Server-Side Local Gradient Averaging and Learning Rate Acceleration for Scalable Split Learning

*Shraman Pal¹, *Mansi Uniyal¹, Jihong Park ^{2,4}, Praneeth Vepakomma ³, Ramesh Raskar ³, Mehdi Bennis ⁵, Moongu Jeon ⁴, Jinho Choi ²

¹IIT Kharagpur, India, ²Deakin University, Australia, ³MIT Media Lab, USA, ⁴GIST, Korea, ⁴University of Oulu, Finland shramanpal@iitkgp.ac.in, mansinumber1@iitkgp.ac.in, jihong.park@{deakin.edu.au, gist.ac.kr}, vepakom@mit.edu, raskar@mit.edu, mehdi.bennis@oulu.fi, mgjeon@gist.ac.kr, jinho.choi@deakin.edu.au

*Equal contribution

Abstract

In recent years, there have been great advances in the field of decentralized learning with private data. Federated learning (FL) and split learning (SL) are two spearheads possessing their pros and cons, and are suited for many user clients and large models, respectively. To enjoy both benefits, hybrid approaches such as SplitFed have emerged of late, yet their fundamentals have still been illusive. In this work, we first identify the fundamental bottlenecks of SL, and thereby propose a scalable SL framework, coined SGLR. The server under SGLR broadcasts a common gradient averaged at the split-layer, emulating FL without any additional communication across clients as opposed to SplitFed. Meanwhile, SGLR splits the learning rate into its server-side and client-side rates, and separately adjusts them to support many clients in parallel. Simulation results corroborate that SGLR achieves higher accuracy than other baseline SL methods including SplitFed, which is even on par with FL consuming higher energy and communication costs. As a secondary result, we observe greater reduction in leakage of sensitive information via mutual information using SLGR over the baselines.

1 Introduction

The recent trend in deep learning has seen exponential growth in terms of architecture sizes Alom et al. (2018). In the computer vision domain, the model sizes over the years have grown larger, as observed in the transition from ResNet and VGG (He et al. 2015; Simonyan and Zisserman 2015) to Inception and DenseNet (Szegedy et al. 2015; Huang et al. 2018). In the field of natural language processing the growth is even more drastic; with introduction of transformer models like BERT (Devlin et al. 2019) that have crossed the 100 million parameter mark; and finally reaching Open AI's recent GPT-3 (Brown et al. 2020) standing at staggering 175 billion parameters. The prerequisites for running these large models are huge training data and computing power, making them still limited to a select few.

Fortunately for such large models, the current trend in data volume and computing power keeps exponentially increasing in both. Nonetheless, the majority of these data and computing power are sourced from mobile edge devices like smartphones, IoT, and e-health wearables, so are inherently dispersed and often privacy-sensitive (Park et al. 2019; Smith et al. 2018). Towards exploiting the indispensable private data and mobile computing power, recent advances in distributed learning have opened new avenues notably in the form of *federated learning (FL)* (McMahan et al. 2016) and *split learning (SL)* (Vepakomma et al. 2018).

In FL, edge devices or *clients* independently train their local models, while periodically sharing their local model parameters through a parameter server (McMahan et al. 2016). Consequently, keeping data private, FL can leverage the parallel computing power and global data of clients. However, FL right off the bat removes the chance of using large models, as edge devices cannot store and communicate them due to the strain on memory, bandwidth, and energy (Park et al. 2021). Alternatively, SL copes with large models by splitting the models into a shared upper segment and distributed lower segments that are respectively stored at the server and clients, respectively. At the split layer or cut layer, in the forward propagation, each client uploads its final hidden representations, so-called smashed data, and downloads its gradient from the server in the backward propagation, thereby keeping local data private.

Notwithstanding, SL has difficulty in achieving scalability, in contrast with FL where the accuracy increases with the number of clients (Zhang, Wei, and Berry 2021). As shown by Fig. 1, the original SL algorithm, vanilla SL or hereafter referred to as *sequential SL (SSL)*, limits the server to supporting clients one by one after each local training completes. Therefore, SSL requires tight pipelining across clients and fails to leverage parallel computing power. By simply allowing multiple client associations per batch, SSL can easily be extended to its parallel version, termed *parallel SL (PSL)*, which however entails a fundamental issue on the *server-client update imbalance* as elaborated in the following two problems.

- 1. Server-Side Large Effective Batch Problem: The server model is updated multiple times per each client model update, resulting in a larger *effective batch size* (i.e., batch size multiplied by the number of clients, when each client has the same batch size) at the server than that of each client (i.e., batch size).
- 2. Backward Client Decoupling Problem: Other clients' impacts are only reflected in the forward propagation

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Schematic illustration of SSL, PSL, SFL, and the proposed *SGLR with a split-layer gradient averaging (SplitAvg) and learning rate splitting (SplitLr)*, where the shaded area widths represent learning rates. The dotted lines and the solid lines imply forward and backward propagation flows respectively.

flows when smashed data pass through the server model, whereas the backward propagation flows across clients are decoupled with each other as illustrated in Fig. 1.

A recently proposed *SplitFed learning (SFL)* algorithm addresses the client decoupling problem by additionally applying FL across client models (Thapa et al. 2021). However, the effectiveness of SFL comes at the cost of large communication overhead, and what is more the server-side effective batch size problem remains unsolved.

To fill the aforementioned void, in this article we propose a novel PSL framework, coined *split learning with a gradient averaging and learning rate splitting (SGLR)*. At its core, SGLR aims to address the server-side effective batch size problem, by separating the learning rate into two parts, and accelerating the server model's learning rate, i.e., *learning rate splitting (SplitLr)*, inspired by the techniques for large batch training (Smith et al. 2018), (Goyal et al. 2018).

Next, for the client decoupling problem, SGLR broadcasts the common gradient averaged at the cut-layer from the server to all clients, i.e., *split-layer gradient averaging* (*SplitAvg*), rather than unicasting a unique gradient to each client. As opposed to SFL, the gradient averaging of SplitAvg is taken at the server, not incurring any additional communication cost. Even compared to PSL that relies on gradient unicasting through orthogonal bandwidth allocations, SplitAvg leverages broadcasting over the entire bandwidth, making SGLR more communication efficient and scalable particularly under limited bandwidth.

Contributions & Organization The major contributions of this work are outlined as follows.

