EFFICIENT AND LIGHT-WEIGHT FEDERATED LEARNING VIA ASYNCHRONOUS DISTRIBUTED DROPOUT

Anonymous Authors[1]

ABSTRACT

Asynchronous learning protocols have regained attention lately, especially in the Federated Learning (FL) setup, where slower clients can severely impede the learning process. Herein, we propose `AsyncDrop`, a novel asynchronous FL framework that utilizes dropout regularization to handle device heterogeneity in distributed settings. Overall, `AsyncDrop` achieves better performance compared to state of the art asynchronous methodologies, while resulting in less communication and training time overheads. We implement our approach and compare it against other asynchronous baselines, both by design and by adapting existing synchronous FL algorithms to asynchronous scenarios. Empirically, `AsyncDrop` reduces the communication cost and training time, while matching or improving the final test accuracy in diverse non-i.i.d. FL scenarios.

1 INTRODUCTION

Background on Federated Learning. Federated Learning (FL) (McMahan et al., 2017; Li et al., 2018; Karimireddy et al., 2019) is a distributed learning protocol that has witnessed fast development the past demi-decade. FL deviates from the traditional distributed learning paradigms and allows the integration of edge devices —such as smartphones (Stojkovic et al., 2022), drones (Qu et al., 2021), and IoT devices (Nguyen et al., 2021)— in the learning procedure.

Yet, such real-life, edge devices are extremely *heterogeneous* (Wang et al., 2021): they have drastically different specifications in terms of compute power, device memory and achievable communication bandwidths. Directly applying common *synchronized* FL algorithms –such as FedAvg and FedProx (Li et al., 2018; McMahan et al., 2017) that require full model broadcasting and global synchronization– results often in a "stragglers" effect (Nguyen et al., 2022; Huba et al., 2022; Tandon et al., 2017).

The ubiquitous synchronous training. One way to handle such issues is by utilizing *asynchrony* instead of *synchrony* in the learning process. To explain the main differences, let us first set up the background. In a synchronous distributed algorithm, a global model is usually stored at a central server and is broadcast periodically to all the participating devices. Then, each device performs local training steps on its own model copy, before the device sends the updated model to the central server. Finally, the central server updates the global model by aggregating the received model copies. This protocol is followed in most FL algorithms, including the well-established FedAvg (McMahan et al., 2017), FedProx (Li et al., 2018), FedNova (Wang et al., 2020) and SCAFFOLD (Karimireddy et al., 2019). The main criticism against synchronous learning could be that it often results in heavy communication/computation overheads and long idle/waiting times for workers.

Asynchrony and its challenges. The deployment of a asynchronous learning method is often convoluted. In the past decade, HogWild! (Niu et al., 2011; Liu et al., 2014) has emerged as a general asynchronous distributed methodology, and has been applied initially in basic ML problems like sparse linear/logistic regression (Zhuang et al., 2013; Yun et al., 2013; Hsieh et al., 2015). Ideally, based on sparsity arguments, each edge device can independently update parts of the global model –that overlap only slightly with the updates of other workers– in a lock-free fashion (Niu et al., 2011; Liu et al., 2014). This way, faster, more powerful edge workers do not suffer from idle waiting due to slower stragglers. Yet, the use of asynchrony has been a topic of dispute in distributed neural network training (Dean et al., 2012; Chen et al., 2016).

Resurgence in asynchrony. Recently, asynchronous methods have regained popularity, mainly due to the interest in applying asynchronous motions within FL on edge devices. Yet, traditional off-the-shelf asynchronous algorithms still have issues, which might be exacerbated in the FL setting. As slower devices take longer local training time, this might result in inconsistent update schedules of the global model, compared to that of faster devices. This might have ramifications: i) For FL on i.i.d. data, this will cause the gradient staleness problem and result in convergence rate decrease; and, ii) on non-i.i.d. data, this will result in a significant drop in global model final accuracy.

As solutions, novel approaches on asynchronous FL propose weighted global aggregation techniques that take into consideration the heterogeneity of the devices (Xie et al., 2019; Chen et al., 2019; Shi et al., 2020); yet, these methods often place a heavy computation/communication burden, as they rely on broadcasting full model updates to all the clients and/or the server. Other works monitor client speed to guide the training assignments (Li et al., 2021; Chai et al., 2020). Finally, recent efforts propose semi-asynchronous methods, where participating devices are selected and buffered in order to complete a semi-synchronous global update periodically (Huba et al., 2022; Wu et al., 2020).

What is different in this work? As most algorithms stem from adapting asynchrony in synchronous FL, *one still needs to broadcast the full model to all devices, following a data parallel distributed protocol (Farber & Asanovic, 1997; Raina et al., 2009), regardless of device heterogeneity*. This inspire us to ask a key question:

> *"Can we select submodels out of the global model and send these instead to each device, taking into account the device heterogeneity?"*

We answer this question affirmatively, by proposing a novel distributed dropout method for FL. We dub our method `AsyncDrop`. Our approach assigns different submodels to each device[1]; empirically, such a strategy decreases the required time to converge to an accuracy level, while preserving favorable final accuracy. Our idea is based on the ideas of HogWild! (Niu et al., 2011; Liu et al., 2014) –in terms of sparse submodels– and Independent Subnetwork Training (IST) (Yuan et al., 2022; Dun et al., 2022; Liao & Kyrillidis, 2021; Wolfe et al., 2021) –where submodels are deliberately created for distribution, in order to decrease both computational and communication requirements.

Overall, the contributions of this work can be summarized as follows:

- We consider and propose *asynchronous distributed dropout* (`AsyncDrop`) for efficient large-scale FL. Our focus is on non-trivial, non-convex ML models –as in neural network training– and our framework provides specific engineering solutions for these cases in practice.
- We theoretically characterize and support our proposal with rigorous and non-trivial convergence rate guarantees. Currently, our theory assumes bounded delays; our future goal is to exploit recent developments that drop such assumptions (Koloskova et al., 2022). Yet, our theory already considers the harder case of neural network training, which is often omitted in existing theory results.
- We provide specific implementation instances and share "best practices" for faster distributed FL in practice. As a side-product, our preliminary results include baseline asynchronous implementations of many synchronous methods (such as FedAvg, FedProx, and more), that are not existent currently, to the best of our knowledge.

[1]We consider both random assignment, as well as structured assignments, based on the computation power of the devices.

2 Problem Setup and Challenges

Optimization in neural network training. We consider FL scenarios over *supervised* neural network training: i.e., we optimize a loss function $\ell(\cdot,\cdot)$ over a dataset, such that the model maps unseen data to their true labels, unless otherwise stated. For clarity, the loss $\ell(\mathbf{W},\cdot)$ encodes both the loss metric and the neural architecture, with parameters $\mathbf{W}$. Formally, given a data distribution $\mathcal{D}$ and $\{\mathbf{x}_i, y_i\} \sim \mathcal{D}$, where $\mathbf{x}_i$ is a data sample, and y_i is its corresponding label, classical deep learning aims in finding $\mathbf{W}^\star$ as in:

$$\mathbf{W}^\star = \underset{\mathbf{W}\in\mathcal{H}}{\operatorname{argmin}} \left\{ \mathcal{L}(\mathbf{W}) := \tfrac{1}{n}\sum_{i=1}^{n} \ell\left(\mathbf{W}, \{\mathbf{x}_i, y_i\}\right) \right\},$$

where $\mathcal{H}$ denotes the model hypothesis class that "molds" the trainable parameters $\mathbf{W}$.

The minimization above can be achieved by using different approaches, but almost all training is accomplished via a variation of *stochastic gradient descent* (SGD) (Robbins & Monro, 1951). SGD modifies the current guess $\mathbf{W}_t$ using stochastic directions $\nabla\ell_{i_t}(\mathbf{W}_i) := \nabla\ell(\mathbf{W}_i, \{\mathbf{x}_{i_t}, y_{i_t}\})$. I.e., $\mathbf{W}_{t+1} \leftarrow \mathbf{W}_t - \eta\nabla\ell_{i_t}(\mathbf{W}_t)$. Here, $\eta > 0$ is the learning rate, and i_t is a single or a mini-batch of examples. Most FL algorithms are based on these basic stochastic motions, like FedAvg (McMahan et al., 2017), FedProx (Li et al., 2018), FedNova (Wang et al., 2020) and SCAFFOLD (Karimireddy et al., 2019).

FL formulation. Let S be the total number of clients in a distributed FL scenario. Each client i has its own local data $\mathcal{D}_i$ such that the whole dataset satisfies $\mathcal{D} = \cup_i \mathcal{D}_i$, and usually $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset, \forall i \neq j$. The goal of FL is to find a global model $\mathbf{W}$ that achieves good accuracy on all data $\mathcal{D}$, by minimizing the following optimization problem:

$$\mathbf{W}^\star = \underset{\mathbf{W}\in\mathcal{H}}{\operatorname{argmin}} \left\{ \mathcal{L}(\mathbf{W}) := \tfrac{1}{S}\sum_{i=1}^{S} \ell\left(\mathbf{W}, \mathcal{D}_i\right) \right\},$$

where $\ell\left(\mathbf{W}, \mathcal{D}_i\right) = \frac{1}{|\mathcal{D}_i|}\sum_{\{\mathbf{x}_j, y_j\}\in\mathcal{D}_i} \ell\left(\mathbf{W}, \{\mathbf{x}_j, y_j\}\right)$. Herein, we consider both i.i.d. and non-i.i.d. cases, since local data distribution $\mathcal{D}_i$ can be heterogeneous and follow a non-i.i.d. distribution.

Details of asynchronous training. An abstract description of how asynchronous FL operates is provided in Algorithm 1. In particular, given a number of server iterations T, each client i gets the updated global model $\mathbf{W}_t$ from the server, and further locally trains it using $\mathcal{D}_i$ for a number of local

Algorithm 1 Meta Asynchronous FL

Parameters: T iters, S clients, l local iters., $\mathbf{W}$ as current global model, $\mathbf{W}_i$ as local model for i-th client, $\alpha \in (0, 1)$, η step size.

∞

$\mathbf{W} \leftarrow$ randomly initialized global model.
`//On each client` i `asynchronously:`
for $t = 0, \ldots, T-1$ **do**
  $\mathbf{W}_{i,t} \leftarrow \mathbf{W}$
  `//Train` $\mathbf{W}_i$ `for` l `iters. via SGD`
  for $j = 1, \ldots, l$ **do**
    $\mathbf{W}_{i,t} \leftarrow \mathbf{W}_{i,t} - \eta \frac{\partial \mathcal{L}}{\mathbf{W}_{i,t}}$
  end for
  `//Write local to global model`
  $\mathbf{W} \leftarrow (1-\alpha) \cdot \mathbf{W} + \alpha \cdot \mathbf{W}_{i,t}$
end for

iterations l. Asynchronous FL assumes each client has different computation power and communication bandwidth; this can be abstracted by different wall-clock times required to finish a number of local training iterations. Thus, when client i has completed its round, the updated model is shared with the server to be aggregated, before the next round of communication and computation starts for client i. *This is different from classical synchronous FL, where the global model is updated only when all participating clients finish (or time-out) certain local training iterations.*

3 Related Work

Challenges. Asynchronous steps often lead to inconsistent update schedules of the global model and are characterized by gradient staleness and drifting.

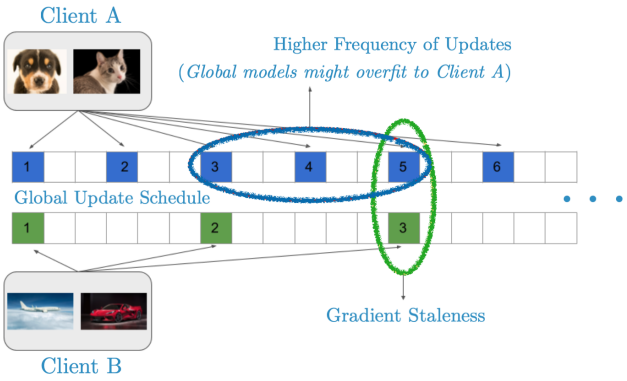

Figure 1. Potential issues in asynchronous FL.

Consider the toy setting in Figure 1. The two clients (Clients A and B) have a significantly different update schedule on the global model: Here, Client A has higher computational power or communication bandwidth –compared to client B– potentially leading to model drifting, lack of fair training and more severe gradient staleness. On top, consider these two clients having different local (non-i.i.d.) data distributions.

Related Work. The issue of model drifting due to data "non-iidness" is a central piece in FL research. Algorithms, such as FedProx (Li et al., 2018), utilize regularization to constrain local parameters "drifting" away.

The gradient staleness problem has been widely studied in asynchronous FL (Xie et al., 2019; Chen et al., 2019; Shi et al., 2020; Li et al., 2021; Chai et al., 2020). In these approaches, the weight of each local client update is proportional to the "capabilities" of the client. This should decrease the negative impact from stale gradients by slower clients. Semi-asynchronous methods have been proposed (Huba et al., 2022; Wu et al., 2020); yet, they require fast clients to wait until all other clients' updates are completed, in order to receive the updated model for the next round.

Finally, numerous quantization (Alistarh et al., 2017; Yu et al., 2019) and sparsification (Aji & Heafield, 2017; Jiang & Agrawal, 2018) techniques have been proposed for reducing computation and communication costs in FL.

4 Asynch. Distributed Dropout

(Distributed) Dropout. Dropout (Wan et al., 2013; Srivastava et al., 2014; Gal & Ghahramani, 2016; Courbariaux et al., 2015) is a widely-accepted regularization technique. The procedure of Dropout is as follows: per training round, a random mask over the parameters is generated; this mask is used to nullify part of the neurons in the neural network for this particular iteration. Variants of dropout include the drop-connect (Wan et al., 2013), multisample dropout (Inoue, 2019), Gaussian dropout (Wang & Manning, 2013), and the variational dropout (Kingma et al., 2015).

The idea of dropout has also been used in efficient distributed and/or FL scenarios. (Horvath et al., 2021a) introduces FjORD and the Ordered Dropout, a *synchronous* distributed dropout technique that leads to ordered, nested representation of knowledge in models, and enables the extraction of lower footprint submodels without the need of retraining. Such submodels are more suitable in client heterogeneity, as they adapt submodel's width to the client's capabilities. See also Nested Dropout (Rippel et al., 2014) and HeteroFL (Diao et al., 2020).

