

Understanding Federated Learning from IID to Non-IID dataset: An Experimental Study

Jungwon Seo, Ferhat Ozgur Catak and Chunming Rong

University of Stavanger, Norway

{jungwon.seo, f.ozgur.catak, chunming.rong}@uis.no

Abstract. As privacy concerns and data regulations grow, federated learning (FL) has emerged as a promising approach for training machine learning models across decentralized data sources without sharing raw data. However, a significant challenge in FL is that client data are often non-IID (non-independent and identically distributed), leading to reduced performance compared to centralized learning. While many methods have been proposed to address this issue, their underlying mechanisms are often viewed from different perspectives. Through a comprehensive investigation from gradient descent to FL, and from IID to non-IID data settings, we find that inconsistencies in client loss landscapes primarily cause performance degradation in non-IID scenarios. From this understanding, we observe that existing methods can be grouped into two main strategies: (i) adjusting parameter update paths and (ii) modifying client loss landscapes. These findings offer a clear perspective on addressing non-IID challenges in FL and help guide future research in the field.

Keywords: Federated Learning · Gradient Descent · Optimization.

1 Introduction

Federated Learning (FL) is a distributed machine learning framework that prioritizes privacy preservation [14]. Unlike traditional centralized methods, where client data is aggregated on a central server for model training, FL allows clients to keep their data locally. Clients collaborate by sharing only their locally trained models with a central server, which aggregates these models to build a global model. This approach enables learning from diverse, distributed datasets without exposing data, making FL especially valuable for privacy-sensitive applications.

Although similar distributed learning frameworks have existed before [3,23,4], the term "Federated Learning" was officially coined by McMahan et al. in 2016 [14]. Since then, particularly in deep learning, FL has gained significant attention in both research and industry [6]. One of the primary challenges in FL is the heterogeneity of data across clients [22]. When client data distributions are non-IID (not Independent and Identically Distributed), the learning direction of individual clients may diverge from that of the global model. This divergence, often referred to as "*client drift*", can degrade the overall performance of the

global model and, in extreme cases, may even hinder convergence [7]. Numerous studies have addressed this issue, primarily from an optimization perspective, proposing both theoretical and empirical methods [13,7,1].

As the field has matured, a variety of new perspectives have emerged. Some researchers argue that FL’s challenges resemble those found in catastrophic forgetting, similar to continual learning [10,9], while others suggest that the smoothness of the loss function, inspired by sharpness-aware training, plays a crucial role [15,20]. Additionally, some point to differences in learned representations, relating to concepts from contrastive learning [12].

Initially, the FL problem was perceived as a straightforward optimization task. However, the incorporation of new perspectives has introduced additional layers of complexity, making it challenging to identify the critical areas of focus. Moreover, most existing studies primarily investigate FL under non-IID settings, leaving a gap in our understanding of how FL behaves from scratch. As a result, FL has become a "black box within a black box," exacerbating the opacity already present in deep learning models. This multitude of perspectives has made it difficult for researchers to interpret proposed methods and to determine the most promising research directions.

To address this, we propose a systematic investigation of FL, starting with IID settings and gradually exploring its behavior across various hyperparameters and conditions. By grounding our analysis in gradient descent and extending it to FL, we aim to uncover the mechanisms driving its performance and explain the unexpected phenomena observed in complex, non-IID environments. Through this approach, we hope to clarify FL’s core challenges and offer direction for future research.

The contributions of this study are as follows:

- **Optimization Algorithm Review:** We provide a clear review of optimization algorithms, from gradient descent to FedAvg.
- **Experiments in IID Settings:** We conduct extensive experiments on FedAvg under IID conditions, exploring various hyperparameters and FL settings to gain insights into its behavior.
- **Non-IID Analysis:** Using insights from the IID experiments, we analyze FL performance in non-IID settings, interpreting the observed phenomena.
- **Categorizing Approaches for Non-IID Challenges:** We classify existing and future methods into two strategies: (a) update path adjustment and (b) loss landscape modification, addressing non-IID data in FL.

2 From Gradient Descent to FedAvg

Gradient Descent (GD) is one of the most fundamental algorithms in machine learning for optimizing model parameters, such as those in neural networks [17]. Consider an objective function $F(\theta)$ that we seek to minimize, where θ represents the model parameters, with $\theta \in \mathbb{R}^d$, and d is the dimensionality of θ . The parameter update rule in GD is given by:

$$\theta^{t+1} = \theta^t - \eta \nabla F(\theta^t; \mathcal{D}) \quad (1)$$

Here, t denotes the iteration step, and $\nabla F(\theta^t; \mathcal{D})$ represents the gradient of the objective function with respect to θ^t , evaluated on the dataset \mathcal{D} . By updating the parameters in the negative direction of the gradient with a learning rate η , the parameters move towards a more optimal solution. This iterative process is repeated until the gradient $\nabla F(\theta^t)$ converges to zero, indicating that the model has reached a local or global minimum.

