
Decentralized federated learning of deep neural networks on non-iid data

Noa Onozko^{*12} Gustav Karlsson^{*12} Olof Mogren¹ Edvin Listo Zec¹³

Abstract

We tackle the non-convex problem of learning a personalized deep learning model in a decentralized setting. More specifically, we study decentralized federated learning, a peer-to-peer setting where data is distributed among many clients and where there is no central server to orchestrate the training. In real world scenarios, the data distributions are often heterogeneous between clients. Therefore, in this work we study the problem of how to efficiently learn a model in a peer-to-peer system with non-iid client data. We propose a method named Performance-Based Neighbor Selection (PENS) where clients with similar data distributions detect each other and cooperate by evaluating their training losses on each other's data to learn a model suitable for the local data distribution. Our experiments on benchmark datasets show that our proposed method is able to achieve higher accuracies as compared to strong baselines.

1. Introduction

Federated learning (FL) (McMahan et al., 2017) is a framework developed to enable learning when data is distributed over several devices or across organizations, typically referred to as nodes or clients. In this framework, the training data never leaves the client, and all computations using the data are performed locally. This is especially useful when data privacy is important, or when collecting and storing data centrally is expensive.

Federated learning can be grouped into one of two categories: **centralized** and **decentralized**. In centralized FL, a central server orchestrates the learning among clients and

is responsible for parameter aggregation, after receiving parameter updates from clients. However, the central server in an FL setup is a potential point of weakness: it could fail or be maliciously attacked, which would make the distributed learning fail. *Decentralized* (peer-to-peer) systems without a central server are not vulnerable to this.

In *decentralized* federated learning, no global model state exists. Instead, the participating clients follow a communication protocol to reach a consensus of a model during training. Standard techniques for decentralized learning include gradient-based algorithms based on gossip learning (Boyd et al., 2006; Jelasity et al., 2007; Ormándi et al., 2013), where clients train their own model based on local data and follow a communication protocol where they randomly communicate (gossip) their model parameters with their neighbors. The goal for the participating clients is to reach a consensus on a good model. In this work, we focus on gradient-based learning algorithms in a decentralized federated setup.

Both centralized and decentralized federated learning approach the important question of how to learn a suitable personalized model when client data distributions differ, i.e. the setting of non-iid data. A lot of research is currently being done regarding this topic in the centralized setting. Meanwhile, this is a relatively understudied problem in the decentralized setting (Kairouz et al., 2019).

The main contribution of this paper is a novel, completely decentralized, federated algorithm for gradient-based methods when client data is non-iid: Performance-Based Neighbor Selection (PENS). In PENS, clients with similar data distributions have a higher probability of collaborating and those with dissimilar data distributions have a lower probability of collaborating. We perform multiple experiments over different non-convex optimization using deep neural networks, and our results show that using PENS leads to a higher performance as compared to all considered baselines.

2. Related work

Gossip learning. Gossip learning has been applied in many different machine learning settings (Kempe et al., 2003; Boyd et al., 2006; Ormándi et al., 2013). However, much of the previous work on gossip learning has been limited

^{*}Equal contribution. Author ordering decided by a coin flip.

¹RISE Research Institutes of Sweden ²Chalmers University of Technology ³KTH Royal Institute of Technology. Correspondence to: Edvin Listo Zec <edvin.listo.zec@ri.se>, Noa Onozko <noa.onozko@gmail.com>.

to settings where each client only stores a single data point. Further, it has been under-explored how non-convex optimization of neural networks works under the gossip learning protocol. In (Giarretta & Girdzijauskas, 2019), the authors study the performance of SVMs and linear regression models on non-iid data in a decentralized gossip learning setup. In (Hegedűs et al., 2019) the authors train and evaluate logistic regression models and compare gossip learning to federated learning with a central server. A gossip-based algorithm for strongly convex functions has been studied in (Koloskova et al., 2019), where the authors prove that their proposed algorithm is linearly convergent with quantized communication.