- Motivated by the server-side large effective batch problem, we propose *SplitLr* so as to separately accelerate the learning rate at the server (see Sec. 3.1).
- Inspired by SFL while identifying its non-negligible communication overhead and possible information leakage incurred by model averaging across clients, we develop *SplitAvg* in which the server broadcasts (or multicasts) averaged gradients to clients without incurring any additional communication overhead (see Sec. 3.2).
- Combining SplitLr and SplitAvg, we finally propose *SGLR* (see Algorithm 1). Simulation results corroborate

the effectiveness of SGLR in terms of accuracy scalability, information leakage, and communication efficiency (see **Sec. 4.2**, **4.3**, and **4.4**, respectively).

The rest of this article is organized as follows. In Sec. 2, details existing works are summarized to identify the serverclient weight imbalance problem. In Sec. 3, the operations of SGLR are described by elaborating the procedures of SplitAvg and SplitLr algorithms. In Sec. 4, the effectiveness of SGLR is validated by simulation and comparison with other baseline frameworks such as *centralized learning (CL)*, FL, sequential/parallel SL, and SFL. Finally, we conclude this article by discussing several future research directions in Sec. 5.

2 Preliminaries: FL, SL, and SFL

Federated learning (McMahan et al. 2016) and Split learning (Gupta and Raskar 2018), (Vepakomma et al. 2018) are two new frameworks that allow training a model effectively from various distributed data sources without sharing the raw data. The device that possesses the data is usually termed as the *client* while the generally computationally powerful device is termed as the *server* which could have multiple functions based on the framework considered. In all the frameworks, there exist a set of clients C, where each *i*-th client has their data D_i and model weights w_i stored locally. $D = \bigcup_{i \in C} D_i$ represents the total data.

We represent the model weights of the *i*-th client as, $\mathbf{w}_{c,i} = [w_{c,i}^{L_c}, w_{c,i}^{L_c-1}, ..., w_{c,i}^1]^T$ where L_c is the cut layer. The shared server model for SL frameworks is denoted by $\mathbf{w}_s = [w_s^L, w_s^{L-1}, ..., w_s^{L_c+1}]^T$ where L is the output layer. In the next subsections, we outline existing frameworks to highlight the benefits and limitations of each method on different aspects of latency, model performance and communication efficiency.

2.1 Sequential Split Learning (SSL)

In SSL, also called vanilla SL, (Vepakomma et al. 2018), (Gupta and Raskar 2018), a lower model segment, not necessarily the same, is present in multiple clients and an upper model segment is present on a shared server. During training, clients are sequentially selected. The selected *i*-th client

selects a mini batch $\mathcal{B}_i \in \mathcal{D}_i$ which consists of $|\mathcal{B}_i|$ inputlabel tuples $(x_{i,j}, y_{i,j})$ from its local data. The *i*-th client produces $|\mathcal{B}_i|$ smashed data $s_{i,j} = f(\mathbf{w}_{c,i}, x_{i,j})$ by passing the *j*-th input data $x_{i,j}$ through it for all $j \in B_i$.

The client uploads the smashed data-label tuples to the server which produces the final predictions $y'_{i,j} = f(\mathbf{w}_s, s_{i,j})$. The loss for the batch is denoted as $L(\mathbf{w}_{c,i}, w_s)$. The optimisation objective therefore simply becomes $min_{\mathbf{w}_{c,i},\mathbf{w}_s} \{L(\mathbf{w}_{c,i}, \mathbf{w}_s)\}$.

Gradients Using the evaluated loss, the server calculates the gradients and backpropagates(BP) for its model layers:

$$\mathbf{g}_{s}^{\mathrm{SSL}} = \begin{bmatrix} g_{s}^{L} \\ g_{s}^{L-1} \\ \dots \\ g_{s}^{L_{c}+1} \end{bmatrix} = \begin{bmatrix} \nabla_{\mathbf{w}_{s}^{L}} L(\mathbf{w}_{s}, \mathbf{w}_{c,i}) \\ \mathrm{BP}(g_{s}^{L}) \\ \dots \\ \mathrm{BP}(g_{s}^{L_{c}+2}) \end{bmatrix}.$$
(1)

The server sends the gradient of the *cut_layer* $g_s^{L_c+1}$ to the *i*-th client. Then the *i*-th client generates the gradient for its own model segment:

$$\mathbf{g}_{c,i}^{\mathrm{SSL}} = \begin{bmatrix} g_{c,i}^{L_c} \\ g_{c,i}^{L_c-1} \\ \dots \\ g_{c,i}^1 \end{bmatrix} = \begin{bmatrix} \mathrm{BP}(g_s^{L_c+1}) \\ \mathrm{BP}(g_{c,i}^{L_c}) \\ \dots \\ \mathrm{BP}(g_{c,i}^2) \end{bmatrix}.$$
(2)

Weight Updates The weight updates for both the client and the server thereafter can be written in a compact form as:

$$\begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} - \eta \begin{bmatrix} \mathbf{g}_s^{\text{SSL}} \\ \mathbf{g}_{c,i}^{\text{SSL}} \end{bmatrix}.$$
(3)

After that the lower model segment $\mathbf{w}_{c,i}$ is sent to the (i+1)-th client $\mathbf{w}_{c,i+1} \leftarrow \mathbf{w}_{c,i}$ where $(i+1) \in C \setminus \{i\}$. Thereafter the *i*+1-th client continues with its local iterations.

Benefits and Limitations SSL can achieve high accuracy as the shared server benefits from multiple local datasets making it robust. Also due to the lower model weight sharing, there is information sharing between the clients directly as well. The latency, however, increases linearly with the increase in the number of clients which is a huge drawback as it fails to capture the potential of parallel computing which motivates the use of PSL.

2.2 Parallel Split Learning (PSL)

In PSL, the server and clients are the same as in SSL but all clients are connected to the server simultaneously. All the clients run their forward passes in parallel and upload their smashed data to the server together. The server generates a loss $L(\mathbf{w}_{c,i}, \mathbf{w}_s)$ corresponding to each client's smashed data in parallel.

