIFAR-10		ImageNet-100		CelebA	
	No C.S.	C.S.						
Conventional FL	$85.33 \pm 0.17$	$84.61\pm0.62$	$86.30\pm0.20$	$86.51\pm0.12$	$76.36\pm0.56$	$76.62\pm0.34$	$91.41 \pm 0.15$	$91.33 \pm 0.17$
PruneFL (ours)	$85.07 \pm 0.31$	$83.90 \pm 0.48$	$85.47 \pm 0.19$	$85.50\pm0.18$	$77.23 \pm 0.25$	$76.07\pm0.31$	$91.38 \pm 0.14$	$91.2\pm0.17$
SNIP	$85.16\pm0.22$	$84.05\pm0.70$	$81.16\pm0.22$	$81.44\pm0.13$	$76.89 \pm 0.27$	$76.76\pm0.60$	$91.48 \pm 0.07$	$89.69 \pm 0.75$
SynFlow	$84.77\pm0.17$	$84.19 \pm 0.21$	$82.41 \pm 0.42$	$82.36 \pm 0.28$	$76.66 \pm 0.28$	$76.60\pm0.24$	$91.24 \pm 0.23$	$90.23 \pm 0.11$
Online learning	$85.31 \pm 0.37$	$84.30 \pm 0.27$	$10.00\pm0.00$	$10.00\pm0.00$	$75.32\pm0.18$	$74.97 \pm 0.15$	$90.39 \pm 0.13$	$88.61 \pm 0.53$
Iterative pruning	$84.87 \pm 0.22$	$84.08 \pm 0.39$	$84.45\pm0.14$	$84.64\pm0.14$	$76.18\pm0.30$	$76.96 \pm 0.26$	$90.14 \pm 0.37$	$88.65 \pm 0.33$

TABLE D.1 Average of the last 5 measured accuracies (%). C.S. stands for client selection.