The first decentralized work on gossip-based optimization for non-convex deep learning studied CNNs and experimentally showed that an asynchronous and decentralized framework achieved high accuracies with low communication costs (Blot et al., 2016). Training CNNs in a decentralized federated learning setting has also been applied for segmentation of brain images (Roy et al., 2019).

Some recent work study communication costs (peer-to-peer communication) in non-convex optimization for different types of network topologies (Assran et al., 2019; Wang et al., 2019a).

In (Kong et al., 2021), the authors identify the changing consensus distance between clients as key to explain the gap between centralized and decentralized training and focus on non-convex optimization.

Non-iid data. All aforementioned works solve important problems. However, a key assumption is made in the studies: that data is independently and identically distributed (iid) over clients. The problem of non-iid is becoming more studied in the case of centralized federated learning. In this setting, solutions for skewed data distributions have been explored in many different ways including fine-tuning a global model locally (Wang et al., 2019b), posing the personalization problem in FL as a meta-learning objective (Jiang et al., 2019), using knowledge distillation techniques (Jeong et al., 2018), by mixing local and global models (Deng et al., 2020; Listo Zec et al., 2020) and with data-sharing methods (Zhao et al., 2018). Meanwhile, similar techniques have not yet been widely applied and researched in decentralized federated learning.

In (Ghosh et al., 2020) they study the problem of covariate shift, similar to us, but in a centralized federated learning setup. In their paper, the authors develop a client clustering framework to learn one global model per cluster with a central parameter server.

Label distribution shift for decentralized deep learning has been studied in (Niwa et al., 2020). In this work, the authors propose to solve a dual problem that seeks to minimize a

linearly constrained cost function. By solving a constrained optimization problem, their method achieves similar models among clients in the non-iid data setting.

In this work, we continue the research on the effectiveness of deep neural networks in decentralized peer-to-peer networks where data is non-iid. More specifically, we study the problem of covariate shift where $\mathcal{D}_i(x)$ varies, but the conditional distributions $\mathcal{D}_i(y|x) = \mathcal{D}_j(y|x)$ for all clients $i \neq j$.

3. Problem formulation

We formulate the problem as an empirical risk minimization (ERM) problem, as commonly used in statistical learning setups. The goal is to learn weights w for a model by optimizing some loss over data. In a decentralized setting, we have k clients that are able to communicate with neighboring clients in a communication network. We assume that each client i has a data distribution $\mathcal{D}_i(x, y)$ over input features x and labels y .

Let $\ell(w; z) : \mathbb{R}^d \rightarrow \mathbb{R}$ be the loss as a function of the model parameters w and data points $z = (x, y)$. Thus, the aim of the optimization is to minimize

$$\mathcal{L}_i(w) := \mathbb{E}_{z \sim \mathcal{D}_i}[\ell(w; z)] \quad \forall i = 1, \dots, k \quad (1)$$

In this work, we study the problem of covariate-shift. To do this, we create different distributions $\mathcal{D}^r(x)$ for each image dataset by rotating the images with r degrees. $\mathcal{D}^0(x)$ is defined as a dataset where images have been rotated with 0 degrees of rotation, and $\mathcal{D}^{180}(x)$ with 180 degrees of rotation.

We perform experiments on two and four different data distributions, where $r \in \{0, 180\}$ or $r \in \{0, 90, 180, 270\}$. The train and test sets for the studied datasets are randomly split into one equally large partition for each value of r , and a rotation of r is applied on each partition thus creating $\mathcal{D}^r(x)$. Then each client is populated with training samples uniformly from one such data distribution. Since the labels are unchanged after rotation, this means that the marginal distributions over input features differ between these groups, but the conditional distributions $\mathcal{D}(y|x)$ are the same for all clients. This way of creating different client distributions has previously been used in (Ghosh et al., 2020) for centralized federated learning.

The main challenge of this paper is that we assume the client distributions \mathcal{D}_i^r are *unknown for each client i* , and our goal is to design an algorithm that can both identify \mathcal{D}_i^r and at the same time perform distributed optimization.