A major limitation of GD lies in its dependence on a static dataset \mathcal{D} . This becomes problematic when the dataset is too large to fit into a single device’s memory or when full-batch training becomes computationally expensive. GD requires computing the gradient over all data points and using the averaged gradient for each update at iteration t . Expanding the update rule in Eq. 1, we obtain:

$$\theta^{t+1} = \theta^t - \eta \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \nabla F(\theta^t; x_i, y_i) \quad (2)$$

Here, x_i and y_i represent the feature and label pairs from the dataset \mathcal{D} in a supervised learning scenario. Eq. 2 highlights that, before each update, the gradient must be computed and averaged across all data points. Each data point contributes a distinct gradient, and the average of these gradients determines the final update direction. This gradient averaging is fundamental to optimization and closely relates to the parameter averaging concept in Federated Averaging (FedAvg). FedAvg can thus be viewed as a generalization of traditional optimization methods, a connection that will be elaborated upon.

To address GD’s inefficiencies, SGD was introduced. Unlike GD, which updates after computing the gradient over the entire dataset, SGD updates the parameters after each individual data sample. This leads to faster updates and reduced memory usage, making it suitable for large-scale problems. However, this method can be unstable, as computing gradients from a single data point may introduce bias and slow convergence in terms of model performance, despite the higher computational efficiency.

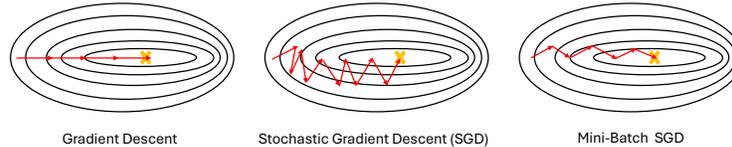


Fig. 1: Illustration of how different GD methods navigate the loss landscape

To balance the trade-offs between GD and SGD, mini-batch SGD was developed. It computes gradients using a batch of size B for each update, reducing SGD’s instability while improving GD’s efficiency. GD, SGD, and mini-batch SGD can be seen as special cases within a broader framework, where the batch size B ranges from 1 (SGD) to the entire dataset (GD), as shown in Algorithm 1, with their respective optimization paths illustrated in Fig. 1.

Algorithm 1 Stochastic Gradient Descent (SGD)

```

1: Input: Learning rate  $\eta$ , Batch size  $B$ , Number of epochs  $E$ 
2: for  $n = 1$  to  $E$  do
3:   for each random mini-batch  $\mathcal{B}$  of size  $B$  with data  $(x_i, y_i)$  do
4:     Compute gradient  $g = \frac{1}{B} \sum_{i \in \mathcal{B}} \nabla_{\theta} F(\theta; x_i, y_i)$ 
5:     Update parameters:  $\theta \leftarrow \theta - \eta \cdot g$ 

```

The concept of gradient averaging for training extends naturally to distributed and parallel computing, as gradient calculations are independent until aggregation. Scaling SGD process across multiple devices (e.g., GPUs or distributed systems) accelerates training and mitigates the limitations of centralized hardware. While there are various approaches to parallel SGD [3,23,4], in this work, we focus on aggregating gradients from each device per batch of each client, as described in Algorithm 2. In this distributed learning approach, we introduce a parameter, K , representing the number of devices. Gradients are computed on each device, followed by averaging locally and then globally.

Algorithm 2 Parallel Stochastic Gradient Descent (Parallel SGD)

```

1: Input:  $\eta$ ,  $B$ ,  $E$ , Number of devices  $K$ 
2: for  $n = 1$  to  $E$  do
3:   for each device  $k = 1$  to  $K$  in parallel do
4:     Uniformly select a random mini-batch  $B_k$  from  $\mathcal{D}_k$ 
5:     Compute gradient on device  $k$ :  $g_k = \frac{1}{|B_k|} \sum_{i \in B_k} \nabla_{\theta} \ell(\theta; x_i, y_i)$ 
6:   Aggregate gradients:  $g = \frac{1}{K} \sum_{k=1}^K g_k$ 
7:   Update parameters:  $\theta \leftarrow \theta - \eta \cdot g$ 

```

Despite the potential benefits of distributed and parallel processing, communication bottlenecks can arise, particularly during gradient aggregation. While intra-device parallelism (e.g., multi-GPU operations) tends to be efficient, frequent gradient communication between devices introduces overhead, offsetting the advantages of parallelization [18].