Our proposal and main hypothesis. We focus on the *asynchronous version of distributed dropout.* In Appendix, we study theoretically whether asynchrony provably works in non-trivial non-convex scenarios –*as in training neural networks*– with random masks that generate submodels for each worker. The algorithm is described in Algorithm 2, dubbed as `AsyncDrop`, and is based on recent distributed protocols (Yuan et al., 2022; Dun et al., 2022; Liao & Kyrillidis, 2021; Wolfe et al., 2021); key features are highlighted

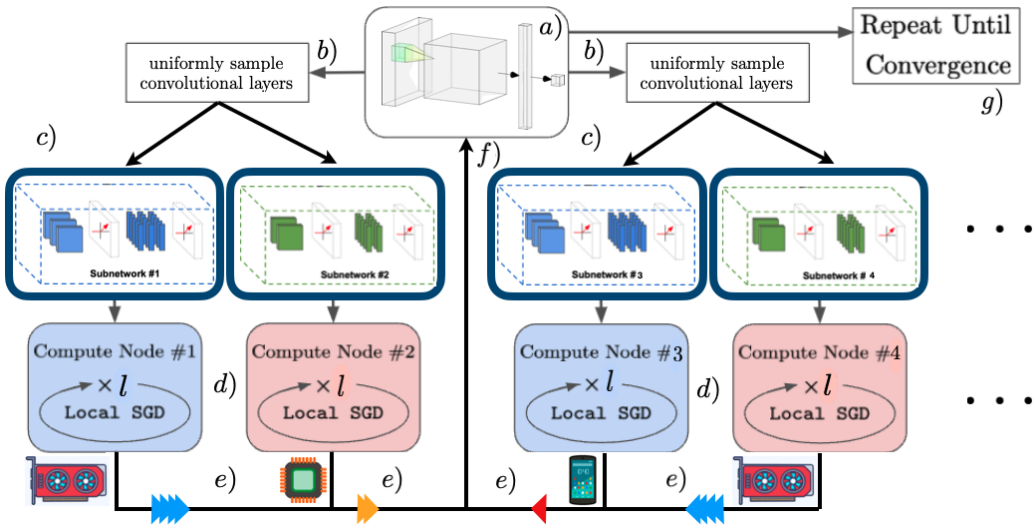

Figure 2. Schematic representation of `AsyncDropout`. $a)$ This is a simple representation of a CNN model. Our algorithm applies for arbitrary depth of CNNs (ResNets) as well as other architectures (MLPs, LSTMs, etc); here we restrict to a shallow CNN for illustration purposes. $b)$ Per request, random sub-sampled CNN models are created that result into different *subnetworks*. $c)$ These submodels are distributed to devices with different computational capabilities (here GPU, CPU, or a smartphone). $d)$ Without loss of generality, we assume that all devices train locally the submodel for l iterations. $e)$ However, each device finishes local training in different timestamps (shown as different colored arrows: red: slow speed; orange: moderate speed; blue: fast speed). $f)$ Yet, the global model is asynchronously updated and new submodels are created without global synchronization. $g)$ The above procedure is repeated till convergence.

Algorithm 2 Asynchronous dropout (`AsyncDrop`)

Parameters: T iters, S clients, l local iters., $\mathbf{W}$ as current global model, $\mathbf{W}_i$ as local model for i-th client, $\alpha \in (0,1)$, η step size.

∞

$\mathbf{W} \leftarrow$ randomly initialized global model.
`//On each client` i `asynchronously:`
for $t = 0, \ldots, T-1$ **do**
  Generate mask $\mathbf{M}_{i,t}$
  $\mathbf{W}_{i,t} \leftarrow \mathbf{W}_t \odot \mathbf{M}_{i,t}$
  `//Train` $\mathbf{W}_{i,t}$ `for` l `iters. via SGD`
  for $j = 1, \ldots, l$ **do**
    $\mathbf{W}_{i,t} \leftarrow \mathbf{W}_{i,t} - \eta \frac{\partial \mathcal{L}}{\mathbf{W}_{i,t}}$
  end for
  `//Write local to global model`
  $\mathbf{W}_{t+1} \leftarrow \mathbf{W}_t \odot (\mathbf{M}_{i,t})^c$
  $+((1-\alpha) \cdot \mathbf{W}_t + \alpha \cdot \mathbf{W}_{i,t}) \odot \mathbf{M}_{i,t}$
end for

in teal-colored text. The main difference from Algorithm 1 is that Algorithm 2 splits the model vertically per iteration, where each submodel contains all layers of the neural network, but only with a (non-overlapping) subset of neurons being active in each layer. Multiple local SGD steps can be performed without the need for the workers to communicate. See also Figure 2 for a schematic representation of asynchronous distributed dropout for training a CNN.

5 Experiments

Setup. We generate simulated FL scenarios with 104 clients/devices of diverse computation and communication capabilities. We implement clients as independent processes, each distributed on different GPUs with access to the same RAM space. We follow HogWild!'s distributed model (Niu et al., 2011): $i)$ we use a shared-memory system to store the global model; $ii)$ each simulated client can update/read the global model in a fully lock free mode; and $iii)$ each client transfers the local model to the assigned GPU for local training. We activate 8 clients at any given moment.

We use 25% dropout rate in (`Hetero`) `AsyncDrop` for CNN and MLP, while we use 12.5% for LSTM. *Even for such low dropout rates, the gains in training are obvious and significant, as we show in the experiments.* In the appendix, we provide ablation studies on how the dropout rate affects the performance of `AsyncDrop`-family of algorithms. We simulate the communication and computation savings by inserting shorter time delay, based on which we estimate the training time and communication cost.

Simulation of heterogeneous computations. We abstract heterogeneous computation and communication capabilities by "forcing" different delays after each training iteration. The delay time is inverse proportional to the intended capacity. In our experiments, we simulated 8 levels of computation and communication capabilities, that are evenly distributed between the slowest client and fastest clients.

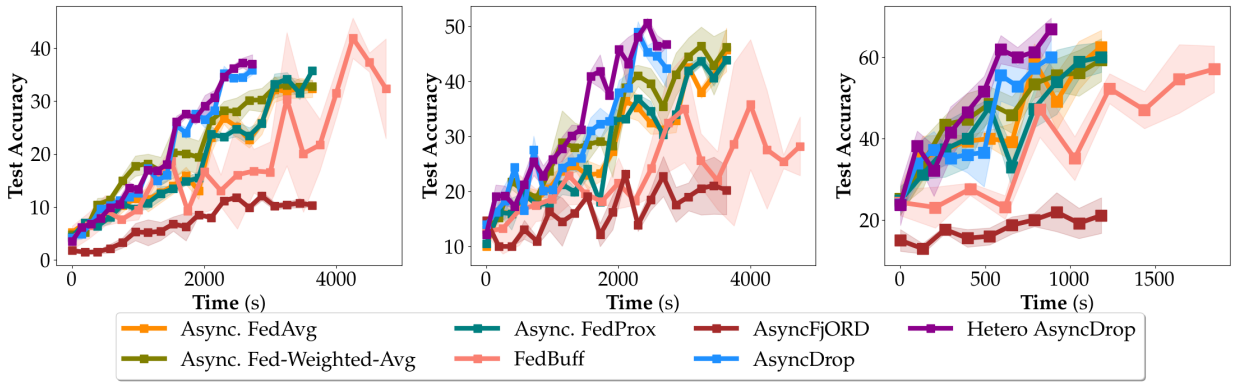

Figure 3. ResNet-based model. ***Left:*** CIFAR100 non-i.i.d.; ***Middle:*** CIFAR10, non-i.i.d.; ***Right:*** FMNIST, non-i.i.d.

The difference between the slowest and the fastest clients is selected to be $\sim 5\times$. We make sure that, at any given moment, clients with diverse capacity are active. Finally, all clients with similar computation power shall have similarly biased local data distribution.[2]

Problem cases. We experiment on diverse neural network architectures and diverse types of learning tasks, including ResNets on Computer Vision datasets (CIFAR10, CIFAR100, FMNIST), MLPs on FMNIST dataset, and LSTMs on sentimental analysis (IMDB).

Baseline methods. For comparison, the baselines we consider are: $i)$ the asynchronous FedAvg is the direct adaptation of FedAvg with asynchronous motions; $ii)$ the asynchronous FedAvg with weighted aggregation represents the general approach of assigning devices different "importance", based on their capabilities (Xie et al., 2019; Chen et al., 2019; Shi et al., 2020); $iii)$ the asynchronous FjORD is our asynchronous adaptation of (Horvath et al., 2021b); $iv)$ the asynchronous FedProx adds an independent proximal loss to the local training loss of each device, in order to control gradient staleness and overfiting (Li et al., 2018); and $v)$ FedBuff is a semi-asynchronous method which uses buffers for stale updates in a synchronized global scheme (Nguyen et al., 2022).

For all baselines and (`Hetero`) `AsyncDrop` on CIFAR10, CIFAR100, FMNIST, we set the local iterations at $l = 50$ while for IMDB, $l = 40$. For FedBuff, we set the buffer size to 4, which is half of the activated clients. We perform 3 trials with different random seeds. We report the maximum test accuracy, as well as the estimated time and communication cost to reach a certain target accuracy: *we select the second lowest test accuracy among all baselines as the target accuracy (third column in result tables).*

CNN-based results. We test (`Hetero`) `AsyncDrop` on ResNet34 and using CIFAR10, CIFAR100 and FMNIST datasets. For the CIFAR10 and CIFAR100 datasets, we train for 320 epochs, while for the FMNIST dataset, we train for 160 epochs. We stop the execution when the fastest client finishes all its epochs. As shown in Table **??**, `Hetero AsyncDrop` shows non-trivial improvements in *final accuracy, training time and communication cost, simultaneously*. `Hetero AsyncDrop` shows lower accuracy compared with FedBuff in the CIFAR100 case; yet, it achieves up to 85% reduction in training time, due to the fact that FedBuff will require faster workers to wait until the buffer is filled to update the global model (this also justifies the up to $\sim 42.22\%$ reduction in total communication cost). Finally, we observe that (`Hetero`) `AsyncDrop` shows quite stable performance; in red color we indicate the variability of results over trials. The similar training time of some baselines to reach target accuracy is caused by epoch-wise testing, using same epoch-wise learning rate schedule and similar convergence rate as shown in Figure 3.

MLP-based and LSTM-based results. We adapt the (`Hetero`) `AsyncDrop` to the MLP model by applying the hidden neuron Dropout, which is similar to channel dropout in CNNs (160 epochs). For the LSTM and the IMDB sentimental analysis dataset (80 epochs), we create non-i.i.d. datasets based on different label distribution in each local training set. As shown in Table 1 in the Appendix, `Hetero AsyncDrop` achieves better performance overall over the MLP model, in terms of accuracy, training time and communication cost. As shown in Table 2 in the Appendix, for the LSTM model, `Hetero AsyncDrop` shows comparable accuracy with respect to other baselines, while achieving reduction in training time and communication cost in most cases. Our conjecture for the lower gain compared with other architectures is that LSTM-based (or even RNN-based) architectures might be difficult to our proposed dropout score mechanism, as the update of each network parameter is the average of several virtual parameters in the unrolled network.

[2]This setting is to avoid the fastest clients with similar computation power cover all the data, which will reduce the problem into a trivial synchronous federated learning problem.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

REFERENCES

Aji, A. F. and Heafield, K. Sparse communication for distributed gradient descent. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 440–445, 2017. [Cited on page 3.]

Alistarh, D., Grubic, D., Li, J., Tomioka, R., and Vojnovic, M. QSGD: Communication-efficient SGD via gradient quantization and encoding. *Advances in Neural Information Processing Systems*, 30, 2017. [Cited on page 3.]

Chai, Z., Chen, Y., Zhao, L., Cheng, Y., and Rangwala, H. Fedat: A communication-efficient federated learning method with asynchronous tiers under non-iid data. *ArXivorg*, 2020. [Cited on pages 2 and 3.]

Chen, J., Pan, X., Monga, R., Bengio, S., and Jozefowicz, R. Revisiting distributed synchronous sgd. *arXiv preprint arXiv:1604.00981*, 2016. [Cited on page 1.]

Chen, Y., Sun, X., and Jin, Y. Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. *IEEE transactions on neural networks and learning systems*, 31(10): 4229–4238, 2019. [Cited on pages 2, 3, and 5.]

Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems*, 28, 2015. [Cited on page 3.]

Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Ranzato, M., Senior, A., Tucker, P., Yang, K., et al. Large scale distributed deep networks. *Advances in neural information processing systems*, 25, 2012. [Cited on page 1.]

Diao, E., Ding, J., and Tarokh, V. HeteroFL: Computation and communication efficient federated learning for heterogeneous clients. *arXiv preprint arXiv:2010.01264*, 2020. [Cited on page 3.]

Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, pp. 1675–1685. PMLR, 2019. [Cited on page 9.]

Dun, C., Wolfe, C. R., Jermaine, C. M., and Kyrillidis, A. ResIST: Layer-wise decomposition of ResNets for distributed training. In *Uncertainty in Artificial Intelligence*, pp. 610–620. PMLR, 2022. [Cited on pages 2 and 3.]

Preliminary work. Under review by the Machine Learning and Systems (MLSys) Conference Workshop on Resource-Constrained Learning in Wireless Networks. Do not distribute.

Farber, P. and Asanovic, K. Parallel neural network training on multi-spert. In *Proceedings of 3rd International Conference on Algorithms and Architectures for Parallel Processing*, pp. 659–666, 1997. doi: 10.1109/ICAPP.1997.651531. [Cited on page 2.]

Gal, Y. and Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016. [Cited on page 3.]

Horvath, S., Laskaridis, S., Almeida, M., Leontiadis, I., Venieris, S., and Lane, N. FjORD: Fair and accurate federated learning under heterogeneous targets with ordered dropout. *Advances in Neural Information Processing Systems*, 34:12876–12889, 2021a. [Cited on page 3.]

Horvath, S., Laskaridis, S., Almeida, M., Leontiadis, I., Venieris, S. I., and Lane, N. D. Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout, 2021b. URL https://arxiv.org/abs/2102.13451. [Cited on page 5.]

Hsieh, C.-J., Yu, H.-F., and Dhillon, I. PaSSCODE: Parallel asynchronous stochastic dual co-ordinate descent. In *International Conference on Machine Learning*, pp. 2370–2379. PMLR, 2015. [Cited on page 1.]

Huba, D., Nguyen, J., Malik, K., Zhu, R., Rabbat, M., Yousefpour, A., Wu, C.-J., Zhan, H., Ustinov, P., Srinivas, H., et al. Papaya: Practical, private, and scalable federated learning. *Proceedings of Machine Learning and Systems*, 4:814–832, 2022. [Cited on pages 1, 2, and 3.]

Inoue, H. Multi-sample dropout for accelerated training and better generalization. *arXiv preprint arXiv:1905.09788*, 2019. [Cited on page 3.]

Jiang, P. and Agrawal, G. A linear speedup analysis of distributed deep learning with sparse and quantized communication. *Advances in Neural Information Processing Systems*, 31, 2018. [Cited on page 3.]

Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S. J., Stich, S. U., and Suresh, A. T. SCAFFOLD: Stochastic controlled averaging for federated learning, 2019. URL https://arxiv.org/abs/1910.06378. [Cited on pages 1 and 2.]

Kingma, D. P., Salimans, T., and Welling, M. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28, 2015. [Cited on page 3.]

Koloskova, A., Stich, S. U., and Jaggi, M. Sharper convergence guarantees for asynchronous SGD for distributed and federated learning. *arXiv preprint arXiv:2206.08307*, 2022. [Cited on page 2.]

Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. Federated optimization in heterogeneous networks, 2018. URL https://arxiv.org/abs/1812.06127. [Cited on pages 1, 2, 3, and 5.]

Li, X., Qu, Z., Tang, B., and Lu, Z. Stragglers are not disaster: A hybrid federated learning algorithm with delayed gradients. *arXiv preprint arXiv:2102.06329*, 2021. [Cited on pages 2 and 3.]

Liao, F. and Kyrillidis, A. On the convergence of shallow neural network training with randomly masked neurons. *arXiv preprint arXiv:2112.02668*, 2021. [Cited on pages 2 and 3.]

Liu, J., Wright, S., Ré, C., Bittorf, V., and Sridhar, S. An asynchronous parallel stochastic coordinate descent algorithm. In *International Conference on Machine Learning*, pp. 469–477. PMLR, 2014. [Cited on pages 1 and 2.]

McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017. [Cited on pages 1 and 2.]

Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., and Poor, H. V. Federated learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 23(3):1622–1658, 2021. [Cited on page 1.]

Nguyen, J., Malik, K., Zhan, H., Yousefpour, A., Rabbat, M., Malek, M., and Huba, D. Federated learning with buffered asynchronous aggregation. In *International Conference on Artificial Intelligence and Statistics*, pp. 3581–3607. PMLR, 2022. [Cited on pages 1 and 5.]

Niu, F., Recht, B., Re, C., and Wright, S. J. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent, 2011. URL https://arxiv.org/abs/1106.5730. [Cited on pages 1, 2, and 4.]