Gradients Then the server computes the gradients for its layers by weighing the individual losses according to the size of the data of each client. The gradient for each layer of the server is cast as:

$$\mathbf{g}_{s}^{\mathrm{PSL}} = \Sigma_{i \in \mathcal{C}} \delta_{i} \begin{bmatrix} g_{s,i}^{L} \\ g_{s,i}^{L-1} \\ \vdots \\ g_{s,i}^{L_{c}+1} \end{bmatrix} = \Sigma_{i \in \mathcal{C}} \delta_{i} \begin{bmatrix} g_{s,i}^{L} \\ \mathrm{BP}(g_{s,i}^{L}) \\ \vdots \\ \mathrm{BP}(g_{s,i}^{L_{c}+2}) \end{bmatrix}, \quad (4)$$

where $\delta_i = \frac{|\mathcal{D}_i|}{|\mathcal{D}|}$. The local gradients are sent to the corresponding clients with which the clients calculate the gradients for the layers of its own model. The client-side gradient is given as:

$$\mathbf{g}_{c,i}^{\mathrm{PSL}} = \begin{bmatrix} g_{c,i}^{L_c} \\ g_{c,i}^{L_c-1} \\ \vdots \\ g_{c,i}^{1} \end{bmatrix} = \begin{bmatrix} \mathrm{BP}(g_{s,i}^{L_c+1}) \\ \mathrm{BP}(g_{c,i}^{L_c}) \\ \vdots \\ \mathrm{BP}(g_{c,i}^{2}) \end{bmatrix}.$$
(5)

Note that while the feed-forward flows propagate in parallel, the backward flows propagate sequentially.

Weight Updates (3) The weight update for PSL remains the same as (3) of SSL, i.e.,

$$\begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} - \eta \begin{bmatrix} \mathbf{g}_s^{\text{PSL}} \\ \mathbf{g}_{c,i}^{\text{PSL}} \end{bmatrix}.$$
(6)

It is worth noting that SSL is a special case of PSL where $\delta_i = 0$ and the rest are 0. In this case only the *i*-th client is given importance for the backward pass and essentially neglecting the rest.

Benefits and Limitations PSL achieves lower latency than SSL as the forward pass happens in parallel utilising the parallel computing of edge devices. However, the performance falls below SSL due to the server-client update imbalance problem. The client's local gradient is dependent only on a single batch of input data while, unlike SSL, the server's gradient depends on multiple batches. The multiple batches at the server also lead to a high effective batch size. Also as the clients only utilise their local gradient, the clients get *detached* during the backward pass, which we term as the *backward client decoupling* problem. We elaborate on these in the later sections.

2.3 Federated Learning (FL)

In FL or specifically FedAvg, every client runs an exact copy of the entire model on its own local data where each client *i* focuses on the local optimization task $\min_{w_{c,i}} \{L(\mathbf{w}_{c,i})\}$. The server in this framework only acts as an aggregator. It receives the local model weights from each client and takes a weighted mean (*FedAvg*) to get the *global weights*. These global weights are downloaded by every client. The clients resume training using these global weights. This process continues till convergence. Only *FedAvg* has been considered as it is the closest comparison to our work.

Gradients FL does not consider model-split architectures, and the server does not train any models. Therefore there are no gradients at the server side. The gradients at the client side are given as

$$\mathbf{g}_{c,i}^{\mathsf{FL}} = \begin{bmatrix} g_{c,i}^{L} \\ g_{c,i}^{L-1} \\ \vdots \\ g_{c,i}^{1} \end{bmatrix} = \begin{bmatrix} \nabla_{w_{c,i}^{L}} L(w_{c,i}) \\ \mathbf{BP}(g_{c,i}^{L}) \\ \vdots \\ \mathbf{BP}(g_{c,i}^{2}) \end{bmatrix}.$$
(7)

Note again that each client under FL stores L layers, i.e., an entire model. This is in contrast to the client under SL storing L_c layers while the server stores the remaining $L - L_c$ layers.

Weight Updates The weight update using these gradients can be written in a compact form as:

$$\mathbf{w}_{c,i} \leftarrow \Sigma_{k \in \mathcal{C}} \delta_k \{ \mathbf{w}_{c,k} - \eta \mathbf{g}_{c,k}^{\mathsf{FL}} \}.$$
(8)

In other words, once each client has updated its local weights $\mathbf{w}_{c,i}$, they upload it to the server which averages them. Afterwards, the clients download the averaged weights and assigns them to their local weights for the next iteration.

Benefits and Limitations FL achieves high accuracy, comparable to CL while maintaining a low latency due to its parallel computing clients. With small models, the communication overhead is low as well. However, with larger models, not only would we require powerful devices with large storage capacity but it would also result in large communication overhead. Such computationally powerful devices are infeasible when we consider edge devices as clients.

2.4 Split Federated Learning (SFL)

SFL (Thapa et al. 2021) combine FL and SL by averaging the model weights after updating the weights similar to FL but only limited to the lower model segments stored on the clients. Local Model Weight Averaging or *LocAvg* is used to solve the backward client decoupling problem. The forward pass, the loss $L(\mathbf{w}_{c,i}, \mathbf{w}_s)$ computation at the server-side, and the backward pass remains the same as PSL. The extra step is to average the client model weights in a layer-wise fashion.

Gradients Similar to PSL, the server gradient remains the same $g_s^{SFL} = g_s^{PSL}$. The client gradients can be written equivalently as a weighted sum of all the gradients instead of a weighted sum of the weights. Accordingly, the client gradients are cast as:

$$\mathbf{g}_{c,i}^{\mathrm{SFL}} = \Sigma_{k\in\mathcal{C}}\delta_{k} \begin{bmatrix} g_{c,k}^{L_{c}} \\ g_{c,k}^{L_{c}-1} \\ \vdots \\ g_{c,k}^{1} \end{bmatrix} = \Sigma_{k\in\mathcal{C}}\delta_{k} \begin{bmatrix} \mathrm{BP}(g_{s,k}^{L_{c}+1}) \\ \mathrm{BP}(g_{c,k}^{L_{c}}) \\ \vdots \\ \mathrm{BP}(g_{c,k}^{2}) \end{bmatrix}.$$
(9)

We emphasize that the actual framework, averages the weights but as long as the weight averaging happens after each client weight update, they can be written as a weight update with averaged gradients.

Weight Updates The weight updates of SFL can be simplified and written in a similar form of (3) as follows:

$$\begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} - \eta \begin{bmatrix} \mathbf{g}_s^{\text{SFL}} \\ \mathbf{g}_{c,i}^{\text{SFL}} \end{bmatrix}.$$
(10)

Note here that the averaging comes after completing the backpropagation across all the layers.