4. Algorithm

In our decentralized peer-to-peer network, we use the gossip protocol for communication between clients. Below we describe the random gossip baseline, and our proposed extension PENS.

4.1. Random gossip communication

In this framework, each client starts with a randomly initialized model that is updated using stochastic gradient descent (SGD) on the local client data for E local epochs. The model parameters w_i of client i are then at a random time communicated to a randomly chosen neighboring client j in the network. This action is denoted as **Send** $_{i \rightarrow j}(\cdot)$. Client j then waits for n_{peers} number of models before it aggregates its own current model with the received ones with a simple average: $\bar{w}_j^{t+1} \leftarrow \frac{1}{n} \sum_i w_i^t$. This is the same aggregation method as commonly used in centralized federated learning. The new aggregated model \bar{w}_j^{t+1} is then trained for E epochs before it is ready to be gossiped again. A summary of the gossip learning protocol is presented in algorithm 1.

4.2. PENS: Performance-based neighbor selection

A problem that arises with random gossip when we have different distributions $\mathcal{D}_i^r(x)$ for clients i is that if two clients with dissimilar distributions (i.e. $r = 0$ for client i and $r = 180$ for client j) communicate, the performance of the learned model is usually negatively effected when their models are aggregated.

To solve this problem, we introduce PENS: Performance-based neighbor selection. PENS consists of two main steps. In the first step, the algorithm finds clients of similar marginal distributions $\mathcal{D}_i(x)$ to communicate with. In the second step, the random gossip protocol is followed for the subset of clients selected from the first step. *Clients of similar distributions are found by evaluating sent models on the receiving client's training data.* The main idea of the proposed method is that the training loss of a sent client model w_i is expected to be lower on the training set of a receiving client j that has a similar data distribution, and a higher loss for those clients that have dissimilar distributions.

First, each client communicates randomly in the network for a pre-defined number of neighbor selection communication rounds T . At a random time, a client performs the **Send** $_{i \rightarrow j}(\cdot)$ operation, after which $\mathcal{L}_j(w_i; z)$ is calculated, where $z \sim \mathcal{D}_j$. This is the loss of client model w_i on the training set of client j . Each client waits for n_{sampled} number of models and saves a list of the losses. The top m best performing (lowest loss) clients are selected as *potential neighbors with similar data distributions*. Then their model parameters are aggregated into a new model. This is repeated for T rounds, after which the clients that were

selected more than the expected amount of times (if the sampling of clients would have been uniform) are identified as *neighbors with similar data distributions*.

A set of neighbors with similar data distributions are now identified for every client and this constitutes step 1 of PENS. This is summarized in algorithm 2. In step 2 of PENS, the gossip learning protocol (algorithm 1) is used for the set of selected neighbors for each client.

Algorithm 1 Gossip learning protocol

```

1: function MAIN
2:   while stopping criterion not met do
3:     WAIT ( $\Delta$ )
4:      $j \leftarrow \text{RANDOMPEER}()$  // select random peer
5:     SEND $_{i \rightarrow j}(w_i, j)$ 
6:   end while
7: end function
8: function ONRECEIVEMODEL( $w_j$ )
9:   SAVE( $w_j$ )
10:  if no. of received models  $\geq n_{\text{peers}}$  then
11:     $w_j \leftarrow \text{MERGE\_SAVED\_MODELS}()$ 
12:     $w_j \leftarrow \text{TRAIN}(x; w_j)$  //update on local data  $x$ 
13:  end if
14: end function
    
```