Qu, Y., Dai, H., Zhuang, Y., Chen, J., Dong, C., Wu, F., and Guo, S. Decentralized federated learning for UAV networks: Architecture, challenges, and opportunities. *IEEE Network*, 35(6):156–162, 2021. [Cited on page 1.]

Raina, R., Madhavan, A., and Ng, A. Large-scale deep unsupervised learning using graphics processors. In *ICML*, pp. 873–880. ACM, 2009. [Cited on page 2.]

Rippel, O., Gelbart, M., and Adams, R. Learning ordered representations with nested dropout. In *International Conference on Machine Learning*, pp. 1746–1754. PMLR, 2014. [Cited on page 3.]

Robbins, H. and Monro, S. A stochastic approximation method. *The annals of mathematical statistics*, pp. 400–407, 1951. [Cited on page 2.]

Shi, G., Li, L., Wang, J., Chen, W., Ye, K., and Xu, C. Hysync: Hybrid federated learning with effective synchronization. In *2020 IEEE 22nd International Conference on High Performance Computing and Communications; IEEE 18th International Conference on Smart City; IEEE 6th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, pp. 628–633, 2020. doi: 10.1109/HPCC-SmartCity-DSS50907.2020.00080. [Cited on pages 2, 3, and 5.]

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. [Cited on page 3.]

Stojkovic, B., Woodbridge, J., Fang, Z., Cai, J., Petrov, A., Iyer, S., Huang, D., Yau, P., Kumar, A. S., Jawa, H., et al. Applied federated learning: Architectural design for robust and efficient learning in privacy aware settings. *arXiv preprint arXiv:2206.00807*, 2022. [Cited on page 1.]

Tandon, R., Lei, Q., Dimakis, A. G., and Karampatziakis, N. Gradient coding: Avoiding stragglers in distributed learning. In *International Conference on Machine Learning*, pp. 3368–3376. PMLR, 2017. [Cited on page 1.]

Wan, L., Zeiler, M., Zhang, S., Le Cun, Y., and Fergus, R. Regularization of neural networks using dropconnect. In *International conference on machine learning*, pp. 1058–1066. PMLR, 2013. [Cited on page 3.]

Wang, J., Liu, Q., Liang, H., Joshi, G., and Poor, H. V. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020. [Cited on pages 1 and 2.]

Wang, J., Charles, Z., Xu, Z., Joshi, G., McMahan, H. B., Al-Shedivat, M., Andrew, G., Avestimehr, S., Daly, K., Data, D., et al. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021. [Cited on page 1.]

Wang, S. and Manning, C. Fast dropout training. In *international conference on machine learning*, pp. 118–126. PMLR, 2013. [Cited on page 3.]

Wolfe, C. R., Yang, J., Chowdhury, A., Dun, C., Bayer, A., Segarra, S., and Kyrillidis, A. GIST: Distributed training for large-scale graph convolutional networks. *arXiv preprint arXiv:2102.10424*, 2021. [Cited on pages 2 and 3.]

Wu, W., He, L., Lin, W., Mao, R., Maple, C., and Jarvis, S. Safa: A semi-asynchronous protocol for fast federated learning with low overhead. *IEEE Transactions on Computers*, 70(5):655–668, 2020. [Cited on pages 2 and 3.]

Xie, C., Koyejo, S., and Gupta, I. Asynchronous federated optimization, 2019. URL https://arxiv.org/abs/1903.03934. [Cited on pages 2, 3, and 5.]

Yu, Y., Wu, J., and Huang, L. Double quantization for communication-efficient distributed optimization. *Advances in Neural Information Processing Systems*, 32, 2019. [Cited on page 3.]

Yuan, B., Wolfe, C. R., Dun, C., Tang, Y., Kyrillidis, A., and Jermaine, C. Distributed learning of fully connected neural networks using independent subnet training. *Proceedings of the VLDB Endowment*, 15(8):1581–1590, 2022. [Cited on pages 2 and 3.]

Yun, H., Yu, H.-F., Hsieh, C.-J., Vishwanathan, S., and Dhillon, I. Nomad: Non-locking, stochastic multi-machine algorithm for asynchronous and decentralized matrix completion. *arXiv preprint arXiv:1312.0193*, 2013. [Cited on page 1.]

Zhuang, Y., Chin, W.-S., Juan, Y.-C., and Lin, C.-J. A fast parallel SGD for matrix factorization in shared memory systems. In *Proceedings of the 7th ACM conference on Recommender systems*, pp. 249–256, 2013. [Cited on page 1.]

Theoretical Results. We are interested in understanding whether such a combination of asynchronous computing and dropout techniques lead to convergence and favorable results: *given the variance introduced by both asynchronous updates and training of submodels, it is not obvious whether –and under which conditions– such a protocol would work.*

For ease of presentation and clarity of results, we analyse a one-hidden-layer CNN, and show convergence with random filter dropout. Consider a training dataset $(\mathbf{X}, \mathbf{y}) = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where each $\mathbf{x}_i \in \mathbb{R}^{\hat{d} \times p}$ is an image and y_i being its label. Here, $\hat{d}$ is the number of input channels and p the number of pixels. Let q denote the size of the filter, and let m be the number of filters in the first layer. Based on previous work (Du et al., 2019), we let $\hat{\phi}(\cdot)$ denote the patching operator with $\hat{\phi}(x) \in \mathbb{R}^{q\hat{d} \times p}$. Consider the first layer weight $\mathbf{W} \in \mathbb{R}^{m \times q\hat{d}}$, and second layer (aggregation) weight $\mathbf{a} \in \mathbb{R}^{m \times p}$. We assume that only the first layer weights $\mathbf{W}$ is trainable. The CNN trained on the means squared error has the form:

$$f(\mathbf{x}, \mathbf{W}) = \left\langle \mathbf{a}, \sigma\left(\mathbf{W}\hat{\phi}(\mathbf{x})\right), ; \right\rangle, \mathcal{L}(\mathbf{W}) = \|f(\mathbf{X}, \mathbf{W}) - \mathbf{y}\|_2^2,$$

where $f(\mathbf{x}, \cdot)$ denotes the output of the one-layer CNN for input $\mathbf{x}$, and $\mathcal{L}(\cdot)$ is the loss function. We use the ℓ_2-norm loss for simplicity. We make the following assumption on the training data and the CNN weight initialization.

Assumption .1 *(Training Data) Assume that for all $i \in [n]$, we have $\|\mathbf{x}_i\|_F = q^{-\frac{1}{2}}$ and $|y_i| \leq C$ for some constant C. Moreover, for all $i, i' \in [n]$ we have $\mathbf{x}_i \neq \mathbf{x}_{i'}$.*

Note that this can be satisfied by normalizing the data. For simplicity of the analysis, let $d := q\hat{d}$.

Assumption .2 *(Initialization) $\mathbf{w}_{0,i} \sim \mathcal{N}\left(0, \kappa^2 \mathbf{I}\right)$ and $a_{i,i'} \sim \left\{\pm \frac{1}{p\sqrt{m}}\right\}$ for $i \in [m]$ and $i' \in [p]$.*

In an asynchronous scenario, the neural network weight is updated with stale gradients due to the asynchronous updates, where δ_t is the delay at training step t. *We assume δ_t is bounded by a constant E.* Then, a simple version of gradient descent under these assumptions looks like:

$$\mathbf{W}_t = \mathbf{W}_t - \eta \nabla_{\mathbf{W}} \mathcal{L}\left(\mathbf{W}_{t-\delta_t}\right), \quad \delta_t \leq E,$$

where $\mathbf{W}_{t-\delta_t}$ indicates that the gradient is evaluated on a earlier version of the model parameters. Given the above, we provide the following guarantees:

Theorem .1 *Let $f(\cdot, \cdot)$ be a one-hidden-layer CNN with the second layer weight fixed. Let $\mathbf{u}_t$ abstractly represent the output of the model after t iterations, over the random selection of the masks. Let E denotes the maximum gradient delay/staleness. Let ξ denote the dropout rate ($\xi = 1$ dictates that all neurons are selected), and denote $\theta = 1 - (1-\xi)^S$ the probability that a neuron is active in at least one subnetwork. Assume the number of hidden neurons satisfies $m = \Omega\left(\max\{\frac{n^4 K^2}{\lambda_0^4 \delta^2} \max\{n, d\}, \frac{n}{\lambda_0}\}\right)$ and the step size satisfies $\eta = O\left(\frac{\lambda_0}{n^2}\right)$. Let κ be a proper initialization scaling factor, and it is considered constant. We use λ_0 to denote the smallest eigenvalue of the Neural Tangent Kernel matrix. Let Assumptions 1 and 2 be satisfied. Then, the following convergence rate guarantee is proved to be satisfied:*

$$\begin{aligned}\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{u}_{t+1} - \mathbf{y}\|_2^2\right] \leq & \left(1 - \frac{\theta\eta\lambda_0}{4}\right)^t \|\mathbf{u}_0 - \mathbf{y}\|_2^2 \\ & + O\left(\frac{\theta\eta\lambda_0^3\xi^2\kappa^2E^2}{n^2} + \frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0} + \frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}\right. \\ & \left. + \frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0} + \frac{n\kappa^2\left(\theta - \xi^2\right)}{S}\right)\end{aligned}$$

Remark #1. This theorem states that the sum of the expected weight differences in the t-th iteration (i.e., $\mathbb{E}_{\mathbf{M}_t}[\|\mathbf{u}_{t+1} - \mathbf{y}\|_2^2]$) converges linearly to zero, as dictated by the red term $-\left(1 - \frac{\theta\eta\lambda_0}{4}\right)^t \|\mathbf{u}_0 - \mathbf{y}\|_2^2$– up to an error neighborhood, denoted with the Big-Oh notation term on the right hand side of the expression. Focusing on the latter, there are two types of additive errors: $i)$ the orange-colored terms origin from the dropout analysis: the term $1 - \xi$ is often called as "dropout rate" (when $\xi = 0$, no neurons are selected and the loss hardly decreases, while when $\xi = 1$, all neurons are selected, and the orange-colored terms disappear). $ii)$ the violet-colored terms origin from the asynchronous analysis: when $E = 0$ (i.e., we boil down to synchronous computations), these terms also disappear).

Remark #2. Beyond the above extreme cases, we observe that the error region terms can be controlled by algorithmic and model-design choices: e.g., when the size of the dataset n increases, the first term $\frac{\theta\eta\lambda_0^3\xi^2\kappa^2E^2}{n^2}$ can be controlled; for

sufficiently wide neural network, the terms with m in the denominator can be made arbitrarily small; finally, notice that increasing the number of subnetworks S will drive the last term in the bound zero.

Supplementary material: Efficient and Light-Weight Federated Learning via Asynchronous Distributed Dropout

1 Experimental Results using `AsyncDrop` for the MLP and LSTM architectures

We adapt the (`Hetero`) `AsyncDrop` to the MLP and LSTM model by applying the hidden neuron Dropout, which is similar to channel dropout in CNNs. For the LSTM and the IMDB sentimental analysis dataset, we create non-i.i.d. datasets based on different label distribution in each local training set. For the FMNIST experiments, we train for 160 epochs, while for the IMDB experiments, we train for 80 epochs. As shown in Table 1, `Hetero AsyncDrop` achieves better performance overall over the MLP model, in terms of accuracy, training time and communication cost. As shown in Table 2, for the LSTM model, `Hetero AsyncDrop` shows comparable accuracy with respect to other baselines, while achieving reduction in training time and communication cost in most cases. Our conjecture for the lower gain compared with other architectures is that LSTM-based (or even RNN-based) architectures might be difficult to our proposed dropout score mechanism, as the update of each network parameter is the average of several virtual parameters in the unrolled network.

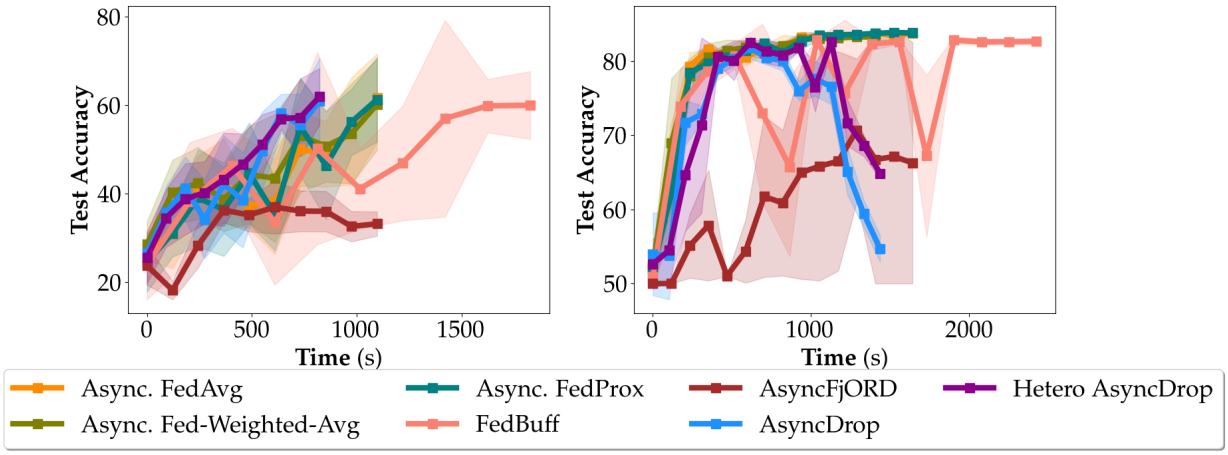

Figure 1: ***Left:*** MLP FMNIST non-i.i.d.; ***Right:*** LSTM IMDB, non-i.i.d.

Table 1: Test accuracy of asynchronous FL baselines vs. (`Hetero`) `AsyncDrop` using a MLP architecture on non-i.i.d MNIST data over > 100 clients. We report the time and communication overhead to reach a certain target accuracy: "Time for `XX`% Accuracy" denotes the second lowest test accuracy among all baselines as the target accuracy.

	Max. Test Accuracy	Time for 59.95% Accuracy	Time Overhead	Comm. Overhead
Async. FedAvg	$61.59 \pm$ **10.02**	1096.1s	+33.33%	+28.73%
Async. Fed-Weighted-Avg	$60.14 \pm$ **10.28**	1096.1s	+33.33%	+28.73%
Async. FedProx	$61.25 \pm$ **9.54**	1096.1s	+33.33%	+28.73%
Async. FjORD	36.92 ± 6.00	N/A	N/A	N/A
FedBuff	$59.96 \pm$ **7.64**	1827.3s	+166.50%	+83.90%
`Hetero AsyncDrop`	61.90 ± 6.39	**822.5s**	**[Best]**	**[Best]**
`AsyncDrop`	$60.85 \pm$ **9.78**	**822.5s**	**[Best]**	**[Best]**

Table 2: Test accuracy of asynchronous FL baselines vs. (`Hetero`) `AsyncDrop` using a LSTM architecture on non-i.i.d IMDB data over > 100 clients. We report the time and communication overhead to reach a certain target accuracy: "Time for `XX`% Accuracy" denotes the second lowest test accuracy among all baselines as the target accuracy.