Benefits and Limitations SFL keeps the latency improvement of PSL by running computations in parallel. By employing the idea of FL, specifically *LocAvg*, it solves the backward client decoupling problem by sharing direct information across the clients. However, this leads to a larger communication overhead than either PSL or FL. Not to mention that the large effective batch size problem at the server remains unattended. In the next section, we aim to tackle all the problems using our proposed novel framework.

3 SGLR: PSL with Split-Layer Gradient Averaging and Learning Rate Splitting

We utilise ideas from existing learning rate acceleration work (Goyal et al. 2018), (Krizhevsky 2014), to derive a new scaling rule that works on the notion of effective batch size at the server-side and we call it *SplitLr*. We also derive ideas from FL to adopt a new method to average gradients only at the cut layer, called *SplitAvg* to solve the backward pass decoupling problem.

The forward pass at the clients is the same as PSL. As the server concatenates the smashed data along the batch dimension, it increases the effective batch size at the server which denotes how many samples pass through the server in one forward pass. The effective batch size at the server increases linearly with the number of clients. The *concatenated* batch at the server denoted by \mathcal{B}_s has an effective batch size of

$$|\mathcal{B}_s| = \Sigma_{i \in \mathcal{C}} |\mathcal{B}_i|, \tag{11}$$

where \mathcal{B}_i denotes the batch for the *i*-th client as used before.

Following the forward pass through the server, we obtain one loss value $L(\mathbf{w}_{c,i}, \mathbf{w}_{c,2}, ..., \mathbf{w}_{c,|\mathcal{C}|}, \mathbf{w}_s)$ for all the clients combined unlike PSL which generated a loss value for each client. The single loss value is used to generate the gradients of the server layers and the local gradients of the clients. Next, we sample a subset of clients and term them *active clients* whose local gradients are averaged which we term as *SplitAvg*. This averaged gradient is substituted for the local gradient of the active clients and the remaining clients utilise their corresponding local gradients.

This tackles the backward pass decoupling problem as the averaged gradient now contains information of multiple clients. This allows information to flow from one client to another indirectly without requiring the communication of client model weights. This framework results in a scalable system with less communication overhead than any SL framework. We present the two algorithms *SplitLr* and *SplitAvg* in the next two subsections along with detailed gradient calculation and weight updates.

3.1 SplitLr: Learning Rate Splitting

In the existing literature, there are several learning rate schemes based on the batch size (Goyal et al. 2018), (Krizhevsky 2014) which outperform the baseline. They show that for optimizers like Stochastic Gradient Descent (SGD) a linear scaling rule $\eta = \eta_0 |\mathcal{B}_{\rangle}|$ is useful whereas for optimizers like Adam a square root scaling rate $\eta = \eta_0 (|\mathcal{B}_{\rangle}|)^{0.5}$ achieves better performance. We do not use different batch sizes at the client side but as shown in (11), the effective batch size in the server is different. Therefore we use the learning rate scaling only for the server-side model.

The utilisation of a different learning rate scheme for a segment of the model is new, which forces us to assume a generic scaling rule and perform experimental simulations to identify the optimal scaling rule. The learning rate scaling rule can be reduced a power-law model:

$$\eta_s = \eta_0 |\mathcal{B}_s|^{\alpha}. \tag{12}$$

The hyperparameter α is iterated with values ranging from 0.0 to 2.0 with a step size of 0.5. The case of $\alpha = 0.0$ is the setting of no SplitLr.

We extend the equation further to include a special case where each client has the same batch size $|\mathcal{B}_i| = b$ which simplifies $|\mathcal{B}_s| = b|\mathcal{C}|$ and the learning rate scheme becomes:

$$\eta_s = \eta_0(b|\mathcal{C}|)^\alpha = \eta_c(|\mathcal{C}|)^\alpha, \tag{13}$$

where $\eta_c = \eta_0 b^{\alpha}$ as *b* is a constant and can be merged into η_0 . This equation shows scaling that is proportional to the number of clients rather than the batch size at the client side. We utilise this equation for our simulation runs. By combining PSL with SplitLr, we name the framework as *SLR*.

The gradient at the server side and client side for SLR is the same as (4) and (5). Therefore $g_s^{SLR} = g_s^{PSL}$ and $g_{c,i}^{SLR} = g_{c,i}^{PSL}$. The modification lies in the weight update step where it becomes:

$$\begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i} \end{bmatrix} - \begin{bmatrix} \eta_s \mathbf{g}_s^{\text{SLR}} \\ \eta_c \mathbf{g}_{c,i}^{\text{SLR}} \end{bmatrix}.$$
(14)

3.2 SplitAvg: Split Layer Gradient Averaging

Unlike SFL, which averages model weights, we aggregate the local gradients of a subset of clients and term it *SplitAvg*. Averaging the gradients decreases the communication overhead compared to any SL framework. The accuracy also increases as we present in Tab. 4. The averaging of the local gradients allows the subset of participating clients to share some information in the form of gradients throughout the training which tackles the backward client decoupling problem. By combining PSL with SplitAvg We name this framework as *SGL*.

The gradients for the server remain as PSL in (4) and can be simply stated as $\mathbf{g}_s^{\text{SGL}} = \mathbf{g}_s^{\text{PSL}}$ but the gradients of the client models are different. For the active clients denoted by $C_a = \phi C$ where ϕ denotes the fraction of total clients which are active, the client-side gradients are given by:

$$\mathbf{g}_{c,i,a}^{SGL} = \begin{bmatrix} g_{c,i,a}^{L_c} \\ g_{c,i,a}^{L_c-1} \\ \vdots \\ g_{c,i,a}^{1} \end{bmatrix} = \begin{bmatrix} \mathsf{BP}(\Sigma_{k\in\mathcal{C}_a}g_{s,k}^{L_c+1}) \\ \mathsf{BP}(g_{c,i,a}^{L_c}) \\ \vdots \\ \mathsf{BP}(g_{c,i,a}^{L_c}) \\ \vdots \\ \mathsf{BP}(g_{c,i,a}^{2}) \end{bmatrix}, \quad (15)$$

where subscript *a* identifies the client activation, and the *i*th client belongs to the set of active clients C_a . For clients which are not active have the same gradients as in PSL with (5). Their gradients $\mathbf{g}_{c,i}^{SGL} = \mathbf{g}_{c,i}^{PSL}$ as they are equal. The weight update for the active clients utilise this av-