Algorithm 2 PENS step 1: find peers with similar data distributions $\mathcal{D}_i^r(x)$

```

1: function MAIN
2:   while stopping criterion not met do
3:     WAIT ( $\Delta$ )
4:      $j \leftarrow \text{RANDOMPEER}()$  // select random peer
5:     SEND $_{i \rightarrow j}(w_i, j)$ 
6:   end while
7:   SELECTNEIGHBORS()
8: end function
9: function ONRECEIVEMODEL( $w_i$ )
10:   $\ell_i \leftarrow \text{CALCULATELOSS}(w_i)$ 
11:  SAVE( $w_i, \ell_i$ )
12:  if no. of received models  $\geq n_{\text{sampled}}$  then
13:     $w_j \leftarrow \text{MERGE}(\text{SELECT\_TOP\_M}(w_i, \ell_i))$ 
14:     $w_j \leftarrow \text{TRAIN}(x; w_j)$  //update on local data  $x$ 
15:    no. of received models  $\leftarrow 0$  //reset
16:  end if
17: end function
18: function SELECTNEIGHBORS
19:  for all peers  $i$  do
20:    if merged with  $i$  more than expected then
21:      NEIGHBORLIST.append( $i$ )
22:    end if
23:  end for
24:  return: NEIGHBORLIST
25: end function
    
```

5. Experimental setup

In this work, we set out to develop an algorithm to solve the problem of non-iid data for gradient-based algorithms in a peer-to-peer network. To do this, we limit the experiments to a peer-to-peer network that is fully connected (i.e. all nodes in the network can communicate with each other). Further, we assume that all clients are able to communicate at any time. We simulated the peer-to-peer network on a computer, and all experiments were performed with a NVIDIA Tesla V100-SXM2-32GB GPU. Our code is available at [github](https://github.com/guskarls/dfl-pens)¹.

5.1. Datasets

Our experiments are performed on two datasets for visual classification, CIFAR-10 (Krizhevsky et al., 2009) and Fashion-MNIST (Xiao et al., 2017). The CIFAR-10 dataset consists of 60 000 32x32 color images in 10 classes, with 6000 images per class. The dataset is split into 50 000 training images and 10 000 test images. The Fashion-MNIST dataset contains 70 000 28x28 gray-scale images of Zalando clothing in 10 classes. It is split into 60 000 training images and 10 000 test images.

5.2. Models and hyperparameters

The CNN used in our experiments consists of three convolutional layers and ReLU activations (with 32 channels in the first layer and 64 channels in the last two layers), each with a kernel size of 3 and each followed by max pooling. This is followed by one fully connected layer of 64 units with a ReLU activation and an output layer with a softmax activation. The size and architecture of this network is not state-of-the-art for visual classification tasks but has sufficient capacity for the comparison that we perform in our experiments. We use SGD as our optimizer with a learning rate $\eta = 10^{-3}$ for PENS. For the random gossip protocol we use $n_{\text{peers}} = 20$ and for PENS $n_{\text{sampled}} = 10$ and $m = 2$, if not explicitly otherwise stated. All hyperparameters were tuned for all baselines and the best ones were chosen with respect to a local validation set on each client.

5.3. Baselines

We compare our proposed algorithm PENS with two baselines: random gossip (Random) and locally trained models without communication (Local) for each client. We also report results for an Oracle, that for each client is given the information of which neighbors have the same distributions $\mathcal{D}_i^r(x)$ for all $i = 1, \dots, k$, and only communicates with these neighbors. Accuracy for a centrally trained model is also presented, denoted by Central, where we train one model on all data non-distributed.

¹<https://github.com/guskarls/dfl-pens>

5.4. Evaluation

For testing of all algorithms, we measure test accuracy for each client i on test data from the client’s own distribution $\mathcal{D}_i^r(x)$. All reported accuracies are averaged over clients for each distribution. We run 4 experiments for each algorithm, with different random seeds, and report the average and a 95% confidence interval. For step 1 of PENS, we let clients communicate $T = 200$ communication rounds. During step 2 of PENS, Random and Oracle, we set $T = 333$ and we perform early stopping on local validation data for each client between client communication rounds. The communication is stopped when the validation loss has converged or when we reach T . The validation sets consist of 100 sample points from \mathcal{D}_i^r per client in all experiments.