	Max. Test Accuracy	Time for 82% Accuracy	Time Overhead	Comm. Overhead
Async. FedAvg	83.73 ± 0.18	1875.6s	+52.38%	+32.06%
Async. Fed-Weighted-Avg	83.26 ± 0.33	1640.6s	+33.33%	+15.56%
Async. FedProx	83.83 ± 0.20	1406.3	+14.28%	**-0.95%**
Async. FjORD	70.68 ± 9.62	N/A	N/A	N/A
FedBuff	82.81 ± 0.30	1038.1s	+68.71%	+14.28%
`Hetero AsyncDrop`	82.60 ± 0.57	**1230.4s**	**[Best]**	-
`AsyncDrop`	81.70 ± 1.45	N/A	N/A	N/A

2 Ablation Study

Dropout Ratio for `Hetero AsyncDrop`. In this section we analyze how dropout ratio affects the final accuracy and training time. As shown in Table 3, low dropout rates result in longer training time and lower final accuracy. Lower final accuracy is mainly due to the fact that `AsyncDrop` with dropout work as a regularization technique that works favorably towards better final accuracy. On the other hand, higher dropout rates cause relative reduction in training time, but significant reduction in the final accuracy. This shows dropping too many parameters causes the training of local model unstable, reducing the final accuracy of the global model.

Table 3: Test accuracy of `Hetero AsyncDrop` with different dropout rate on non-i.i.d. CIFAR100 datasets.

Dropout Ratio	0%	25%	50%
Test Accuracy	32.47 ± 1.89	37.26 ± 0.93	23.95 ± 1.07
Time for Accuracy to 23%	2105.4s	1579.1s	1531.2s

Local Iteration Length for `Hetero AsyncDrop`. In this section we analyze how the number of local iterations (l) affect the final accuracy and training time. Ideally, longer local training iteration will decrease the communication frequency and, thus, the communication and time cost. [5, 2] However, as shown in Figure 2, increasing local training iterations cause significant decrease in final accuracy, which might be due to i) increased gradient staleness from slower clients; ii) more severe model drifting. In future research, we will study how to improve `Hetero AsyncDrop` in order to be more robust for higher l values.

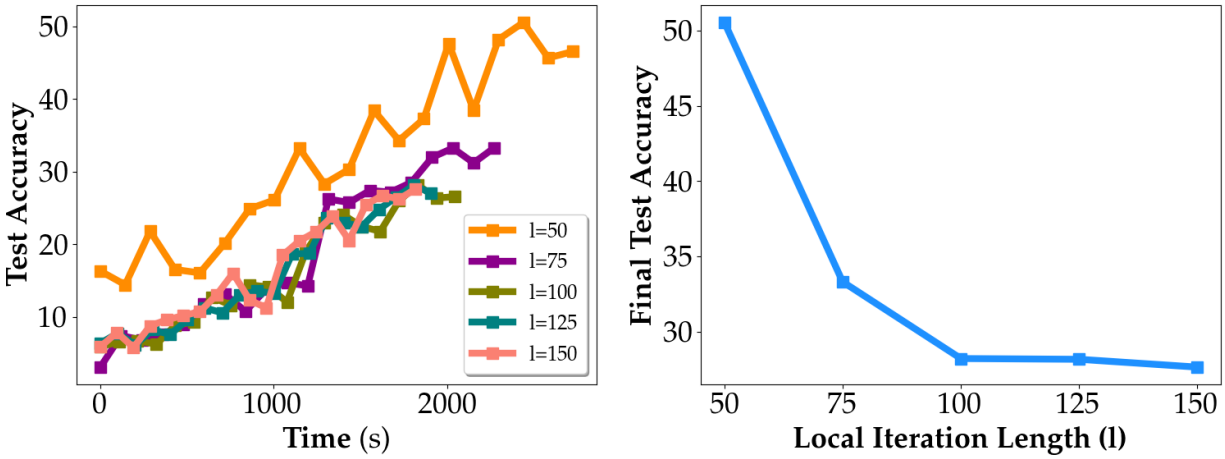

Figure 2: Test Accuracy of `Hetero AsyncDrop` with different local iteration length on CIFAR100, non i.i.d

The effectiveness of Smart Partition in `Hetero AsyncDrop`. As shown in Table 4, on all datasets and all network architectures, `Hetero AsyncDrop` provides improvements on the final accuracy, while decreasing the variance. This indicates that `Hetero AsyncDrop` could be a more stable method that balances heterogeneous clients and data. By

interpreting current experiments, `Hetero AsyncDrop` gives the most gain on ResNet (or CNN model family), which utilizes a channel dropout method.

On the other hand, in both MLP and LSTM, we dropout output hidden neurons and the network parameters accordingly. As filter channels in ResNet are considered to be more semantically disentangled from each other, compared to hidden neurons in fully connected layers, `Hetero AsyncDrop`'s score for filters in ResNet capture better which filters are more frequently/significantly updated (and thus biased towards the faster clients). Thus, our scoring function $q(\cdot)$ targets which parameters to drop in ResNet architectures. We leave to future research on how to improve our proposed mechanism on MLP and LSTM architectures, potentially by using a more structured dropout strategy.

Table 4: Test accuracy of `Hetero AsyncDrop` versus `AsyncDrop`

	`Hetero AsyncDrop`	`AsyncDrop`	Relative Accuracy Gain
ResNet+CIFAR10	50.67 ± 1.75	48.98 ± 3.87	+3.45%
ResNet+CIFAR100	37.26 ± 0.93	35.93 ± 1.07	+2.98%
ResNet+FMNIST	66.89 ± 5.36	60.02 ± 10.38	+11.45%
MLP+FMNIST	61.90 ± 6.39	60.85 ± 9.78	+1.72%
LSTM+IMDB	82.60 ± 0.57	81.70 ± 1.45	+1.10%

Layerwise Smart Partition for `AsyncDrop` on ResNet. For ResNet architectures, another commonly used dropout strategy is that of layerwise dropout [3, 2]. We adapt `Hetero AsynDrop` using layerwise Smart Parition and dropout. The main difference from original `Hetero AsyncDrop` algorithm is that we consider each layer as a single unit for dropout. The Smart Partition score function calculates the total norm change of all parameters in each layer and ranks layers accordingly. Finally, for the i-th fastest worker, we drop layers with the i-th largest q score. In experiments, we set the dropout rate to be 25% to keep the submodel size similar to the original `AsyncDrop` experimebts. As shown in Figure 3, Layerwise `Hetero AsyncDrop` results in a non-negligible decease in performance, with respect to the final accuracy. This might show layer dropping causes increased variance and instability, especially in asynchronous training.

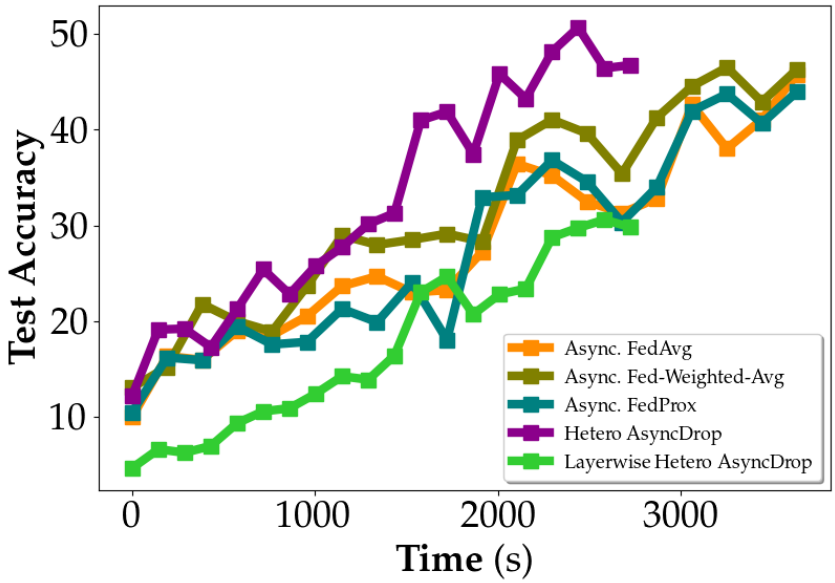

Figure 3: Test Accuracy of Layerwise `Hetero AsyncDrop` versus all baselines on CIFAR100 non- i.i.d.

3 Detailed Mathematical Formulation of `AsyncDrop`

For a vector $\mathbf{v}$, $\|\mathbf{v}\|_2$ denotes its Euclidean (ℓ_2) norm. For a matrix $\mathbf{V}$, $\|\mathbf{V}\|_F$ denotes its Frobenius norm. We use $\mathcal{P}(\cdot)$ to denote the probability of an event, and $\mathbb{I}\{\cdot\}$ to denote the indicator function. For two vectors $\mathbf{v}_1, \mathbf{v}_2$, we use the simplified notation $\mathbb{I}\{\mathbf{v}_1; \mathbf{v}_2\} := \mathbb{I}\{\langle \mathbf{v}_1, \mathbf{v}_2 \rangle \geq 0\}$. Given that the mask in iteration t is $\mathbf{M}_t$, we denote $\mathbb{E}_{[\mathbf{M}_t]}[\cdot] = \mathbb{E}_{\mathbf{M}_0,\ldots,\mathbf{M}_t}[\cdot]$.

Recall that the CNN considered in this paper has the form:

$$f(\mathbf{x}, \theta) = \left\langle \mathbf{a}, \sigma\left(\mathbf{W}_1 \hat{\phi}(\mathbf{x})\right)\right\rangle.$$

Denote $\hat{\mathbf{x}} = \hat{\phi}(\mathbf{x})$. Essentially, this patching operator applies to each channel, with the effect of extending each pixel to a set of pixels around it. So we denote $\hat{\mathbf{x}}_i^{(j)} \in \mathbb{R}^{q\hat{d}}$ as the extended jth pixel across all channels in the ith sample. For each transformed sample, we have that $\|\hat{\mathbf{x}}_i\|_F \leq \sqrt{q}\,\|\mathbf{x}_i\|_F$. We simplify the CNN output as¨

$$f(\hat{\mathbf{x}}, \mathbf{W}) = \langle \mathbf{a}, \sigma(\mathbf{W} \otimes \hat{\mathbf{x}})\rangle = \sum_{r=1}^{m_1}\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}^{(j)}, \mathbf{w}_r\right\rangle\right).$$

In this way, the formulation of CNN reduces to MLP despite a different form of input data $\hat{\mathbf{x}}$ and an additional dimension of aggregation in the second layer. We consider training the neural network f on the mean squared error (MSE):

$$\mathcal{L}(\mathbf{W}) = \left\|f\left(\hat{\mathbf{X}}, \mathbf{W}\right) - \mathbf{y}\right\|_2^2 = \sum_{i=1}^{n}\left(f(\hat{\mathbf{x}}_i, \mathbf{W}) - y_i\right)^2.$$

Similar to HogWild! training is fully lock free and thus cannot prevent several clients from updating the global model at same time, in order to be closer to real-world case, in theoretical analysis, we extend `AsyncDrop` by allowing $S \geq 1$ clients –with same computation speed and thus same global model update schedule– to update the global model simultaneously by averaging their updates. This can be reduced back to original `AsyncDrop`, by simply setting $S = 1$ without changing any part of the analysis.

We assume that we can group all clients by their computation/communication power, such that, in each group, there are exactly $S \geq 1$ clients with the same power; thus, with the same global model update schedule. Each worker will be independently assigned to different dropout masks to be applied on the global model. All updates from local clients in the same group will be averaged.

In math, the subnetwork on client s by filter-wise partition is given by:

$$f_{\mathbf{m}^{(s)}}(\hat{\mathbf{x}}, \mathbf{W}) = \sum_{r=1}^{m_1}\sum_{j=1}^{p} m_r^{(s)} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}^{(j)}, \mathbf{w}_r\right\rangle\right).$$

Trained on the regression loss, the surrogate gradient on current network parameters is given by:

$$\nabla_{\mathbf{w}_r}\mathcal{L}_{\mathbf{m}^{(s)}}(\mathbf{W}) = m_r^{(s)}\sum_{i=1}^{n}\sum_{j=1}^{p}\left(f_{\mathbf{m}^{(s)}}(\hat{\mathbf{x}}_i, \mathbf{W}) - y_i\right) a_{rj}\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_r\right\}.$$

We correspondingly scale the whole network function:

$$f(\hat{\mathbf{x}}, \mathbf{W}) = \xi\sum_{r=1}^{m}\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_r\right\rangle\right).$$

Assuming it is training on the MSE, we write out its gradient as:

$$\nabla_{\mathbf{w}_r}\mathcal{L}(\mathbf{W}) = \xi\sum_{i=1}^{n}\sum_{j=1}^{n}\left(f(\hat{\mathbf{x}}_i, \mathbf{W}) - y_i\right) a_{rj}\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_r\right\}.$$

In this work, we consider the ASYNDROP with one step of local iteration training and $\delta_t \leq E$ gradient delay caused by asynchronous training, given by:

$$\mathbf{w}_{r,t+1} = \mathbf{w}_{r,t} - \eta\frac{N_{r,t-\delta_t}^{\perp}}{N_{r,t-\delta_t}}\sum_{s=1}^{S}\nabla_{\mathbf{w}_r}\mathcal{L}_{\mathbf{m}_{t-\delta_t}^{(s)}}\left(\mathbf{W}_{r,t-\delta_t}\right).$$

Here, let $N_{r,t} = \max\left\{\sum_{s=1}^{S} m_{r,t}^{(s)}, 1\right\}$, and $N_{r,t}^{\perp} = \min\left\{\sum_{s=1}^{S} m_{r,t}^{(s)}, 1\right\}$. Intuitively, $N_{r,t}$ denote the "normalizer" that we will divide the sum of the gradients from all subnetworks with, and $N_{r,t}^{\perp}$ denote the indicator of whether filter r is trained

in at least one subnetwork. Let $\theta = \mathcal{P}\left(N_{r,t}^{\perp} = 1\right) = 1 - (1-\xi)^p$, denoting the probability that at least one of $\left\{m_{r,t}^{(s)}\right\}_{s=1}^{S}$ is one. Denote $u_t^{(i)} = f\left(\hat{\mathbf{x}}_i, \mathbf{W}_t\right)$. For further convenience of our analysis, we define

$$\tilde{u}_{r,t}^{(i)} = \frac{N_{r,t}^{\perp}}{N_{r,t}} \sum_{s=1}^{S} m_{r,t}^{(s)} \hat{u}_t^{(s,i)}; \quad \mathbf{g}_{r,t} = \frac{N_{r,t}^{\perp}}{N_{r,t}} \sum_{s=1}^{S} \nabla_{\mathbf{w}_r} \mathcal{L}_{\mathbf{m}_t^{(s)}}\left(\mathbf{W}_t\right).$$

Then the `AsyncDrop` training has the form

$$\mathbf{w}_{r,t+1} = \mathbf{w}_{r,t} - \eta \mathbf{g}_{r,t-\delta_t}; \quad \mathbf{g}_{r,t} = \sum_{i=1}^{n} \sum_{j=1}^{p} a_{rj} \left(\tilde{u}_{r,t}^{(i)} - N_{r,t}^{\perp} y_i\right) \hat{\mathbf{x}}_i^{(j)} \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}.$$

Suppose that Assumptions 1 and 2 in main text hold. Then for all $i \in [n]$ we have $\|\mathbf{x}_i\|_F = q^{\frac{1}{2}}$, and for all $i, i' \in [n]$ such that $i \neq i'$, we have $\mathbf{x}_i \neq \mathbf{x}_{i'}$. As in previous work [1], we have $\|\hat{\mathbf{x}}_i\|_F =\leq \sqrt{q}\|\mathbf{x}_i\|_F \leq 1$. Thus, for all $j \in [p]$ we have $\left\|\hat{\mathbf{x}}_i^{(j)}\right\| \leq 1$. Moreover, since $\mathbf{x}_i \neq \mathbf{x}_{i'}$ for $i \neq i'$, we then have $\hat{\mathbf{x}}_i \neq \hat{\mathbf{x}}_{i'}$, which implies that $\hat{\mathbf{x}}_i^{(j)} \neq \hat{\mathbf{x}}_{i'}^{(j)}$ for all $i \neq i'$ and $j \in [p]$.