The weight update for the active clients utilise this averaged gradient while the rest use the corresponding local gradients:

$$\begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i,a} \\ \mathbf{w}_{c,i} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w}_s \\ \mathbf{w}_{c,i,a} \\ \mathbf{w}_{c,i} \end{bmatrix} - \eta \begin{bmatrix} \mathbf{g}_s^{\text{SGL}} \\ \mathbf{g}_{c,i,a}^{\text{SGL}} \\ \mathbf{g}_{c,i}^{\text{SGL}} \end{bmatrix}.$$
(16)

It is worth noting that in SFL the averaging occurs after the total backpropagation has taken place while in SGL the gradient is averaged before the backpropagation at the clientside takes place. Taking the fraction $\phi = 1.0$ can lead to the training not converging as we show in the results section. Therefore we also use a phased training strategy that utilises the averaged gradient for a certain phase before switching to the local gradients completely or vice-versa. SplitAvg also acts as an inherent regularizer which helps prevent over-fitting. In the Fig. 4 we can observe that as we increase the fraction of clients which participate in the local gradient averaging, the edges or the activation of the image become less sharp i.e the model becomes more robust capable of classifying unseen images better.

3.3 SGLR: PSL with SplitLr & SplitAvg

The two algorithms combined give the final framework SGLR. The SplitLr method addresses the large effective batch size problem while the SplitAvg addresses the backward decoupling problem. The final weight update by combining both the methods is given as:

$$\begin{bmatrix} \mathbf{w}_{s} \\ \mathbf{w}_{c,i,a} \\ \mathbf{w}_{c,i} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w}_{s} \\ \mathbf{w}_{c,i,a} \\ \mathbf{w}_{c,i} \end{bmatrix} - \begin{bmatrix} \eta_{s} \mathbf{g}_{s}^{\text{SGL}} \\ \eta_{c} \mathbf{g}_{c,i,a}^{\text{SGL}} \\ \eta_{c} \mathbf{g}_{c,i}^{\text{SGL}} \end{bmatrix} .$$
(17)

Here, the *i*-th client is not active for all iterations.

The crux of our changes has only been in the gradients of the clients and the learning rates schemes at the server. The performance improvement of our method over the other existing works barring FL as shown in Tab. 4, highlights the problem of backward decoupling existing in SL. It might be interesting to note that it is quite possible that averaging only a subset of client model weights in SFL might lead to higher performance. The detailed algorithm combining both the methods has been outlined in 1.

The next section presents our experimental results in terms of accuracy, information leakage, and communication efficiency.

4 Simulation Results

This section contains the results for the different runs we perform with our architecture. The section validates that our architecture and training strategies are scalable while maintaining a very low computational and memory requirement.

4.1 Experimental Settings

The default setting for all the experiments in this section is using a Resnet16 architecture where the first 9 layers are stored on each of the client devices and the last 7 layers are stored on the server. The division has been done keeping in mind that the later layers have a much larger number of channels and require larger memory which is suitable for the server. We conduct our runs on the Fashion MNIST dataset (Xiao, Rasul, and Vollgraf 2017) which contains 60000 training images and 10000 test images.

We separate 10000 images for the validation set from the training set such that it is i.i.d in nature. Each client possesses 1000 i.i.d samples which are randomized on each run. We use Adam (Kingma and Ba 2017) with a client learning rate of 10^{-3} and default momentum values of $\beta_1 = 0.9$, $\beta_2 = 0.999$ as our optimizer and vary the server learning

Algorithm 1: SGLR: Parallel SL with SplitLr & SplitAvg

1: SplitLr 2: $\eta_c = \eta_0$ {Client Lr} 3: $\eta_s = \eta_0 C^{\alpha}$ {Server Lr} 4: for $epoch \leftarrow (0 - E)$ do 5: /*Runs on clients*/ 6: for each client $i \in C$ in parallel do 7: $s_i \leftarrow f(\mathbf{w}_{c,i}, \mathcal{B}_i)$ 8: Upload: (s_i, Y_i) from client_i to server 9: end for 10: /*Runs on server*/ $s \leftarrow [s_1, s_2, \cdots, s_{|\mathcal{C}|}] \\ Y \leftarrow [Y_1, Y_2, \cdots, Y_{|\mathcal{C}|}]$ 11: 12: $Y_p \leftarrow f(\mathbf{w}_s, s)$ 13: Calculate loss $L(\mathbf{w}_{c,1},\cdots,\mathbf{w}_{c,|\mathcal{C}|},\mathbf{w}_s)$ 14: Calculate gradient $\mathbf{g}_s^{\text{SGLR}}$ 15: Weight Update $\mathbf{w}_s \leftarrow \mathbf{w}_s - \eta_s \mathbf{g}_s^{\text{SGLR}}$ 16: 17: SplitAvg 18: $C_a \leftarrow \phi C$ {Sampled active clients} Download average gradient $\Sigma_{k \in C_a} \mathbf{g}_{s,k}^{L_c+1}$ 19: Download local gradient $\mathbf{g}_{s,k}^{L_c+1}$ for $k \in \mathcal{C} \setminus \mathcal{C}_a$ 20: /*Runs on clients*/ 21: 22: for each client $i \in C$ in parallel do 23: if $i \in C_a$ then Calculate gradient $\mathbf{g}_{c,i,a}^{\text{SGLR}}$ 24: Weight Update $\mathbf{w}_{c,i} \leftarrow \mathbf{w}_{c,i} - \eta_c \mathbf{g}_{c,i,a}^{\text{SGLR}}$ 25: else 26: Calculate gradient $\mathbf{g}_{c,i}^{\mathrm{SGLR}}$ 27: Weight Update $\mathbf{w}_{c,i} \leftarrow \mathbf{w}_{c,i} - \eta_c \mathbf{g}_{c,i}^{\text{SGLR}}$ 28: 29: end if end for 30: 31: end for

rate according to the SplitLr setting. We utilise a batch size of 8 taking into consideration that edge devices do not possess much RAM. Having laid down the basic settings of our experimental runs we move to the actual results.