6. Results and discussion

In table 1, results on CIFAR-10 are shown for $r = \{0, 180\}$. Accuracies are reported both for independent and common weight initialization for the client models. In centralized federated learning, it is known that a common initialization of client models is important for federated averaging to work (McMahan et al., 2017). Meanwhile, our results suggest that a common initialization is not necessary in order for the different algorithms to reach a high accuracy in decentralized federated learning. Further, the proposed method PENS achieves an accuracy that is higher than all baselines and close to the performance of Oracle with perfect information of client data distributions $\mathcal{D}_i^r(x)$. In table 2 we see that PENS outperforms the baselines also in the case of $r = \{0, 90, 180, 270\}$.


In table 3, accuracies for all algorithms on Fashion-MNIST with $r = \{0, 180\}$ are presented for 100 and 500 training samples per client, with the number of clients set to 100. Although the difference of test accuracy between the baselines is smaller as compared to CIFAR-10 (since Fashion-MNIST is an easier problem), our proposed method PENS outperforms both baselines in this setting as well.

6.1. Impact of training set size

In figure 1a we compare PENS to the baseline algorithms on the CIFAR-10 dataset in a setting where we fix the number of clients to 100 while at the same time varying the size of the local train sets. The results show that by increasing the size of the local train set on each client, performance increases for all compared algorithms. We further see that PENS consistently outperforms both baselines. In a low data setting, with 100 training samples per client, PENS is closer to Random in performance as compared to Oracle. However, when the training size increases, the difference between Random and PENS increases, as PENS manages to find the correct neighbors with similar data distributions



(a) Number of clients fixed to 100. Training set size per client varying.



(b) Train set size per client fixed to 150. Number of clients varying.

Figure 1. Test accuracy on CIFAR-10 with $r = \{0, 180\}$ as a function of (a) training samples per client and (b) number of clients, while fixing the other. Oracle has perfect information of client distributions, as opposed to the other methods.

for each client. This is further visualized in figure 2, where a heatmap over the communication pattern is plotted for Oracle, PENS, and Random. Here we see that for each client, PENS manages to find almost all clients with similar data distributions.

In table 5 we report for precision (the fraction of clients with the same distribution \mathcal{D}^r among the selected clients) and recall (the fraction of clients with the same distribution \mathcal{D}^r) for the peers that PENS selects for each client. Here we report experiments for 100 clients for a different number of training samples per client. We note that the precision of our method is very high and robust to the number of training samples per client. The recall is lower than the precision, but also robust to the number of training samples.

6.2. Impact of number of clients

Figure 1b shows results for experiments for CIFAR-10 where the number of samples per client is fixed to 250 (150 train and 100 validation samples), but with a varying number of clients in the peer-to-peer network. Since no communication is allowed for the local baseline, the performance is constant with respect to the number of clients. Meanwhile, for the other methods, we see that by adding more clients (and therefore also the total amount of data in the system) the performance increases. Further, our proposed method PENS consistently outperforms the random baseline.

6.3. Impact of n_{sampled} and top m performers

The parameter n_{sampled} in step 1 of PENS decides for each client how many other client models to sample at every

Table 1. Test accuracy reported on CIFAR-10. Independent and common model weight initialization. 100 clients, 400 training samples.

Method	Acc. (independent) %	Acc. (common) %
Central	65.5 ± 0.3	65.5 ± 0.3
Oracle	58.1 ± 0.4	57.7 ± 0.6
PENS	57.8 ± 0.4	57.2 ± 1.3
Random	54.7 ± 0.8	54.1 ± 0.8
Local	33.4 ± 0.2	33.3 ± 0.7

communication round, and m decides how many of the top-performing (lowest loss) to merge with. Experiments were carried out to study how sensitive PENS is to the choice of these hyperparameters. In table 4 we summarize the results for varying values of these hyperparameters in the setting of 100 clients and 400 training samples per client for CIFAR-10. We note that the test accuracies are relatively stable for different values of n_{sampled} and m . Further, our results suggest that the ratio $\frac{n_{\text{sampled}}}{m}$ should not be too large, i.e. if n_{sampled} is increased m should be higher as well. We have noticed in our experiments that if m is set too low relative to n_{sampled} , PENS will collapse into always choosing the same few clients and miss to find other peers of the same distribution.