4 Proof of `AsyncDrop` Convergence

In this section, our goal is to prove the convergence of extended `AsyncDrop`. We first state the extended version here, and proceed proving it.

Theorem 4.1 *Let $f(\cdot,\cdot)$ be a one-hidden-layer CNN with the second layer weight fixed. Let $\mathbf{u}_t$ abstractly represent the output of the model after t iterations, over the random selection of the masks. Let E denotes the maximum gradient delay/staleness. Let ξ denote the dropout rate ($\xi = 1$ dictates that all neurons are selected), and denote $\theta = 1 - (1-\xi)^S$ the probability that a neuron is active in at least one subnetwork. Assume the number of hidden neurons satisfies $m = \Omega\left(\max\{\frac{n^4K^2}{\lambda_0^4\delta^2}\max\{n,d\}, \frac{n}{\lambda_0}\}\right)$ and the step size satisfies $\eta = O\left(\frac{\lambda_0}{n^2}\right)$. Let κ be a proper initialization scaling factor, and it is considered constant. We use λ_0 to denote the smallest eigenvalue of the Neural Tangent Kernel matrix. Let Assumptions 1 and 2 be satisfied. Then, the following convergence rate guarantee is proved to be satisfied:*

$$\begin{aligned}\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{u}_{t+1} - \mathbf{y}\|_2^2\right] \leq & \left(1 - \frac{\theta\eta\lambda_0}{4}\right)^t \|\mathbf{u}_0 - \mathbf{y}\|_2^2 \\ & + O\left(\frac{\theta\eta\lambda_0^3\xi^2\kappa^2E^2}{n^2} + \frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0} + \frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}\right. \\ & \left. + \frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0} + \frac{n\kappa^2\left(\theta - \xi^2\right)}{S}\right).\end{aligned}$$

We care about the MSE computed on the scaled full network:

$$u_k^{(i)} = \xi \sum_{r=1}^{m} \sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t}\right\rangle\right); \quad \mathcal{L}\left(\mathbf{W}_t\right) = \|\mathbf{u}_t - \mathbf{y}\|_2^2.$$

Performing gradient descent on this scaled full network involves computing:

$$\nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right) = \xi \sum_{i=1}^{n}\sum_{j=1}^{p}\left(u_t^{(i)} - y_i\right) a_{rj}\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}.$$

We will prove Theorem 4.1 by induction. We assume therom 4.1 is true for all $t' < t$ and $\|\mathbf{w}_{r,t'} - \mathbf{w}_{r,0}\|_2 \leq R := O\left(\frac{\kappa\lambda_0}{n}\right)$.

4.1 Change of Activation Pattern

Let R be some fixed scale. For convenience, we denote:

$$A_{ir}^{(j)} = \left\{\exists \mathbf{w} : \|\mathbf{w} - \mathbf{w}_{r,0}\|_2 \leq R; \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}\right\} \neq \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,0}\right\}\right\}.$$

Note that $A_{ir}^{(j)}$ happens if and only if $\left|\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,0} \right\rangle\right| < R$. Therefore $\mathbb{P}\left(A_{ir}^{(j)}\right) < \frac{2R}{\kappa\sqrt{2\pi}}$. Denote:

$$P_{ij} = \left\{ r \in [m] : \neg A_{ir}^{(j)} \right\}; \quad P_{ij}^{\perp} = [m] \setminus P_{ij}.$$

The next lemma shows the magnitude of $P_{ij}^{\perp}$ is upper bounded by a controlled quantity.

Lemma 4.1 *Let $m = \Omega\left(R^{-1} \log \frac{np}{\delta}\right)$. Then with probability at least $1 - O(\delta)$ it holds for all $i \in [n]$ and $j \in [p]$ that*

$$\left|P_{ij}^{\perp}\right| \leq 3m\kappa^{-1}R.$$

Proof: The magnitude of $P_{ij}^{\perp}$ satisfies:

$$\left|P_{ij}^{\perp}\right| = \sum_{r=1}^{m} \mathbb{I}\left\{A_{ir}^{(j)}\right\}.$$

The indicator function $\mathbb{I}\left\{A_{ir}^{(j)}\right\}$ has bounded first and second moment:

$$\mathbb{E}_{\mathbf{W}}\left[\mathbb{I}\left\{A_{ir}^{(j)}\right\}\right] = \mathbb{P}\left(A_{ir}^{(j)}\right) \leq \frac{2R}{\kappa\sqrt{2\pi}},$$
$$\mathbb{E}_{\mathbf{W}}\left[\left(\mathbb{I}\left\{A_{ir}^{(j)} - \mathbb{E}_{\mathbf{W}}\left[\mathbb{I}\left\{A_{ir}^{(j)}\right\}\right]\right\}\right)^2\right] \leq \mathbb{E}_{\mathbf{W}}\left[\mathbb{I}\left\{A_{ir}^{(j)}\right\}^2\right] \leq \frac{2R}{\kappa\sqrt{2\pi}}.$$

This allows us to apply the Berstein Inequality to get that:

$$\mathbb{P}\left(\sum_{r=1}^{m} \mathbb{I}\left\{A_{ir}^{(j)}\right\} > \frac{2mR}{\kappa\sqrt{2\pi}} + mt\right) < \exp\left(-\frac{m\kappa t^2\sqrt{2\pi}}{8\left(1+\frac{t}{3}\right)R}\right).$$

Therefore, with probability at least $1 - np\exp\left(-m\kappa^{-1}R\right)$ it holds for all $i \in [n]$ and $j \in [p]$ that:

$$\left|P_{ij}^{\perp}\right| = \sum_{r=1}^{m} \mathbb{I}\left\{A_{ir}^{(j)}\right\} \leq 3m\kappa^{-1}R.$$

Letting $m = \Omega\left(R^{-1}\log\frac{np}{\delta}\right)$ gives that the success probability is at least $1 - O(\delta)$.

4.2 Initialization Scale

Let $\mathbf{w}_{0,r} \sim \mathcal{N}\left(0, \kappa^2\mathbf{I}\right)$ and $a_j \sim \left\{-\frac{1}{p\sqrt{m}}, \frac{1}{p\sqrt{m}}\right\}$ for all $r \in [m]$ and $j \in [p]$. The following lemmas hold true:

Lemma 4.2 *Suppose $\kappa \leq 1, R \leq \kappa\sqrt{\frac{d}{32}}$. With probability at least $1 - e^{md/32}$ we have that:*

$$\|W_0\|_F \leq \kappa\sqrt{2md} - \sqrt{m}R.$$

Lemma 4.3 *Assume $\kappa \leq 1$ and $R \leq \frac{\kappa}{\sqrt{2}}$. With probability at least $1 - ne^{-\frac{m}{32}}$ over initialization, it holds for all $i \in [n]$ that:*

$$\sum_{r=1}^{m} \left\langle \mathbf{w}_{0,r}, \mathbf{x}_i \right\rangle^2 \leq 2m\kappa^2 - mR^2$$
$$\sum_{i=1}^{n}\sum_{r=1}^{m} \left\langle \mathbf{w}_{0,r}, \mathbf{x}_i \right\rangle^2 \leq 2mn\kappa^2 - mnR^2.$$

Moreover, we can bound the initial MSE, based on the lemma below:

Lemma 4.4 *Assume that for all $i \in [n]$, y_i satisfies $|y_i| \leq C$ for some $C > 0$. Then, we have*

$$\mathbb{E}_{\mathbf{W}_0,\hat{\mathbf{a}}}\left[\|\mathbf{y}-\mathbf{u}_0\|_2^2\right] \leq \left(p^{-1}+C^2\right)n.$$

Proof: It is obvious that $\mathbb{E}_{\mathbf{W}_0,\hat{\mathbf{a}}}\left[u_0^{(i)}\right] = 0$ for all $i \in [n]$. Moreover:

$$\begin{aligned}
\mathbb{E}_{\mathbf{W}_0,\hat{\mathbf{a}}}\left[u_0^{(i)2}\right] &= \sum_{r,r'=1}^{m}\sum_{j,j'=1}^{p}\mathbb{E}_{\hat{\mathbf{a}}}\left[a_{rj}a_{r'j'}\right]\mathbb{E}_{\mathbf{W}_0}\left[\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{0,r}\right\rangle\right)\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j')},\mathbf{w}_{0,r}\right\rangle\right)\right] \\
&= \frac{1}{p^2m}\sum_{r=1}^{m}\sum_{j=1}^{p}\mathbb{E}_{\mathbf{W}_0}\left[\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{0,r}\right\rangle\right)^2\right] \\
&\leq \frac{1}{p^2m}\sum_{r=1}^{m}\sum_{j=1}^{p}\mathbb{E}_{\mathbf{W}_0}\left[\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{0,r}\right\rangle\right] \\
&\leq p^{-1}
\end{aligned}$$

Therefore:

$$\mathbb{E}_{\mathbf{W}_0,\hat{\mathbf{a}}}\left[\|\mathbf{u}_0-\mathbf{y}\|_2^2\right] = \sum_{i=1}^{n}\mathbb{E}_{\mathbf{W}_0,\hat{\mathbf{a}}}\left[\left(u_0^{(i)}-y_i\right)_2^2\right] = \sum_{i=1}^{n}\left(\mathbb{E}_{\mathbf{W}_0,\hat{\mathbf{a}}}\left[u_0^{(i)2}\right]+y_i^2\right) \leq \left(p^{-1}+C^2\right)n$$

4.3 Kernel Analysis

The neural tangent kernel is defined to be the inner product of the gradient with respect to the neural network output. We let the finite-width NTK be defined as:

$$\mathbf{H}(t)_{ii'} = \sum_{r=1}^{m}\sum_{j,j'=1}^{p}a_{rj}a_{rj'}\left\langle\hat{\mathbf{x}}_i^{(j)},\hat{\mathbf{x}}_{i'}^{(j')}\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j')};\mathbf{w}_{r,t}\right\}.$$

Moreover, let the infinite width NTK be defined as

$$\mathbf{H}_{ii'}^{\infty} = \frac{1}{p^2}\sum_{j=1}^{p}\left\langle\hat{\mathbf{x}}_i^{(j)},\hat{\mathbf{x}}_{i'}^{(j)}\right\rangle\mathbb{E}_{\mathbf{w}\sim\mathcal{N}(0,\mathbf{I})}\left[\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j)};\mathbf{w}\right\}\right].$$

Let $\lambda_0 = \lambda_{\min}\left(\mathbf{H}^{\infty}\right)$. Note that since $\hat{\mathbf{x}}_i \not\parallel \hat{\mathbf{x}}_{i'}$ for $i \neq i'$. Thus $\hat{\mathbf{x}}_i^{(j)} \not\parallel \hat{\mathbf{x}}_{i'}^{(j)}$ for $i \neq i'$. () shows that the matrix $\hat{\mathbf{H}}(j)$, as defined below, is positive definite for all $j \in [p]$:

$$\hat{\mathbf{H}}(j)_{ii'}^{\infty} = \left\langle\hat{\mathbf{x}}_i^{(j)},\hat{\mathbf{x}}_{i'}^{(j)}\right\rangle\mathbb{E}_{\mathbf{w}\sim\mathcal{N}(0,\mathbf{I})}\left[\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j)};\mathbf{w}\right\}\right].$$

Since $\mathbf{H}^{\infty} = p^{-2}\sum_{j=1}^{p}\hat{\mathbf{H}}(j)^{\infty}$, we have that $\mathbf{H}^{\infty}$ is positive definite and thus $\lambda_0 > 0$. The following lemma shows that the NTK remains positive definite throughout training.

Lemma 4.5 *Let $m = \Omega\left(\lambda_0^{-2}n^2\log\frac{n}{\delta}\right)$. If for all $r \in [m]$ and all t we have $\|\mathbf{w}_{r,t}-\mathbf{w}_{r,0}\|_2 \leq R := O\left(\frac{\kappa\lambda_0}{n}\right)$. Then with probability at least $1-\delta$ we have that for all t:*

$$\lambda_{\min}\left(\mathbf{H}(t)\right) \geq \frac{\lambda_0}{2}.$$

Proof: To start, we notice that for all $r \in [m]$:

$$\begin{aligned}
&\mathbb{E}_{\mathbf{W}_0,\mathbf{a}}\left[\sum_{j,j'=1}^{p}a_{rj}a_{rj'}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,0}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j')};\mathbf{w}_{r,0}\right\}\right] \\
&\qquad = \frac{1}{p^2m}\mathbb{E}_{\mathbf{w}\sim\mathcal{N}(0,\mathbf{I})}\left[\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j)};\mathbf{w}\right\}\right]
\end{aligned}$$

Moreover, we have that:

$$\left|\sum_{j,j'=1}^{p} a_{rj}a_{rj'}\left\langle \hat{\mathbf{x}}_i^{(j)}, \hat{\mathbf{x}}_{i'}^{(j')}\right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,0}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j')}; \mathbf{w}_{r,0}\right\}\right| \leq 1.$$

Thus, we can apply Hoeffding's inequality with bounded random variable to get that

$$\mathbb{P}\left(\left|\mathbf{H}(0)_{i,i'} - \mathbf{H}_{i,i'}^{\infty}\right| \geq t\right) \leq 2\exp\left(-mt^2\right).$$

Therefore, with probability at least $1 - O(\delta)$ it holds that for all $i, i' \in [n]$:

$$\left|\mathbf{H}(0)_{i,i'} - \mathbf{H}_{i,i'}^{\infty}\right| \leq \frac{\log\frac{n}{\delta}}{\sqrt{m}},$$

which implies that:

$$\|\mathbf{H}(0) - \mathbf{H}^{\infty}\| \leq \|\mathbf{H}(0) - \mathbf{H}^{\infty}\|_F \leq \frac{n\left(\log\delta^{-1} + \log n\right)}{\sqrt{m}}.$$

As long as $m = \Omega\left(\lambda_0^{-2}n^2\log\frac{n}{\delta}\right)$, we will have:

$$\|\mathbf{H}(t) - \mathbf{H}^{\infty}\| \leq \frac{\lambda_0}{4}.$$

Now, we move on to bound $\|\mathbf{H}(t) - \mathbf{H}(0)\|$. We have that:

$$\mathbf{H}(t)_{i,i'} - \mathbf{H}(0)_{i,i'} = \sum_{r=1}^{m}\sum_{j,j'=1}^{p} a_{rj}a_{rj'}\left\langle \hat{\mathbf{x}}_i^{(j)}, \hat{\mathbf{x}}_{i'}^{(j')}\right\rangle z_{r,i,i'}^{(j,j')},$$

with

$$z_{r,i,i'}^{(j,j')} = \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j')}; \mathbf{w}_{r,t}\right\} - \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,0}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j')}; \mathbf{w}_{r,0}\right\}.$$

We observe that $|z_{r,i,i'}^{(j,j')}|$ only if $A_{ir}^{(j)} \vee A_{i'r}^{(j')}$. Therefore:

$$\mathbb{E}_{\mathbf{W}}\left[z_{r,i,i'}^{(j,j')}\right] \leq \mathbb{P}\left(A_{ir}^{(j)}\right) + \mathbb{P}\left(A_{i'r}^{(j')}\right) \leq \frac{4R}{\kappa\sqrt{2\pi}}.$$

For the case $j = j'$, we first notice that

$$\mathbb{E}_{\mathbf{W}}\left[\left(z_{r,i,i'}^{(j,j')} - \mathbb{E}_{\mathbf{W}}\left[z_{r,i,i'}^{(j,j')}\right]\right)^2\right] \leq \mathbb{E}_{\mathbf{W}}\left[z_{r,i,i'}^{(j,j')2}\right] \leq \frac{4R}{\kappa\sqrt{2\pi}}.$$