4.2 Accuracy Scalability

As the number of clients increases, we present the top-1 accuracy across different settings, where the baseline is when no algorithm, SplitAvg or SplitLr, is used. We look at the individual impact of SplitLr and SplitAvg before observing how combining the two algorithms performs.

Impact of SplitLr The SplitLr algorithm aims to utilise the large effective batch size at the server to its advantage. The use of SplitLr improves on the results providing higher accuracy across the board leading to nearly a 1.5% increase for 20 clients. The results for learning rate splitting have been shown in Tab. 1. As observed in Tab. 1, large values of α do well only for the small number of clients while smaller values perform well for a large number of clients. This is



Figure 2: Accuracy convergence with varied learning rate power for 6 worker clients.



Figure 3: Convergence Analysis for different settings of SplitAvg. Final scores are higher than baseline ($\phi = 0.0$) as well as the score at any given point of training.

quite understandable as (Kingma and Ba 2017) states that Adam is quite sensitive to learning rate and too large values often leads to instability. As an example using α close to 2.0, for 20 clients, the learning rate becomes 400 times larger. For a client learning rate of 10^{-3} , the server learning rate becomes 0.4 which is very large for Adam and is not used in the existing literature. Therefore for large clients, we suggest using smaller values of α while for a smaller number of clients larger values can be utilised.

С	w.o. SplitLr	SplitLr					
		α=2.0	α=1.5	α =1.0	α =0.5		
1	82.45	82.45	82.45	82.45	82.45		
2	82.04	82.11	82.00	81.68	81.76		
4	82.68	82.14	83.42	82.93	82.46		
6	82.84	83.56	84.21	84.58	83.30		
8	83.49	58.41	83.87	84.52	84.23		
10	83.49	34.37	84.21	84.98	84.78		
15	84.34	-	84.08	85.79	85.30		
20	84.88	-	-	86.00	86.25		

Table 1: Top-1 accuracy for different SplitLr settings. Lower values of α perform better as the number of clients increase.

Impact of SplitAvg The SplitAvg algorithm solves the backward client decoupling problem. We provide results for a fraction ϕ of 0.25, 0.50 and 0.75. As observed in the Tab. 2, using SplitAvg outperforms the baseline. For 20 clients,

С	C Constant learning rate			SplitLr $\alpha = 0.5$			SplitLr $\alpha = 1.0$					
	w.o. SplitAvg	SplitAvg		w.o. SplitAvg	SplitAvg		w.o. SplitAvg	SplitAvg				
		$\phi = 0.25$	$\phi = 0.5$	$\phi = 0.75$		$\phi = 0.25$	$\phi = 0.5$	$\phi = 0.75$		$\phi = 0.25$	$\phi = 0.5$	$\phi = 0.75$
1	82.45	82.45	82.45	82.45	82.45	82.45	82.45	82.45	82.45	82.45	82.45	82.45
2	82.04	81.91	81.84	81.94	81.76	81.88	81.93	81.92	81.68	81.95	82.10	81.72
4	82.68	82.90	83.16	83.81	82.46	82.78	83.24	84.22	82.93	83.30	83.66	84.45
6	82.84	83.20	84.02	84.08	83.30	83.06	84.89	84.62	84.28	83.70	84.32	84.24
8	83.49	83.73	84.27	85.14	84.23	84.33	85.04	85.29	84.52	84.58	85.41	85.68
10	83.49	84.35	84.74	85.49	84.78	84.65	85.22	85.93	84.98	84.82	85.89	86.11
15	84.34	85.87	85.86	85.86	85.30	85.89	86.71	87.22	85.79	85.19	86.15	86.20
20	84.88	86.33	86.59	86.82	86.25	86.96	87.48	87.47	86.00	86.54	87.11	86.67

Table 2: Top 1 Accuracy on FashionMNIST for different settings of SGLR.

the accuracy increases by nearly 2%. Using $\phi = 0.75$ performs, in general, the best for the setting with no SplitLr. It is also interesting to note that the increase in the accuracy is more as the number of clients increases when SplitAvg is used. This points towards the potential of SplitAvg when the number of clients is large.

Integrated Impact of SplitLr & SplitAvg Tab. 2 shows the results by combining the two methods of SplitLr and SplitAvg. Both algorithms show scalability individually as well as collectively. For the SplitLr with $\alpha = 0.5$ SplitAvg with $\phi = 0.75$ achieves the highest accuracy in nearly all the experiments. For SplitLr with $\alpha = 1.0$ the SplitAvg with $\phi = 0.75$ again nearly achieves highest accuracy for all experiments. It is interesting to note that for the larger number of clients, $\phi = 0.50$ performs on par with $\phi = 0.75$ which might indicate that as the number of clients increases, it is better to reduce ϕ to keep the number of active clients in a suitable range. The accuracy when combining SplitLr and SplitAvg is higher than the set with only one algorithm. This indicates that they can be stacked on top of one another. To draw a comparison, for 20 clients, without either of the methods, the accuracy is 84.88%, with SplitAvg($\phi = 0.5$) the top-1 accuracy increases to 86.59% which is +1.71%. Further using SplitLr($\alpha = 0.5$) the performance increases to 87.48% which is +0.89% over the setting of only using SplitAvg. Therefore there is a total increase of 2.60% by using both SplitLr and SplitAvg. In the special case, when all the local gradients are averaged, the training is unstable and does not converge as seen in Tab. 3. Using the phased training approach is better but the performance improvement is still lacking. Future work on improving gradient averaging to incorporate more clients will be useful in this respect.

С	w.o. SplitAvg	All SplitAvg						
_		Initial 60%	Initial 40%	Final 60%	Final 40%	Full		
1	82.45	82.45	82.45	82.45	82.45	82.45		
2	81.68	81.99	81.79	82.21	81.47	81.82		
4	82.93	83.49	82.96	83.22	82.83	82.07		
8	84.52	84.12	84.97	84.48	83.56	83.62		
10	84.98	84.76	85.32	84.47	84.31	81.61		
15	85.59	85.12	85.20	85.05	85.68	69.45		
20	86.00	85.23	86.12	85.62	86.53	60.67		

Table 3: SplitAvg with phased training for SplitLr with $\alpha = 1.0$. SplitAvg increases the performance only with phased training but fails to converge without it.