A natural approach to mitigate this issue is to reduce the communication frequency. Several studies have investigated methods where each device updates parameters independently for a set duration, merging gradients less frequently. In LocalSGD [19], each device performs training for I local epochs, followed by gradient merging every $I + 1$ epochs. The globally updated model is then used as the starting point for subsequent training, as described in Algorithm 3.

Algorithm 3 Local Stochastic Gradient Descent (Local SGD)

```

1: Input:  $\eta$ ,  $B$ ,  $E$ ,  $K$ , Synchronization period  $I$ 
2: for  $n = 1$  to  $E$  do
3:   for each device  $k = 1$  to  $K$  in parallel do
4:     Initialize local parameters:  $\theta_k \leftarrow \theta$ 
5:     for  $t = 1$  to  $I$  do
6:       Uniformly select a random mini-batch  $B_k^t$  from  $\mathcal{D}_k$ 
7:       Compute gradient on device  $k$ :  $g_k = \frac{1}{|B_k^t|} \sum_{i \in B_k^t} \nabla_{\theta} \ell(\theta_k; x_i, y_i)$ 
8:     Update local parameters:  $\theta_k \leftarrow \theta_k - \eta \cdot g_k$ 
9:   Synchronize: Aggregate parameters across devices:  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^K \theta_k$ 

```

LocalSGD can be viewed as a special case of FedAvg, more aligned with distributed learning. It uses centrally managed, evenly partitioned IID data across all devices, reducing gradient variance and closely approximating centralized training performance. While sharing FedAvg’s core mechanics, LocalSGD assumes balanced IID data, making it more suited for distributed learning.

In contrast, FedAvg (Algorithm. 4) assumes partial participation of client (Line 3) and non-IID data across clients, requiring adjustments in gradient aggregation. Client updates are weighted by dataset size, with larger datasets contributing more to the global model, as noted in line 12.

Algorithm 4 Federated Averaging (FedAvg)

```

1: Input: Global  $\eta_g$ , Local  $\eta_l$ ,  $B$ ,  $E$ ,  $K$ , Communication rounds  $R$ , Fraction of clients  $C$ 
2: for each round  $r = 1$  to  $R$  do
3:   Randomly select a subset  $\mathcal{S}$  of  $\max(C \cdot K, 1)$  clients
4:   for each client  $k \in \mathcal{S}$  in parallel do
5:     Initialize local model:  $\theta_k \leftarrow \theta$ 
6:     for each local epoch  $n = 1$  to  $E$  do
7:       for each mini-batch  $B_k$  from local dataset  $\mathcal{D}_k$  do
8:         Compute gradient:  $g_k = \frac{1}{|B_k|} \sum_{(x_i, y_i) \in B_k} \nabla_{\theta} \ell(\theta_k; x_i, y_i)$ 
9:         Update local model:  $\theta_k \leftarrow \theta_k - \eta_l \cdot g_k$ 
10:        Send back to server: (Option I)  $\Delta\theta_k = \theta_k - \theta$  or (Option II)  $\theta_k$ 
11:       Aggregate the total number of data points:  $M = \sum_{k \in \mathcal{S}} |\mathcal{D}_k|$ 
12:       Update global model: (I)  $\theta \leftarrow \theta + \eta_g \cdot \sum_{k \in \mathcal{S}} \frac{|\mathcal{D}_k|}{M} \cdot \Delta\theta_k$  or (II)  $\theta \leftarrow \sum_{k \in \mathcal{S}} \frac{|\mathcal{D}_k|}{M} \cdot \theta_k$ 

```

Originally, FedAvg exchanges local model parameters (Option II) instead of updates (Option I), which are mathematically equivalent. However, sending updates allows for more flexibility in modifying the aggregation process, enabling techniques such as global learning rate adjustments and optimizers [7,16]. Moreover, this formulation provides a clearer conceptual framework for understanding FL. From the perspective of Option I in line 12, the updates from each client resemble gradients from individual data samples in GD. Therefore, FL is not merely parameter mixing but remains a structured optimization process. The following sections provide an empirical analysis of how FL works in practice.