Thus, applying Berstein Inequality to the case $j = j'$ we have that:

$$\mathbb{P}\left(\sum_{r=1}^{m} z_{r,i,i'}^{(j,j)} \geq m\left(\mathbb{E}_{\mathbf{W}}\left[z_{r,i,i'}^{(j,j')}\right] + t\right)\right) \leq \exp\left(-\frac{\kappa m t^2\sqrt{2\pi}}{8\left(1 + \frac{t}{3}\right)R}\right).$$

For the case $j \neq j'$, we notice that:

$$\mathbb{E}_{\mathbf{W},\mathbf{a}}\left[a_{rj}a_{rj'}z_{r,i,i'}^{(j,j')}\right] = 0.$$

Moreover:

$$\left|a_{rj}a_{rj'}z_{r,i,i'}^{(j,j')}\right| \leq \frac{4R}{p^2 m\kappa\sqrt{2\pi}},$$

$$\mathbb{E}_{\mathbf{W},\mathbf{a}}\left[\left(a_{rj}a_{rj'}z_{r,i,i'}^{(j,j')}\right)^2\right] = \frac{1}{p^4m^2}\mathbb{E}_{\mathbf{W}}\left[z_{r,i,i'}^{(j,j')2}\right] \leq \frac{4R}{p^4m^2\kappa\sqrt{2\pi}}.$$

Applying the Berstein Inequality to the case $j \neq j'$, we have that:

$$\mathbb{P}\left(\sum_{r=1}^{m} a_{rj}a_{rj'}z_{r,i,i'}^{(j,j')} \geq \frac{t}{p^2}\right) \leq \exp\left(-\frac{m\kappa t^2\sqrt{2\pi}}{8\left(1+\frac{t}{3}\right)R}\right).$$

Combining both cases, we have that with probability at least $1 - p^2\exp\left(-\frac{m\kappa t^2\sqrt{2\pi}}{8\left(1+\frac{t}{3}\right)R}\right)$, it holds that:

$$|\mathbf{H}(t)_{i,i'} - \mathbf{H}(0)_{i,i'}| \leq p^{-1}\mathbb{E}_{\mathbf{W}}\left[z_{r,i,i'}^{(j,j')}\right] + t \leq \frac{2R}{p\kappa} + t^2.$$

Choose $t = \kappa^{-1}R$. Then as long as $m = \frac{\log\frac{np}{\delta}}{R}$, it holds that with probability at least $1 - O(\delta)$:

$$|\mathbf{H}(t)_{i,i'} - \mathbf{H}(0)_{i,i'}| \leq 3\kappa^{-1}R.$$

This implies that:

$$\|\mathbf{H}(t) - \mathbf{H}(0)\|_2 \leq \|\mathbf{H}(t) - \mathbf{H}(0)\|_F \leq 3n\kappa^{-1}R.$$

Thus, $\|\mathbf{H}(t) - \mathbf{H}(0)\|_2 \leq \frac{\lambda_0}{4}$ as long as $R = O\left(\frac{\kappa\lambda_0}{n}\right)$. This shows that $\lambda_{\min}\left(\mathbf{H}(t)\right) \geq \frac{\lambda_0}{2}$ for all t with probability at least $1 - O(\delta)$.

4.4 Surrogate Gradient Bound

As we see in previous section, the `AsyncDrop` scheme can be written as

$$\mathbf{w}_{r,t+1} = \mathbf{w}_{r,t} - \eta\mathbf{g}_{r,t-\delta_t}; \quad \mathbf{g}_{r,t} = \sum_{i=1}^{n}\sum_{j=1}^{p} a_{rj}\left(\tilde{u}_{r,t}^{(i)} - N_{r,t}^{\perp}y_i\right)\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\}$$

with $\tilde{u}_{r,t}^{(i)}$ defined as

$$\tilde{u}_{r,t}^{(i)} = \frac{N_{r,t}^{\perp}}{N_{r,t}}\sum_{s=1}^{S} m_{r,t}^{(s)}\hat{u}_t^{(s,i)} = \sum_{r'=1}^{m}\sum_{j=1}^{p}\underbrace{\left(\frac{N_{r,t}^{\perp}}{N_{r,t}}\sum_{s=1}^{S} m_{r,t}^{(s)}m_{r',t}^{(s)}\right)}_{\nu_{r,r',t}} a_{rj}\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t}\right\rangle\right)$$

The mixing of the surrogate function $\tilde{u}_{r,t}^{(i)}$ can be bounded by

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\tilde{u}_{r,t}^{(i)}\right] &= \sum_{s=1}^{S}\sum_{r'=1}^{m}\sum_{j=1}^{p}\mathbb{E}_{\mathbf{M}_t}\left[m_{r,t}^{(s)}m_{r,t'}^{(s)}\cdot\frac{N_{r,t}^{\perp}}{N_{r,t}}\right]a_{r'j}\sigma\left(\left\langle\hat{\mathbf{x}}^{(j)},\mathbf{w}_{r,t'}\right\rangle\right) \\
&= \xi\theta\sum_{r'=1}^{m}\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle\hat{\mathbf{x}}^{(j)},\mathbf{w}_{r,t}\right\rangle\right) + (1-\xi)\theta\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t}\right\rangle\right) \\
&= \theta u_t^{(i)} + (1-\xi)\theta\underbrace{\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t}\right\rangle\right)}_{\hat{\epsilon}_{r,t}^{(i)}}
\end{aligned}$$

Therefore,

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\mathbf{g}_{r,t}\right] &= \sum_{i=1}^{n}\sum_{j=1}^{n} a_{rj}\mathbb{E}_{\mathbf{M}_t}\left[\tilde{u}_t^{(i)} - N_{r,t}^{\perp}y_i\right]\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&= \xi^{-1}\theta\nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right) + (1-\xi)\theta\underbrace{\sum_{i=1}^{n}\sum_{j'=1}^{p} a_{rj'}\hat{\epsilon}_{r,t}^{(i)}\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\}}_{\boldsymbol{\epsilon}_{r,t}}
\end{aligned}$$

Now, we have

$$\left|\hat{\epsilon}_{r,t}^{(i)}\right| \leq \frac{1}{\sqrt{m}} \left\|\mathbf{w}_{r,t}\right\|_2 ; \quad \left\|\boldsymbol{\epsilon}_{r,t}\right\|_2 \leq \frac{n}{\sqrt{m}} \left|\hat{\epsilon}_{r,t}^{(i)}\right| \leq \frac{n}{m} \left\|\mathbf{w}_{r,t}\right\|_2$$

Moreover, we would like to investigate the norm and norm squared of the gradient. In particular, we first notice that, under the case of $N_{t,r}^{\perp} = 1$, we have

$$\mathbf{g}_{r,t} = \xi^{-1} \nabla_{\mathbf{w}_r} \mathcal{L}\left(\mathbf{W}_t\right) + \sum_{i=1}^{n} \sum_{j=1}^{p} a_{rj} \left(\tilde{u}_{r,t}^{(i)} - u_t^{(i)}\right) \hat{\mathbf{x}}_i^{(j)} \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}$$

Thus, we are interested in $\left\|\tilde{\mathbf{u}}_{r,t} - \mathbf{u}_t\right\|_2$. Following from previous work [4] (lemma 19, 20, and 21), we have that

$$\mathbb{E}_{\mathbf{M}_t}\left[\nu_{r,r',t} \mid N_{r,t}^{\perp} = 1\right] \begin{cases} \xi & \text{if } r \neq r' \\ 1 & \text{if } r = r' \end{cases} \qquad \text{Var}\left(\nu_{r,r',t} \mid N_{r,t}^{\perp} = 1\right) = \begin{cases} \frac{\theta - \xi^2}{S} & \text{if } r \neq r' \\ 0 & \text{if } r = r' \end{cases}$$

Therefore

$$\begin{aligned}
&\mathbb{E}_{\mathbf{M}_t}\left[\left\|\tilde{\mathbf{u}}_{r,t} - \mathbf{u}_t\right\|_2^2 \mid N_{r,t}^{\perp} = 1\right] \\
&= \sum_{i=1}^{n} \mathbb{E}_{\mathbf{M}_t}\left[\left(\sum_{r'=1}^{m} \sum_{j=1}^{p} \left(\nu_{r,r',t} - \xi\right) a_{rj} \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r',t} \right\rangle\right)\right)^2 \mid N_{r,t}^{\perp} = 1\right] \\
&= \sum_{i=1}^{n} \sum_{r'=1}^{m} \mathbb{E}_{\mathbf{M}_t}\left[\left(\nu_{r,r',t} - \xi\right)^2 \mid N_{r,t}^{\perp} = 1\right] \left(\sum_{j=1}^{p} a_{rj} \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r',t} \right\rangle\right)\right)^2 \\
&\leq \frac{1}{mp} \sum_{i=1}^{n} \sum_{r'=1}^{m} \mathbb{E}_{\mathbf{M}_t}\left[\left(\nu_{r,r',t} - \xi\right)^2 \mid N_{r,t}^{\perp} = 1\right] \sum_{j=1}^{p} \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r',t} \right\rangle\right)^2 \\
&\leq \frac{1}{mp} \sum_{i=1}^{n} \sum_{r' \neq r} \text{Var}\left(\nu_{r,r',t} \mid N_{r,t}^{\perp} = 1\right) \sum_{j=1}^{p} \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r',t} \right\rangle\right)^2 + \\
&\qquad \frac{1}{mp} \sum_{i=1}^{n} \mathbb{E}_{\mathbf{M}_t}\left[\left(\nu_{r,r,t} - \xi\right)^2 \mid N_{r,t}^{\perp} = 1\right] \sum_{j=1}^{p} \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r',t} \right\rangle\right)^2 \\
&\leq \frac{1}{mpS} (\theta - \xi) \sum_{r' \neq r} \sum_{i=1}^{n} \sum_{j=1}^{p} \left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r',t} \right\rangle^2 + \frac{1}{mp} (\theta - \xi^2) \sum_{i=1}^{n} \sum_{j=1}^{p} \left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t} \right\rangle^2
\end{aligned}$$

With high probabilty it holds that

$$\sum_{r=1}^{m} \sum_{i=1}^{n} \sum_{j=1}^{p} \left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,0} \right\rangle^2 \leq 2mnp\kappa^2 - mnpR^2$$

Thus, with sufficiently large m, the second term is always smaller than the first term, and we have

$$\mathbb{E}_{\mathbf{M}_t}\left[\left\|\tilde{\mathbf{u}}_{r,t} - \mathbf{u}_t\right\|_2^2 \mid N_{r,t}^{\perp} = 1\right] \leq 8n\kappa^2(\theta - \xi^2) S^{-1}$$

Now, we can compute that

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{g}_{r,t}\right\|_2 \mid N_{r,t}^{\perp} = 1\right] &= \mathbb{E}_{\mathbf{M}_t}\left[\left\|\sum_{i=1}^{n} \sum_{j=1}^{p} a_{rj} \left(\tilde{u}_{r,t}^{(i)} - u_t^{(i)}\right) \hat{\mathbf{x}}_i^{(j)} \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}\right\| \mid N_{r,t} = 1\right] + \\
&\qquad \xi^{-1} \theta \left\|\nabla_{\mathbf{w}_r} \mathcal{L}\left(\mathbf{W}_t\right)\right\|_2 \\
&= \sqrt{\frac{n}{m}} \mathbb{E}_{\mathbf{M}_t}\left[\left\|\tilde{\mathbf{u}}_{r,t} - \mathbf{u}_t\right\|_2 \mid N_{r,t} = 1\right] + \theta \sqrt{\frac{n}{m}} \left\|\mathbf{u}_t - \mathbf{y}\right\|_2 \\
&\leq n\kappa \sqrt{\frac{\theta - \xi^2}{mS}} + \theta \sqrt{\frac{n}{m}} \left\|\mathbf{u}_t - \mathbf{y}\right\|_2
\end{aligned}$$

And we know that $\mathbf{g}_{r,t} = 0$ when $N_{r,t}^{\perp} = 0$. Therefore,

$$\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{g}_{r,t}\|_2\right] \leq \theta n \kappa \sqrt{\frac{\theta - \xi^2}{mS}} + \theta^2 \sqrt{\frac{n}{m}} \|\mathbf{u}_t - \mathbf{y}\|_2$$

Similarly

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{g}_{r,t}\|_2^2 \mid N_{r,t}^{\perp} = 1\right] &= \frac{2n}{m}\mathbb{E}_{\mathbf{M}_t}\left[\|\tilde{\mathbf{u}}_{r,t} - \mathbf{u}_t\|_2^2 \mid N_{r,t} = 1\right] + \frac{2\theta^2 n}{m}\|\mathbf{u}_t - \mathbf{y}\|_2^2 \\
&\leq \frac{16 n^2 \kappa^2 \left(\theta - \xi^2\right)}{mS} + \frac{2\theta^2 n}{m}\|\mathbf{u}_t - \mathbf{y}\|_2^2
\end{aligned}$$

and thus

$$\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{g}_{r,t}\|_2^2\right] \leq \frac{16\theta n^2 \kappa^2 \left(\theta - \xi^2\right)}{mS} + \frac{2\theta^3 n}{m}\|\mathbf{u}_t - \mathbf{y}\|_2^2$$

4.5 Step-wise Convergence

Consider

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{u}_{t+1} - \mathbf{y}\|_2^2\right] &= \|\mathbf{u}_t - \mathbf{y}\|_2^2 - 2\left\langle \mathbf{u}_t - \mathbf{u}_{t+1}, \mathbf{u}_t - \mathbf{y}\right\rangle + \|\mathbf{u}_t - \mathbf{u}_{t+1}\|_2^2 \\
&= \|\mathbf{u}_t - \mathbf{y}\|_2^2 - 2\left\langle \mathbf{I}_{1,t} + \mathbf{I}_{2,t}, \mathbf{u}_t - \mathbf{y}\right\rangle + \|\mathbf{u}_t - \mathbf{u}_{t+1}\|_2^2 \\
&\leq \|\mathbf{u}_t - \mathbf{y}\|_2^2 - 2\left\langle \mathbf{I}_{1,t}, \mathbf{u}_t - \mathbf{y}\right\rangle + 2\|\mathbf{I}_{2,t}\|_2 \|\mathbf{u}_t - \mathbf{y}\|_2 + \|\mathbf{u}_t - \mathbf{u}_{t+1}\|_2^2
\end{aligned}$$

Let $P_{ij} = \left\{r \in [m] : \neg A_{ir}^{(j)}\right\}$. Here $\mathbf{I}_{1,t}$ and $\mathbf{I}_{2,t}$ are characterized as in previous work.

$$\begin{aligned}
I_{1,t}^{(i)} &= \sum_{j=1}^{p} \sum_{r \in P_{ij}} a_{rj}\left(\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t}\right\rangle\right) - \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t+1}\right\rangle\right)\right) \\
I_{2,t}^{(i)} &= \sum_{j=1}^{p} \sum_{r \in P_{ij}^{\perp}} a_{rj}\left(\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t}\right\rangle\right) - \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t+1}\right\rangle\right)\right)
\end{aligned}$$

We first bound the magnitude of $\mathbf{I}_{2,t}$

$$\begin{aligned}
\left|I_{2,t}^{(i)}\right| &= \frac{1}{p\sqrt{m}} \sum_{j=1}^{p} \sum_{r \in P_{ij}^{\perp}} \left|\left(\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t}\right\rangle\right) - \sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t+1}\right\rangle\right)\right)\right| \\
&\leq \frac{1}{p\sqrt{m}} \sum_{j=1}^{p} \sum_{r \in P_{ij}} \|\mathbf{w}_{r,t} - \mathbf{w}_{r,t+1}\|_2 \\
&\leq \frac{\eta}{p\sqrt{m}} \sum_{j=1}^{p} \sum_{r \in P_{ij}} \|\mathbf{g}_{r,t-\delta_t}\|_2 \\
&\leq \frac{\eta}{\sqrt{m}} \cdot 3m\kappa^{-1}R \cdot \left(n\kappa\sqrt{\frac{\theta - \xi^2}{mS}} + \theta\sqrt{\frac{n}{m}}\|\mathbf{u}_{t-\delta_t} - \mathbf{y}\|_2\right) \\
&= 3\kappa^{-1}\theta\eta\sqrt{n}R\|\mathbf{u}_{t-\delta_t} - \mathbf{y}\|_2 + 3\eta n R\sqrt{\frac{\theta - \xi^2}{S}}
\end{aligned}$$