7. Future work

There are several interesting research directions left to explore which we had to limit ourselves from including in this paper. First, we assumed that all clients were able to communicate equally fast and at all times. This is a strong assumption that is not always true in many real world applications,




Figure 2. Heatmaps over the communication pattern between clients on CIFAR-10 with two rotated \mathcal{D}^r for (a) Oracle (b) PENS and (c) random. The reported values are normalized counts: a value close to 1 means that client j received a model from client i frequently. The clients are sorted so that clients with ID 0-99 have data distributions $\mathcal{D}^0(x)$ and clients with ID 100-199 have $\mathcal{D}^{180}(x)$.

Table 2. Test accuracies on CIFAR-10 with 4 rotated distributions $\mathcal{D}^r(x)$, $r = \{0, 90, 180, 270\}$. 100 clients and 400 training samples per client.

Method	Accuracy (%)
Central	60.3 ± 0.8
Oracle	53.2 ± 1.1
PENS	53.4 ± 0.7
Random	49.3 ± 1.0
Local	33.4 ± 0.7

Table 3. Test accuracy reported on Fashion-MNIST for 100 clients with 100 and 500 training samples per client and $r = \{0, 180\}$.

Method	100	500
Central	85.3 ± 0.3	87.3 ± 0.3
Oracle	79.4 ± 0.4	84.9 ± 0.2
PENS	78.8 ± 0.4	84.7 ± 0.3
Random	77.5 ± 0.5	83.7 ± 0.2
Local	63.9 ± 0.6	74.4 ± 0.5

and it would therefore be interesting to study **system heterogeneity** in a decentralized network where clients have different hardware and computational budgets. Second, we assumed that the decentralized **network topology** was fully connected. As future work, it would be interesting to study how PENS performs on other types of network topologies and how these effect the learning among clients.

8. Conclusions

In this work we have studied non-convex optimization of decentralized federated learning using deep neural networks in a non-iid data setting. Our experiments show that our proposed method PENS is able to efficiently aid clients to identify neighboring peers with similar data distributions

Table 4. Accuracies for varying n_{sampled} and m_{selected} on CIFAR-10 for 100 clients with 150 training samples per client, $r = \{0, 180\}$.

n_{sampled}	m	n_{sampled}/m	Accuracy (%)
20	6	3.3	48.4 ± 0.5
50	15	3.3	47.8 ± 1.0
10	2	5.0	47.7 ± 1.4
10	3	3.3	47.2 ± 1.1
50	10	5.0	46.9 ± 1.4
20	4	5.0	46.9 ± 0.8
10	5	2.0	46.4 ± 0.9
5	2	2.5	46.2 ± 0.8
10	1	10.0	46.1 ± 0.7
5	1	5.0	45.4 ± 0.7
20	10	2.0	45.4 ± 1.8
20	2	10.0	45.4 ± 2.0
5	3	1.7	45.2 ± 2.0
50	25	2.0	45.1 ± 2.5
50	5	10.0	44.5 ± 2.0

Table 5. Precision and recall for 100 clients for different number of training samples and $r = \{0, 180\}$.

Train set size	Precision (%)	Recall (%)
400	99.7 ± 0.1	70.8 ± 0.8
300	99.7 ± 0.1	71.1 ± 0.9
200	99.7 ± 0.1	70.1 ± 0.7
100	95.8 ± 1.5	67.9 ± 1.5

in a fully decentralized FL setting and in that way guide the learning of the clients to achieve high performance in a non-iid data setting. PENS works by using the training loss to find clients in the network that share similar data distributions, which then focus on communicating with each other instead of random neighbors. Our results (figures 1 and 2) suggest that given enough training data per client,

PENS will reach the same accuracy as an oracle that is given *perfect information* of the client data distributions. We have limited ourselves to the non-iid setting of covariate shift in this work. Meanwhile, we hypothesize that our proposed method also works well on other types of non-iid data, such as label distribution skew, concept shift or concept drift.