3 Experimental Setup

3.1 Basic Configuration

We conduct our experiments using the CIFAR-10 dataset (50,000 instances in the training set and 10,000 in the test set), a standard benchmark in FL research. The model architecture is a convolutional neural network (CNN) with two convolutional layers followed by ReLU activations and max-pooling. It includes two fully connected layers, leading to a 10-class output, as described in [14]. The task is a classification problem, with cross-entropy loss serving as the objective function. To ensure a controlled experimental environment, we deliberately exclude additional techniques such as weight decay, data augmentation, and advanced optimizers (e.g., momentum or adaptive techniques). This approach allows us to isolate the impact of the hyperparameters B , E , and η on the SGD process.

3.2 FL Setup

Evaluation We evaluate model performance by analyzing both convergence speed and generalization quality. While rapid convergence is desirable, it does not inherently guarantee strong generalization, making it crucial to track both aspects throughout training. In our setup, global model accuracy (Top-1) is measured server-side using a dedicated test set. Test loss is also computed server-side to assess generalization, while training loss is computed locally by clients. The server averages client-side training losses to monitor overall optimization progress.

Transition from Centralized to Federated Learning To set initial hyperparameters, we first establish reasonable baseline performance in a centralized learning (CL) setup, corresponding to $K = 1$. We use $B = 500$, $\eta = 0.005$, and $E = 1$, yielding a final accuracy of around 68%. These hyperparameters are not fine-tuned for peak performance in CL; rather, the aim is to observe how performance evolves as we transition from CL to FL.

When K changes, the dataset is proportionally divided among K clients. For $K = 1$ (CL), a single client holds the entire training dataset of 50,000 samples, while for $K = 10$, each client is allocated 5,000 unique training samples. We initialize with a balanced IID dataset, ensuring that all clients have an equal distribution of labels and an identical number of samples. The effect of varying the number of clients on FL training dynamics is shown in Fig. 2.

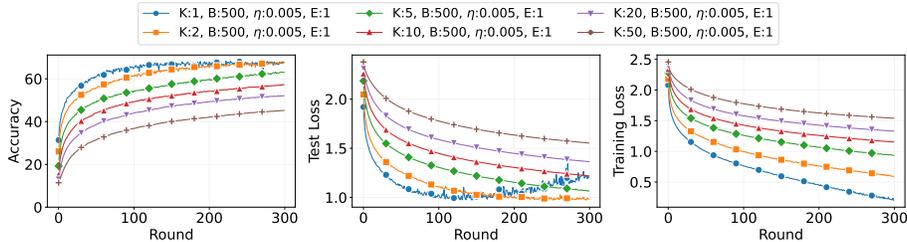


Fig. 2: Impact of varying K with consistent hyperparameters.

In this setup, as K increases, there is a noticeable decline in performance within a fixed number of communication rounds, highlighting a common oversight. When transitioning from CL to FL with a limited dataset, the critical factor that must be consistently controlled for a fair comparison is the **amount** of effective updates per round, u , rather than the **number**, due to the influence of η . It is defined as:

$$u = \eta \cdot \frac{E \cdot |\mathcal{D}|}{B \cdot K} \quad (3)$$

Therefore, local epoch E , batch size B , and learning rate η must be adjusted based on each client’s dataset size $\frac{|\mathcal{D}|}{K}$ to ensure all scenarios achieve an equal u .

In Fig. 3, we present three accuracy charts corresponding to adjustments in B , E , and η . Unlike the previous fixed hyperparameter setting, these adjustments

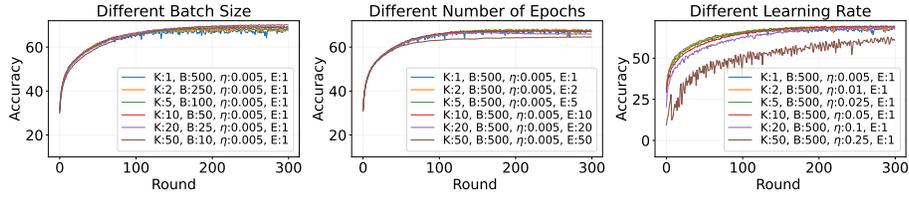


Fig. 3: Impact of matched B , E , and η on performance with different K .

align performance and convergence rates across different client counts. When E is increased to 50 and η is set to 0.25, it results in an extreme case, which we disregard in our analysis. Unless specified otherwise, the base setting for all experiments is $K = 10$, $E = 1$, $B = 50$, and $\eta = 0.005$, with variations applied to a single parameter in each experiment.