Thus

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{I}_{2,t}\|_2\right] &\leq \mathbb{E}_{\mathbf{M}_t}\left[\sqrt{n}\max_{i\in[n]}\left|I_{2,t}^{(i)}\right|\right] \\
&\leq \frac{\eta}{p}\sqrt{\frac{n}{m}}\max_{i\in[n]}\sum_{j=1}^{p}\sum_{r\in P_{ij}}\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{g}_{r,t-\delta_t}\|_2\right] \\
&\leq \eta\sqrt{\frac{n}{m}}\cdot 3m\kappa^{-1}R\cdot\left(n\kappa\theta\sqrt{\frac{\theta-\xi^2}{mS}}+\theta^2\sqrt{\frac{n}{m}}\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\|_2\right) \\
&= 3\kappa^{-1}\theta^2\eta nR\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\|_2+3\theta\eta n^{\frac{3}{2}}R\sqrt{\frac{\theta-\xi^2}{S}}
\end{aligned}$$

Therefore, given assumption $\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\|_2 \geq \|\mathbf{u}_t-\mathbf{y}\|_2$

$$\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{I}_{2,t}\|_2\|\mathbf{u}_t-\mathbf{y}\|_2\right] \leq 6\theta\eta n\kappa^{-1}R\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\|_2^2+3S^{-1}\theta\left(\theta-\xi^2\right)\eta n^2\kappa R$$

Letting $R=O\left(\frac{\kappa\lambda_0}{n}\right)$ gives that

$$\mathbb{E}_{\mathbf{M}_t}\left[\|\mathbf{I}_{2,t}\|_2\|\mathbf{u}_t-\mathbf{y}\|_2\right] = O\left(\theta\eta\lambda_0\right)\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\|_2^2+O\left(\theta\eta\lambda_0(\theta-\xi^2)S^{-1}n\kappa^2\right)$$

As in previous work, $I_{1,t}^{(i)}$ can be written as

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[I_{1,t}^{(i)}\right] &= \xi\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\mathbb{E}_{\mathbf{M}_t}\left[\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t}\right\rangle\right)-\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t+1}\right\rangle\right)\right] \\
&= \xi\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbb{E}_{\mathbf{M}_t}\left[\mathbf{w}_{r,t}-\mathbf{w}_{r,t+1}\right]\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&= \xi\eta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbb{E}_{\mathbf{M}_t}\left[\mathbf{g}_{r,t-\delta_t}\right]\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&= \xi\eta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbb{E}_{\mathbf{M}_t}\left[\mathbf{g}_{r,t-\delta_t}-\mathbf{g}_{r,t}+\mathbf{g}_{r,t}\right]\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&= \xi\eta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbb{E}_{\mathbf{M}_t}\left[\mathbf{g}_{r,t}\right]\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\}+ \\
&\qquad \xi\eta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbb{E}_{\mathbf{M}_t}\left[\mathbf{g}_{r,t-\delta_t}-\mathbf{g}_{r,t}\right]\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&= \eta\theta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right)\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&\qquad +\xi(1-\xi)\theta\eta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\boldsymbol{\epsilon}_{r,t}\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&\qquad +\eta\theta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_{t-\delta_t}\right)-\nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right)\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\} \\
&\qquad +\xi(1-\xi)\theta\eta\sum_{j=1}^{p}\sum_{r\in P_{ij}}a_{rj}\left\langle\hat{\mathbf{x}}_i^{(j)},\boldsymbol{\epsilon}_{r,t-\delta_t}-\boldsymbol{\epsilon}_{r,t}\right\rangle\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\}
\end{aligned}$$

$$
\begin{aligned}
&= \eta\theta\xi \sum_{i'=1}^{n} \sum_{j,j'=1}^{p} \sum_{r\in P_{ij}} \left(u_t^{(i')} - y_{i'}\right) a_{rj} a_{rj'} \left\langle \hat{\mathbf{x}}_i^{(j)}, \hat{\mathbf{x}}_{i'}^{(j')} \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\} \cdot \mathbb{I}\left\{\hat{\mathbf{x}}_{i'}^{(j')}; \mathbf{w}_{r,t}\right\} \\
&\quad + \xi(1-\xi)\theta\eta \sum_{j=1}^{p} \sum_{r\in P_{ij}} a_{rj} \left\langle \hat{\mathbf{x}}_i^{(j)}, \boldsymbol{\epsilon}_{r,t} \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\} \\
&\quad + \eta\theta \sum_{j=1}^{p} \sum_{r\in P_{ij}} a_{rj} \left\langle \hat{\mathbf{x}}_i^{(j)}, \nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_{t-\delta_t}\right) - \nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right) \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\} \\
&\quad + \xi(1-\xi)\theta\eta \sum_{j=1}^{p} \sum_{r\in P_{ij}} a_{rj} \left\langle \hat{\mathbf{x}}_i^{(j)}, \boldsymbol{\epsilon}_{r,t-\delta_t} - \boldsymbol{\epsilon}_{r,t} \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\} \\
&= \eta\theta \sum_{i'=1}^{n} \left(\mathbf{H}(t)_{ii'} - \mathbf{H}(t)_{ii'}^{\perp}\right) \left(u_t^{(i')} - y_{i'}\right) \\
&\quad \underbrace{+ \xi(1-\xi)\theta\eta \sum_{j=1}^{p} \sum_{r\in P_{ij}} a_{rj} \left\langle \hat{\mathbf{x}}_i^{(j)}, \boldsymbol{\epsilon}_{r,t} \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}}_{\gamma_{1,i,t}} \\
&\quad \underbrace{+ \eta\theta \sum_{j=1}^{p} \sum_{r\in P_{ij}} a_{rj} \left\langle \hat{\mathbf{x}}_i^{(j)}, \nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_{t-\delta_t}\right) - \nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right) \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}}_{\gamma_{2,i,t}} \\
&\quad \underbrace{+ \xi(1-\xi)\theta\eta \sum_{j=1}^{p} \sum_{r\in P_{ij}} a_{rj} \left\langle \hat{\mathbf{x}}_i^{(j)}, \boldsymbol{\epsilon}_{r,t-\delta_t} - \boldsymbol{\epsilon}_{r,t} \right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}}_{\gamma_{3,i,t}}
\end{aligned}
$$

Note that:

$$
\begin{aligned}
|\gamma_{1,i,t}| &\le \xi(1-\xi)\theta\eta m^{-\frac{1}{2}} \sum_{r=1}^{m} \|\boldsymbol{\epsilon}_{r,t}\| \\
&\le \xi(1-\xi)\theta\eta n m^{-1} \|\mathbf{W}_t\|_F \\
&\le O\left(\xi(1-\xi)\theta\eta n\kappa\sqrt{\frac{d}{m}}\right)
\end{aligned}
$$

Given that:

$$
\begin{aligned}
\left\|\left(u_{t-\delta_t}^{(i)} - u_t^{(i)}\right)\right\|_2 &= \left\|\xi \sum_{r=1}^{m}\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t-\delta_t}\right\rangle\right) - \xi \sum_{r=1}^{m}\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t}\right\rangle\right)\right\|_2 \\
&\le \xi \sum_{r=1}^{m}\sum_{j=1}^{p} \left\|a_{rj}\hat{\mathbf{x}}_i^{(j)}\right\|_2 \|\mathbf{w}_{r,t-\delta_t} - \mathbf{w}_{r,t}\|_2 \le \xi\frac{1}{\sqrt{m}}R_E
\end{aligned}
$$

We show:

$$
\begin{aligned}
&\|\nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_{t-\delta_t}\right) - \nabla_{\mathbf{w}_r}\mathcal{L}\left(\mathbf{W}_t\right)\|_2 \\
&\quad = \left\|\xi \sum_{i=1}^{n}\sum_{j=1}^{p} \left(u_{t-\delta_t}^{(i)} - y_i\right) a_{rj}\hat{\mathbf{x}}_i^{(j)} \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t-\delta_t}\right\} - \xi \sum_{i=1}^{n}\sum_{j=1}^{p} \left(u_t^{(i)} - y_i\right) a_{rj}\hat{\mathbf{x}}_i^{(j)} \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}\right\|_2 \\
&\quad \le \xi \sum_{i=1}^{n}\sum_{j=1}^{p} \left\|a_{rj}\hat{\mathbf{x}}_i^{(j)}\right\|_2 \left\|\left(u_{t-\delta_t}^{(i)} - y_i\right) \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t-\delta_t}\right\} - \left(u_t^{(i)} - y_i\right) \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}\right\|_2 \\
&\quad \le \xi \sum_{i=1}^{n}\sum_{j=1}^{p} \left\|a_{rj}\hat{\mathbf{x}}_i^{(j)}\right\|_2 \left(\left\|\left(u_{t-\delta_t}^{(i)} - u_t^{(i)}\right)\right\|_2 + \left\|\left(u_t^{(i)} - y_i\right)\right\|_2\right) \\
&\quad \le \xi^2 \frac{n}{m} R_E + \xi\sqrt{\frac{n}{m}} \|\mathbf{u}_t - \mathbf{y}\|_2
\end{aligned}
$$

Thus,

$$\|\gamma_{2,i,t}\|_2 \leq 2\eta\theta n m^{-2}\xi^2 R_E + \eta\theta\xi n^{\frac{1}{2}} m^{-\frac{3}{2}} \|\mathbf{u}_t - \mathbf{y}\|_2$$

Given that:

$$\begin{aligned}
\left\|\hat{\epsilon}_{r,t-\delta_t}^{(i)} - \hat{\epsilon}_{r,t}^{(i)}\right\|_2 &= \left\|\sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t-\delta_t}\right\rangle\right) - \sum_{j=1}^{p} a_{rj}\sigma\left(\left\langle \hat{\mathbf{x}}_i^{(j)}, \mathbf{w}_{r,t}\right\rangle\right)\right\|_2 \\
&\leq \sum_{j=1}^{p} \left\|a_{rj}\hat{\mathbf{x}}_i^{(j)}\right\|_2 \left(\|\mathbf{w}_{r,t-\delta_t} - \mathbf{w}_{r,t}\|_2 + \|\mathbf{w}_{r,t}\|_2\right) \\
&\leq \kappa\sqrt{\frac{2d}{m}} + \frac{1}{\sqrt{m}} R_E
\end{aligned}$$

We show:

$$\begin{aligned}
\left\|\boldsymbol{\epsilon}_{r,t-\delta_t}^{(i)} - \boldsymbol{\epsilon}_{r,t}^{(i)}\right\|_2 &= \left\|\sum_{i=1}^{n}\sum_{j'=1}^{p} a_{rj'}\hat{\epsilon}_{r,t-\delta_t}^{(i)}\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t-\delta_t}\right\} - \sum_{i=1}^{n}\sum_{j'=1}^{p} a_{rj'}\hat{\epsilon}_{r,t}^{(i)}\hat{\mathbf{x}}_i^{(j)}\mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)}; \mathbf{w}_{r,t}\right\}\right\|_2 \\
&\leq \sum_{i=1}^{n}\sum_{j'=1}^{p} \left\|a_{rj'}\hat{\mathbf{x}}_i^{(j)}\right\|_2 \left(\left\|\left(\hat{\epsilon}_{r,t-\delta_t}^{(i)} - \hat{\epsilon}_{r,t}^{(i)}\right)\right\|_2 + \left\|\left(\hat{\epsilon}_{r,t}^{(i)}\right)\right\|_2\right) \\
&\leq \frac{n}{\sqrt{m}}\left(\kappa\sqrt{\frac{2d}{m}} + \frac{1}{\sqrt{m}}R_E + \frac{1}{\sqrt{m}}\kappa\sqrt{2dm}\right)
\end{aligned}$$

Thus,

$$\|\gamma_{3,i,t}\|_2 \leq \sqrt{2}\xi(1-\xi)\theta\eta n m^{-\frac{3}{2}}\kappa\sqrt{d} + 2\xi(1-\xi)\theta\eta n m^{-\frac{3}{2}} R_E + \sqrt{2}\xi(1-\xi)\theta\eta n m^{-1}\kappa\sqrt{d})$$

This implies that:

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\langle \mathbf{I}_{1,t}, \mathbf{u}_t - \mathbf{y}\rangle\right] &= \eta\theta \sum_{i,i'=1}^{n} \left(u_t^{(i)} - y_i\right)\left(\mathbf{H}(t)_{ii'} - \mathbf{H}(t)_{ii'}^{\perp}\right)\left(u_t^{(i')} - y_{i'}\right) \\
&\quad + \sum_{i=1}^{n}\gamma_{1,i,t}\left(u_t^{(i)} - y_i\right) + \sum_{i=1}^{n}\gamma_{2,i,t}\left(u_t^{(i)} - y_i\right) + \sum_{i=1}^{n}\gamma_{3,i,t}\left(u_t^{(i)} - y_i\right) \\
&= \eta\theta\left\langle \mathbf{u}_t - \mathbf{y}, \left(\mathbf{H}(t) - \mathbf{H}(t)^{\perp}\right)(\mathbf{u}_t - \mathbf{y})\right\rangle \\
&\quad + \sum_{i=1}^{n}\gamma_{1,i,t}\left(u_t^{(i)} - y_i\right) + \sum_{i=1}^{n}\gamma_{2,i,t}\left(u_t^{(i)} - y_i\right) + \sum_{i=1}^{n}\gamma_{3,i,t}\left(u_t^{(i)} - y_i\right) \\
&\geq \eta\theta\left(\lambda_{\min}\left(\mathbf{H}(t)\right) - \lambda_{\max}\left(\mathbf{H}(t)^{\perp}\right)\right)\|\mathbf{u}_t - \mathbf{y}\|_2^2 \\
&\quad - \sum_{i=1}^{n}|\gamma_{1,i,t}|\cdot\left|u_t^{(i)} - y_i\right| - \sum_{i=1}^{n}|\gamma_{2,i,t}|\cdot\left|u_t^{(i)} - y_i\right| - \sum_{i=1}^{n}|\gamma_{3,i,t}|\cdot\left|u_t^{(i)} - y_i\right| \\
&\geq \eta\theta\left(\lambda_{\min}\left(\mathbf{H}(t)\right) - \lambda_{\max}\left(\mathbf{H}(t)^{\perp}\right)\right)\|\mathbf{u}_t - \mathbf{y}\|_2^2 \\
&\quad - \sqrt{n}\max_i|\gamma_{i,t}|\,\|\mathbf{u}_t - \mathbf{y}\|_2 \\
&\quad 2\eta\theta - n^{\frac{3}{2}}m^{-2}\xi^2 R_E\|\mathbf{u}_t - \mathbf{y}\|_2 - \eta\theta\xi n m^{-\frac{3}{2}}\|\mathbf{u}_t - \mathbf{y}\|_2^2 \\
&\quad - \sqrt{2}\xi(1-\xi)\theta\eta n m^{-\frac{3}{2}}\kappa\sqrt{d}\,\|\mathbf{u}_t - \mathbf{y}\|_2 - 2\xi(1-\xi)\theta\eta n m^{-\frac{3}{2}}R_E\|\mathbf{u}_t - \mathbf{y}\|_2 \\
&\quad \sqrt{2} - \xi(1-\xi)\theta\eta n m^{-1}\kappa\sqrt{d}\,\|\mathbf{u}_t - \mathbf{y}\|_2 \\
&\geq \eta\theta\left(\lambda_{\min}\left(\mathbf{H}(t)\right) - \lambda_{\max}\left(\mathbf{H}(t)^{\perp}\right) - O\left(\lambda_0\right) - n m^{-\frac{3}{2}}\right)\|\mathbf{u}_t - \mathbf{y}\|_2^2 \\
&\quad - O\left(\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0} + \frac{\eta^2\theta^2 n^3\xi^4 R_E^2}{m^4\lambda_0} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2 n^2\kappa^2 d}{m^3\lambda_0} + \right.
\end{aligned}$$

$$\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2R_E^2}{m^3\lambda_0}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2d}{m^2\lambda_0}\Bigg)$$