Accuracy Comparison with Existing Methods We now compare our proposed framework, with other existing works which we had outlined in the Preliminaries section. We present the top 1 accuracy and comput. energy per client in Tab. 4 on the Fashion MNIST dataset (Xiao, Rasul, and Vollgraf 2017) by using AlexNet. The splits of the model are the same for every SL framework where the client model only comprises of 1 convolutional layer and 1 max-pooling layer while the remaining layers are stored in the server. The number of clients has been set to 5 and other hyperparameters have been kept the same for a fair comparison. Compared to the SL frameworks, as show by Tab. 4, our method achieves superior accuracy. FL and SGLR are comparable in terms of accuracy to CL which is often considered to be the ceiling. This motivates the question of which method to use. The answer is guided by the communication efficiency and training time which is explored in the Communication Efficiency subsection. The comput. energy was calculated in peta-FLOPS (PFLOPS) as suggested in OpenAI (Dario Amodei 2018). In terms of comput. energy per client, SL frameworks require nearly 94.7% less energy per client than FL due to the minimal operations carried out at the client side. Therefore CL and FL achieve higher performance at the cost of higher energy consumption.

Method	CL	FL	SGLR	SFL	PSL
Top 1 Acc.	90.5	89.7	89.3	86.0	84.7
Comput. per Client (PFLOPS)	2.944	0.588	0.031	0.031	0.031

Table 4: Top 1 Accuracy on Fashion MNIST dataset with AlexNet. FL and SGLR perform the closest to the CL. FL and CL have large comput. energy cost compared to SL frameworks.

4.3 Information Leakage

Another aspect that we analyse is privacy or information leakage. As described in (Geiping et al. 2020), one may extract sensitive information from the intermediate layers to reconstruct and obtain the original data, breaching the privacy of the data set. Therefore we assume a *honest-but-curious* server which attempts to reconstruct the image from the received smashed data after training. We state that SplitAvg creates models whose smashed data are harder to reconstruct the original image with thereby increasing the data privacy.



Figure 4: Intermediate Representations of the trained ResNet model. The different variations of using Local Gradient Averaging for constant server learning rate.

Fraction ϕ	0.00	0.25	0.50	0.75	1.00
Loss	0.2153	0.2142	0.2166	0.2296	0.2759

Table 5: Value of cross-entropy loss between the reconstructed image and original image for different values of ϕ .

To validate our claim, we use mutual information as a proxy. For two discrete variables X and Y, whose joint probability is given by $P_{XY}(x, y)$, the mutual information denoted by I(X;Y) can be calculated as

$$I(X;Y) = \Sigma_{x,y} \mathsf{P}_{X,Y}(x,y) \log \frac{\mathsf{P}_{X,Y}(x,y)}{\mathsf{P}_X(x)\mathsf{P}_Y(y)}$$
(18)

It can also be expressed as the expectation of $\log \frac{P_{X,Y}(x,y)}{P_X(x)P_Y(y)}$. Similar to (Wang et al. 2021), we utilise the decoder to find the loss between the original image and the reconstructed image formed from the smashed data. We utilise the decoder of a Resnet AutoEncoder which was pretrained on CIFAR-10 from PyTorch Lightning bolts(Falcon and Cho 2020).

We calculate the loss for different variations of SplitAvg by changing the value of ϕ and selecting 1 client in random out of 6 clients. We can thus infer from the loss values are tabulated in Tab. 5, that with an increase in the fraction ϕ there is, an increase in the loss indicating the increase in hardness to generate an accurate reconstructed image. This points towards less information leakage and an increase in privacy. The initial slow increase can be credited to the small number of clients as with $\phi = 0.25$, only 1 client is an active client which practically has no effect. With $\phi = 0.5$, 3 clients are active leading to a marginal increase in loss. With $\phi = 1.0$ all 6 clients are active when there is quite a large jump in the loss indicating less information leakage.

4.4 Communication Efficiency

We now look at the communication efficiency of the different frameworks by making a quantitative analysis of the different frameworks in terms of Training Time and Communication per Epoch in Tab. 6 where the cut layer output size (MB) is denoted as S(L), the client model size (MB) in FL is denoted by S(w) and in SL is denoted by $S(w_c)$. T is the time required for a forward and backward pass and R is the rate at which the communication occurs.

In SGLR, the number of data samples per client is $\frac{|\mathcal{D}|}{|\mathcal{C}|}$ and the payload size of the smashed data per client per iteration is $\frac{|\mathcal{D}|S(L)}{|\mathcal{C}|}$. On the other hand, the gradient payload size per iteration is $(1 - \phi) \frac{|\mathcal{D}|S(L)}{|\mathcal{C}|}$ for the clients that use



Figure 5: Training Time for SGLR, FL, SL, and SFL for different datasets sizes including 25000 (Cityscapes), 50000 (MNIST), 500000 (IMDB-Wiki) and 2000000 (ImageNet). 3 settings varying model sizes and number of clients has been considered for each framework.

Method	Comm. per Client	Tot. Comm.	Tot. Training Time
FL	$2 \mathbf{S}(\mathbf{w}) $	$2 \mathcal{C} \mathbf{S}(\mathbf{w}) $	$T + 2 \frac{ \mathbf{S}(\mathbf{w}) }{R}$
PSL	$\frac{2 \mathcal{D} S(L)}{ \mathcal{C} }$	$2 \mathcal{D} \mathbf{S}(L)$	$T + 2 \frac{ \mathcal{D} S(L)}{R} +$
SFL	$\frac{2 \mathcal{D} \mathbf{S}(L)}{ \mathcal{C} } + 2 \mathbf{S}(\mathbf{w}_c) $	$2 \mathcal{D} \mathbf{S}(L) + 2 \mathcal{C} \mathbf{S}(\mathbf{w}_c) $	$T + 2\frac{ \mathcal{D} \mathbf{S}(L) }{ \mathcal{C} R} + 2\frac{ \mathbf{S}(\mathbf{w}_c) }{R}$
SGLR	$\frac{(2-\phi) \mathcal{D} \mathbf{S}(L)+\mathbf{S}(L)}{ \mathcal{C} }$	$(2-\phi) \mathcal{D} \mathbf{S}(L) + \mathbf{S}(L)$	$T + \frac{(2-\phi) \mathcal{D} \mathbf{S}(L) + \mathbf{S}(L)}{ \mathcal{C} R}$