References

- Assran, M., Loizou, N., Ballas, N., and Rabbat, M. Stochastic gradient push for distributed deep learning. In *International Conference on Machine Learning*, pp. 344–353. PMLR, 2019.
- Blot, M., Picard, D., Cord, M., and Thome, N. Gossip training for deep learning. *arXiv preprint arXiv:1611.09726*, 2016.
- Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. Randomized gossip algorithms. *IEEE transactions on information theory*, 52(6):2508–2530, 2006.
- Deng, Y., Kamani, M. M., and Mahdavi, M. Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461*, 2020.
- Ghosh, A., Chung, J., Yin, D., and Ramchandran, K. An efficient framework for clustered federated learning. *arXiv preprint arXiv:2006.04088*, 2020.
- Giaretta, L. and Girdzijauskas, Š. Gossip learning: Off the beaten path. In *2019 IEEE International Conference on Big Data (Big Data)*, pp. 1117–1124. IEEE, 2019.
- Hegedűs, I., Danner, G., and Jelasity, M. Gossip learning as a decentralized alternative to federated learning. In *IFIP International Conference on Distributed Applications and Interoperable Systems*, pp. 74–90. Springer, 2019.
- Jelasity, M., Voulgaris, S., Guerraoui, R., Kermarrec, A.-M., and Van Steen, M. Gossip-based peer sampling. *ACM Transactions on Computer Systems (TOCS)*, 25(3):8–es, 2007.
- Jeong, E., Oh, S., Kim, H., Park, J., Bennis, M., and Kim, S.-L. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. *arXiv preprint arXiv:1811.11479*, 2018.
- Jiang, Y., Konečný, J., Rush, K., and Kannan, S. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*, 2019.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- Kempe, D., Dobra, A., and Gehrke, J. Gossip-based computation of aggregate information. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pp. 482–491. IEEE, 2003.
- Koloskova, A., Stich, S., and Jaggi, M. Decentralized stochastic optimization and gossip algorithms with compressed communication. In *International Conference on Machine Learning*, pp. 3478–3487. PMLR, 2019.
- Kong, L., Lin, T., Koloskova, A., Jaggi, M., and Stich, S. U. Consensus control for decentralized deep learning. *arXiv preprint arXiv:2102.04828*, 2021.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Listo Zec, E., Mogren, O., Martinsson, J., Sützfled, L. R., and Gillblad, D. Specialized federated learning using a mixture of experts. *arXiv preprint arXiv:2010.02056*, 2020.
- 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.
- Niwa, K., Harada, N., Zhang, G., and Kleijn, W. B. Edge-consensus learning: Deep learning on p2p networks with nonhomogeneous data. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 668–678, 2020.
- Ormándi, R., Hegedűs, I., and Jelasity, M. Gossip learning with linear models on fully distributed data. *Concurrency and Computation: Practice and Experience*, 25(4):556–571, 2013.
- Roy, A. G., Siddiqui, S., Pölsterl, S., Navab, N., and Wachinger, C. Braintorrent: A peer-to-peer environment for decentralized federated learning. *arXiv preprint arXiv:1905.06731*, 2019.
- Wang, J., Sahu, A. K., Yang, Z., Joshi, G., and Kar, S. Matcha: Speeding up decentralized sgd via matching decomposition sampling. In *2019 Sixth Indian Control Conference (ICC)*, pp. 299–300. IEEE, 2019a.
- Wang, K., Mathews, R., Kiddon, C., Eichner, H., Beaufays, F., and Ramage, D. Federated evaluation of on-device personalization. *arXiv preprint arXiv:1910.10252*, 2019b.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.