For $\mathbf{H}(t)^\perp$ we have that

$$\begin{aligned}
\lambda_{\max}\left(\mathbf{H}(t)^\perp\right)^2 &\leq \left\|\mathbf{H}(t)^\perp\right\|_F^2 \\
&\leq \sum_{i,i'=1}^{n}\left(\sum_{j,j'=1}^{p}\sum_{r\in P_{ij}} a_{rj}a_{rj'}\left\langle\hat{\mathbf{x}}_i^{(j)},\hat{\mathbf{x}}_{i,i'}^{(j,j')}\right\rangle \mathbb{I}\left\{\hat{\mathbf{x}}_i^{(j)};\mathbf{w}_{r,t}\right\}\mathbb{I}\left\{\hat{\mathbf{x}}_{i,i'}^{(j,j')};\mathbf{w}_{r,t}\right\}\right)^2 \\
&\leq \frac{n^2}{m^2}\left(\max_{ij}|P_{ij}|\right)^2 \\
&\leq n^2\kappa^{-2}R^2
\end{aligned}$$

Choosing $R=O\left(\frac{\kappa\lambda_0}{n}\right)$ gives

$$\lambda_{\max}\left(\mathbf{H}(t)^\perp\right)\leq O\left(\lambda_0\right)$$

We require $m\geq\Omega\left(\frac{n}{\lambda_0}\right)$, such that

$$nm^{-\frac{3}{2}}\leq O\left(\lambda_0\right)$$

Thus, Plugging in $\lambda_{\min}\left(\mathbf{H}(t)\right)\geq\frac{\lambda_0}{2}$, we have

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\langle\mathbf{I}_{1,t},\mathbf{u}_t-\mathbf{y}\rangle\right] \geq \eta\theta\lambda_0\left(\frac{1}{2}-O\left(1\right)\right)\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2 \\
-O\Bigg(\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2d}{m\lambda_0}+\frac{\eta^2\theta^2n\kappa^2\lambda_0\xi^4E^2}{m^4}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2d}{m^3\lambda_0}+ \\
\frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2d}{m^2\lambda_0}\Bigg).
\end{aligned}$$

Lastly, we analyze the last term in the quadratic expansion

$$\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{u}_t-\mathbf{u}_{t+1}\right\|_2^2\right] &= \sum_{i=1}^{n}\mathbb{E}_{\mathbf{M}_t}\left[\left(u_t^{(i)}-u_{t+1}^{(i)}\right)^2\right] \\
&\leq p^{-1}\sum_{i=1}^{n}\sum_{j=1}^{p}\sum_{r=1}^{m}\mathbb{E}_{\mathbf{M}_t}\left[\left(\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t}\right\rangle\right)-\sigma\left(\left\langle\hat{\mathbf{x}}_i^{(j)},\mathbf{w}_{r,t+1}\right\rangle\right)\right)^2\right] \\
&\leq p^{-1}\eta^2\sum_{i=1}^{n}\sum_{j=1}^{p}\sum_{r=1}^{m}\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{g}_{r,t-\delta_t}\right\|_2^2\right] \\
&\leq O\left(\theta^3\eta^2n^2\right)\left\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\right\|_2^2+O\left(\theta\eta^2n^2\kappa^2\left(\theta-\xi^2\right)S^{-1}\right)
\end{aligned}$$

Letting $\eta=O\left(\frac{\lambda_0}{n^2}\right)$ gives

$$\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{u}_t-\mathbf{u}_{t+1}\right\|_2^2\right]\leq O\left(\theta\eta\lambda_0\right)\left\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\right\|_2^2+O\left(\theta\eta\lambda_0\kappa^2\left(\theta-\xi^2\right)S^{-1}\right)$$

Putting all three terms together and as $\left\|\mathbf{u}_{t-\delta_t}-\mathbf{y}\right\|_2^2\leq\left\|\mathbf{u}_{t-E}-\mathbf{y}\right\|_2^2$, we have that

$$
\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{u}_{t+1}-\mathbf{y}\right\|_2^2\right] \leq{}& \left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2-\eta\theta\lambda_0\left(1-O\left(1\right)\right)\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2+O\left(\theta\eta\lambda_0\right)\left\|\mathbf{u}_{t-E}-\mathbf{y}\right\|_2^2 \\
&+O\left(\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0}+\frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}+\right. \\
&\qquad\left.\frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0}+\frac{n\kappa^2\left(\theta-\xi^2\right)}{S}\right) \\
\leq{}& \left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2-\eta\theta\lambda_0\left(1-O\left(1\right)\right)\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2+O\left(\theta\eta\lambda_0\right)\left(2\left\|\mathbf{u}_{t-E}-\mathbf{u}_t\right\|_2^2+2\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2\right) \\
&+O\left(\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0}+\frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}+\right. \\
&\qquad\left.\frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0}+\frac{n\kappa^2\left(\theta-\xi^2\right)}{S}\right) \\
\leq{}& \left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2-\eta\theta\lambda_0\left(1-O\left(1\right)\right)\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2+O\left(\theta\eta\lambda_0\right)\left(8\xi^2R_E^2+2\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2\right) \\
&+O\left(\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0}+\frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}+\right. \\
&\qquad\left.\frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0}+\frac{n\kappa^2\left(\theta-\xi^2\right)}{S}\right)
\end{aligned}
$$

Letting $O\left(\theta\eta\lambda_0\right)=\frac{1}{4}\theta\eta\lambda_0$,

$$
\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{u}_{t+1}-\mathbf{y}\right\|_2^2\right] \leq{}& \left(1-\frac{\theta\eta\lambda_0}{4}\right)\left\|\mathbf{u}_t-\mathbf{y}\right\|_2^2+ \\
&O\left(\frac{\theta\eta\lambda_0^3\xi^2\kappa^2E^2}{n^2}+\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0}+\frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}+\right. \\
&\qquad\left.\frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0}+\frac{n\kappa^2\left(\theta-\xi^2\right)}{S}\right)
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\mathbb{E}_{\mathbf{M}_t}\left[\left\|\mathbf{u}_{t+1}-\mathbf{y}\right\|_2^2\right] \leq{}& \left(1-\frac{\theta\eta\lambda_0}{4}\right)^t\left\|\mathbf{u}_0-\mathbf{y}\right\|_2^2 \\
&+O\left(\frac{\theta\eta\lambda_0^3\xi^2\kappa^2E^2}{n^2}+\frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0}+\frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4E^2}{m^4}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^3\lambda_0}+\right. \\
&\qquad\left.\frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0E^2}{m^3}+\frac{\xi^2(1-\xi)^2\theta^2\eta^2n^2\kappa^2 d}{m^2\lambda_0}+\frac{n\kappa^2\left(\theta-\xi^2\right)}{S}\right)
\end{aligned}
$$

4.6 Bounding Weight Perturbation

Next we show that $\|\mathbf{w}_{r,t} - \mathbf{w}_{r,0}\|_2 \leq R$ under sufficient over-parameterization. To start, we notice that

$$
\begin{aligned}
\mathbb{E}_{[\mathbf{M}_{t-1}]}[\|\mathbf{w}_{r,t} - \mathbf{w}_{r,0}\|_2] &\leq \sum_{t'=0}^{t-1} \mathbb{E}_{[\mathbf{M}_{t'}]}\left[\|\mathbf{w}_{r,t'+1} - \mathbf{w}_{r,t'}\|_2\right] \\
&\leq \eta \sum_{t'=0}^{t-1} \mathbb{E}_{[\mathbf{M}_{t'}]}\left[\left\|\mathbf{g}_{r,t'-\delta_{t'}}\right\|_2\right] \\
&\leq \eta \sum_{t'=0}^{t-1} \left(\theta^2 \sqrt{\frac{n}{m}} \mathbb{E}_{[\mathbf{M}_{t'-1}]}\left[\left\|\mathbf{u}_{t'-\delta_{t'}} - \mathbf{y}\right\|_2\right] + \theta n \kappa \sqrt{\frac{\theta - \xi^2}{mS}}\right) \\
&\leq \eta \theta^2 \sqrt{\frac{n}{m}} \sum_{t'=0}^{t-1} \mathbb{E}_{[\mathbf{M}_{t'-1}]}\left[\left\|\mathbf{u}_{t'-\delta_{t'}} - \mathbf{y}\right\|_2\right] + \eta t \theta n \kappa \sqrt{\frac{\theta - \xi^2}{mS}} \\
&\leq \eta \theta \sqrt{\frac{n}{m}} \|\mathbf{u}_0 - \mathbf{y}\|_2 \sum_{t'=0}^{t-1} \left(1 - \frac{\eta \theta \lambda_0}{4}\right)^{t'} + \eta T \theta n \kappa \sqrt{\frac{\theta - \xi^2}{mS}} \\
&\quad + \eta \theta T \cdot O\left(\frac{\theta \eta \lambda_0^3 \xi^2 \kappa^2}{n^2} + \frac{\xi^2 (1-\xi)^2 \theta \eta n^3 \kappa^2 d}{m \lambda_0} + \frac{\eta^2 \theta^2 n \kappa^2 \lambda_0 \xi^4}{m^4} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^3 \lambda_0} + \right. \\
&\qquad \left. \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 \kappa^2 \lambda_0}{m^3} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^2 \lambda_0}\right) \\
&\leq \eta \theta \sqrt{\frac{n}{m}} \frac{1 - \left(1 - \frac{\eta \theta \lambda_0}{4}\right)^t}{\frac{\eta \theta \lambda_0}{4}} \|\mathbf{u}_0 - \mathbf{y}\|_2 \\
&\quad + \eta \theta T \cdot O\left(\frac{\theta \eta \lambda_0^3 \xi^2 \kappa^2}{n^2} + \frac{\xi^2 (1-\xi)^2 \theta \eta n^3 \kappa^2 d}{m \lambda_0} + \frac{\eta^2 \theta^2 n \kappa^2 \lambda_0 \xi^4}{m^4} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^3 \lambda_0} + \right. \\
&\qquad \left. \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 \kappa^2 \lambda_0}{m^3} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^2 \lambda_0}\right) \\
&\leq \eta \theta \sqrt{\frac{n}{m}} \frac{1 - \left(1 - \frac{\eta \theta \lambda_0 t}{4}\right)}{\frac{\eta \theta \lambda_0}{4}} \|\mathbf{u}_0 - \mathbf{y}\|_2 \\
&\quad + \eta \theta T \cdot O\left(\frac{\theta \eta \lambda_0^3 \xi^2 \kappa^2}{n^2} + \frac{\xi^2 (1-\xi)^2 \theta \eta n^3 \kappa^2 d}{m \lambda_0} + \frac{\eta^2 \theta^2 n \kappa^2 \lambda_0 \xi^4}{m^4} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^3 \lambda_0} + \right. \\
&\qquad \left. \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 \kappa^2 \lambda_0}{m^3} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^2 \lambda_0}\right) \\
&\leq O\left(\eta \theta \lambda_0 t \sqrt{\frac{n}{m}}\right) \|\mathbf{u}_0 - \mathbf{y}\|_2 \\
&\quad + \eta \theta T \cdot O\left(\frac{\theta \eta \lambda_0^3 \xi^2 \kappa^2}{n^2} + \frac{\xi^2 (1-\xi)^2 \theta \eta n^3 \kappa^2 d}{m \lambda_0} + \frac{\eta^2 \theta^2 n \kappa^2 \lambda_0 \xi^4}{m^4} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^3 \lambda_0} + \right. \\
&\qquad \left. \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 \kappa^2 \lambda_0}{m^3} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^2 \lambda_0}\right)
\end{aligned}
$$

where the last inequality follows from the geometric sum and $\beta \leq O\left(p^{-1}\right)$. Using the initialization scale, we have that

$$
\begin{aligned}
\mathbb{E}_{\mathbf{W},\mathbf{a},[\mathbf{M}_t]}\left[\|\mathbf{w}_{r,t} - \mathbf{w}_{r,0}\|_2\right] \leq O\left(\eta \theta \lambda_0 t n m^{-\frac{1}{2}}\right) + \eta \theta T \cdot O\left(\frac{\theta \eta \lambda_0^3 \xi^2 \kappa^2}{n^2} + \frac{\xi^2 (1-\xi)^2 \theta \eta n^3 \kappa^2 d}{m \lambda_0} + \frac{\eta^2 \theta^2 n \kappa^2 \lambda_0 \xi^4}{m^4} + \right. \\
\left. \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^3 \lambda_0} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 \kappa^2 \lambda_0}{m^3} + \frac{\xi^2 (1-\xi)^2 \theta^2 \eta^2 n^2 \kappa^2 d}{m^2 \lambda_0}\right)
\end{aligned}
$$

With probability $1-\delta$, it holds for all $t \in [T]$ that

$$\|\mathbf{w}_{r,t} - \mathbf{w}_{r,0}\|_2 \leq O\left(\frac{t\lambda_0^2 K}{\delta n\sqrt{m}}\right) + \eta\theta T \cdot O\left(\frac{\theta\eta\lambda_0^3\xi^2\kappa^2}{n^2} + \frac{\xi^2(1-\xi)^2\theta\eta n^3\kappa^2 d}{m\lambda_0} + \frac{\eta^2\theta^2 n\kappa^2\lambda_0\xi^4}{m^4} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2 n^2\kappa^2 d}{m^3\lambda_0} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2\kappa^2\lambda_0}{m^3} + \frac{\xi^2(1-\xi)^2\theta^2\eta^2 n^2\kappa^2 d}{m^2\lambda_0}\right)$$

To enforce $\|\mathbf{w}_{r,t} - \mathbf{w}_{r,0}\|_2 \leq R := O\left(\frac{\kappa\lambda_0}{n}\right)$, we then require

$$m = \Omega\left(\frac{n^3 K^2}{\lambda_0^4\delta^2\kappa^2}\max\{n, d\}\right)$$

More specific,

$$\|\mathbf{w}_{r,t} - \mathbf{w}_{r,0}\|_2 \leq R_t := O\left(\frac{t\kappa\lambda_0}{n}\right)$$
$$\|\mathbf{w}_{r,t} - \mathbf{w}_{r,t-E}\|_2 \leq R_E := O\left(\frac{E\kappa\lambda_0}{n}\right)$$

References

[1] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, pages 1675–1685. PMLR, 2019. [Cited on page 5]

[2] Chen Dun, Cameron R Wolfe, Christopher M Jermaine, and Anastasios Kyrillidis. ResIST: Layer-wise decomposition of ResNets for distributed training. In *Uncertainty in Artificial Intelligence*, pages 610–620. PMLR, 2022. [Cited on pages 2 and 3]

[3] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016. [Cited on page 3]

[4] Fangshuo Liao and Anastasios Kyrillidis. On the convergence of shallow neural network training with randomly masked neurons, 2021. [Cited on page 10]

[5] Sebastian U Stich. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018. [Cited on page 2]