Table 6: Communication overhead and total training time analysis of the different training methods of FL, PSL, SFL, and SGLR.

the non averaged gradients and S(L) for the clients that use the averaged gradient. The total communicated overhead per iteration *e* per client *i* denoted by γ is

$$\gamma_{i,e} = \frac{(2-\phi)|\mathcal{D}|\mathbf{S}(L) + \mathbf{S}(L)}{|\mathcal{C}|}$$
(19)

and the communication overhead per epoch for all the client becomes

$$\gamma_e = \sum_{i \in \mathcal{C}} \gamma_{i,e} = (2 - \phi) |\mathcal{D}| \mathbf{S}(L) + \mathbf{S}(L)$$
(20)

Analyzing the decrease in total communication overhead of SGLR to SFL for a given setting of 100 clients, $\phi = 0.5$, $|\mathcal{D}| = 50000$ samples and the AlexNet model where S(L) =0.024MB and $S(\mathbf{w}) = 200$ MB and $S(\mathbf{w}_c) = 67$ MB, the percentage reduction in communication of SGLR with respect to SFL is 88.6%. Similarly for the case of FL, the percentage decrease is 95.499%. There we achieve comparable performance at higher communication efficiency. It is more efficient than FL when model sizes are large and local data are not too big.

Another important aspect is the training time. Fig. 5 demonstrates the training time by varying environment parameters. We assumed S(L), and the size of client model to be the same as with the analysis of communication efficiency. With large model sizes, SGLR is more efficient than FL while having comparable accuracy but FL wins when size of the data grows large (ImageNet) though it can be offset when the number of clients is large as well which is usually the case in practical cases when edge devices are considered.

5 Conclusion

In this work, we developed novel learning rate splitting (SplitLr) and split-layer gradient averaging (SplitAvg) algorithms to resolve the two-fold fundamental problem of parallel SL. In the forward propagation, a parallel SL architecture serves multiple clients using a common server model at which the effective batch size increases proportionally with the number of clients, making the server's learning rate slower than desired. In the backward propagation, each client model update is decoupled with other client models, yielding fewer gains from the federation as compared with other scalable distributed learning frameworks such as FL. Motivated by this, we proposed a scalable SL framework SGLR laid by SplitLr and SplitAvg that address such serverside large batch problems and the backward client decoupling problem, respectively. Simulation results corroborated that SGLR yields less information leakage and higher communication efficiency than the standard parallel SL. This advantage comes mainly from SplitAvg allowing to multicast averaged gradients to multiple clients in common. Combining SplitAvg with SplitLr, SGLR outshines the standard parallel SL in terms of accuracy and scalability for different numbers of clients. The accuracy of SGLR is even on par with FL that may incur larger communication costs and huge information leakage into model-inversion attackers. Extending SplitLr, it is worth developing a learning rate splitting method for cyclic learning rates. Convergence and differential privacy analysis on SplitAvg could also be interesting topics for future work.

Acknowledgement

This research was supported in part by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (DP200100391), in part by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2014-3-00077, AI National Strategy Project), and in part by EU-CHISTERA project LeadingEdge, CONNECT, and 6G Flagship (6GENESIS).

References

Alom, M. Z.; Taha, T. M.; Yakopcic, C.; Westberg, S.; Sidike, P.; Nasrin, M. S.; Esesn, B. C. V.; Awwal, A. A. S.; and Asari, V. K. 2018. The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches. arXiv:1803.01164.

Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. arXiv:2005.14165.

Dario Amodei, D. H. 2018. AI and Compute. https://openai. com/blog/ai-and-compute/. Accessed: 2021-11-30.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805. Falcon, W.; and Cho, K. 2020. A Framework For Contrastive Self-Supervised Learning And Designing A New Approach. *arXiv preprint arXiv:2009.00104*.

Geiping, J.; Bauermeister, H.; Dröge, H.; and Moeller, M. 2020. Inverting Gradients – How easy is it to break privacy in federated learning? arXiv:2003.14053.

Goyal, P.; Dollár, P.; Girshick, R.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2018. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. arXiv:1706.02677.

Gupta, O.; and Raskar, R. 2018. Distributed learning of deep neural network over multiple agents. arXiv:1810.06060.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep Residual Learning for Image Recognition. arXiv:1512.03385.

Huang, G.; Liu, Z.; van der Maaten, L.; and Weinberger, K. Q. 2018. Densely Connected Convolutional Networks. arXiv:1608.06993.

Kingma, D. P.; and Ba, J. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980.

Krizhevsky, A. 2014. One weird trick for parallelizing convolutional neural networks. arXiv:1404.5997.

McMahan, H. B.; Moore, E.; Ramage, D.; and y Arcas, B. A. 2016. Federated Learning of Deep Networks using Model Averaging. *ArXiv*, abs/1602.05629.

Park, J.; Samarakoon, S.; Bennis, M.; and Debbah, M. 2019. Wireless Network Intelligence at the Edge. *Proceedings of the IEEE*, 107(11): 2204–2239.

Park, J.; Samarakoon, S.; Elgabli, A.; Kim, J.; Bennis, M.; Kim, S.; and Debbah, M. 2021. Communication-Efficient and Distributed Learning Over Wireless Networks: Principles and Applications. *Proceedings of the IEEE*, 109(5): 796–819.

Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556.

Smith, S. L.; Kindermans, P.-J.; Ying, C.; and Le, Q. V. 2018. Don't Decay the Learning Rate, Increase the Batch Size. arXiv:1711.00489.

Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2015. Rethinking the Inception Architecture for Computer Vision. arXiv:1512.00567.

Thapa, C.; Chamikara, M. A. P.; Camtepe, S.; and Sun, L. 2021. SplitFed: When Federated Learning Meets Split Learning. arXiv:2004.12088.

Vepakomma, P.; Gupta, O.; Swedish, T.; and Raskar, R. 2018. Split learning for health: Distributed deep learning without sharing raw patient data. arXiv:1812.00564.

Wang, Y.; Ni, Z.; Song, S.; Yang, L.; and Huang, G. 2021. Revisiting Locally Supervised Learning: an Alternative to End-to-end Training. In *International Conference on Learning Representations*.

Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. arXiv:1708.07747.

Zhang, M.; Wei, E.; and Berry, R. 2021. Faithful Edge Federated Learning: Scalability and Privacy. arXiv:2106.15905.