4 Experimental Results

4.1 Impact of Hyperparameters under IID Condition

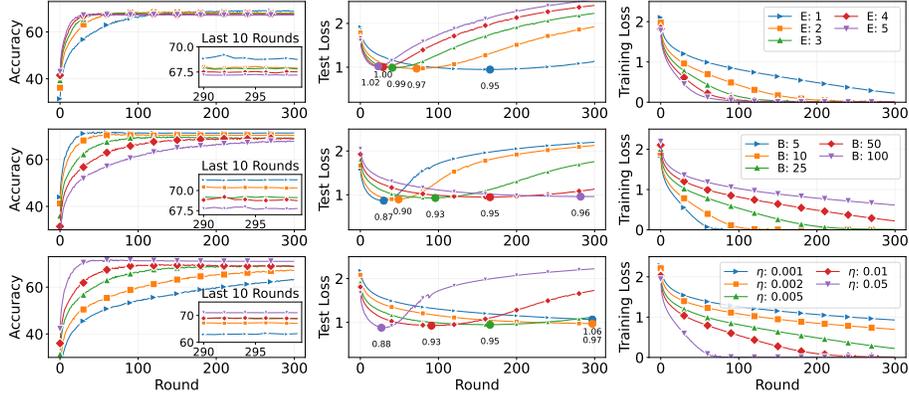


Fig. 4: Training dynamics when $K = 10$. The first row shows results for varying E , the second for different B , and the third for various η .

We first assess how three hyperparameters— E , η , and B —affect test accuracy, test loss, and training loss. As illustrated in Fig. 4, increasing E or η , or decreasing B , speeds up convergence. This is evident from the curvature of the accuracy plots and the more rapid decline in training loss, which is consistent with our update amount equation u (Eq. 3). Notably, the point of overfitting occurs earlier with faster convergence, as marked by the minimum test loss (indicated by the circle).

However, some notable observations can be made. While higher E values result in faster convergence, they lead to suboptimal minimum test loss and

final accuracy, with lower epochs (e.g., $E = 1$) providing better performance (first row in Fig. 4). In contrast, the smallest batch size ($B = 5$) achieves the highest final accuracy and the lowest test loss (second row in Fig. 4). Similarly, higher η produce the best final accuracy and the lowest test loss (third row in Fig. 4) while both of them provide faster convergence. Although the differences are not substantial under IID conditions, they become more pronounced in non-IID settings, which we will discuss later.

4.2 Partial Participation and Imbalanced Data under IID Condition

This section explores the effects of partial client participation and imbalanced data in the context of FL under IID conditions. While these are not hyperparameters, they are key aspects of the experimental setup. In FL, the server cannot always ensure that all clients participate in every training round. As a result, the server either waits for all clients or proceeds with only those that respond in time, leading to partial participation (PP), where only a subset of clients contribute to model updates in each round.

Even under IID conditions, where each client has an equal proportion of labels, the number of data samples per client can vary, creating imbalanced datasets. Some clients may have significantly more data than others.

We now examine how partial participation and data imbalance influence training dynamics, particularly with regard to convergence and model performance.

Impact of Partial Participation To evaluate the impact of partial participation (PP), we vary the number of clients involved in each communication round, randomly selecting 1, 2, 5, or 10 out of a total of 10 clients. While one might assume that involving more clients per round would result in better performance, our experiments under IID conditions reveal minimal performance differences across varying levels of participation, as shown in Fig. 5.

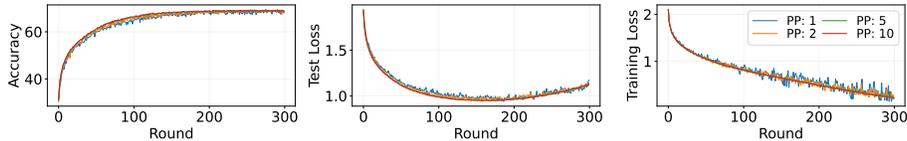


Fig. 5: Training Dynamics under Partial Participation with IID Condition

This result suggests that in the IID setting, an update from a single client remains valid for others, allowing them to resume training locally from that state without significant performance degradation. Although fewer clients participating per round introduce more fluctuations, the final accuracy still converges to nearly the same level.

Impact of Imbalanced Data. In our experiments with imbalanced datasets, we construct the environment using a Standard Gaussian Mixture (SGM) model. Here, an SGM value of 0 corresponds to a balanced dataset, while higher SGM values indicate increasing levels of data imbalance, as illustrated in Fig. 6.

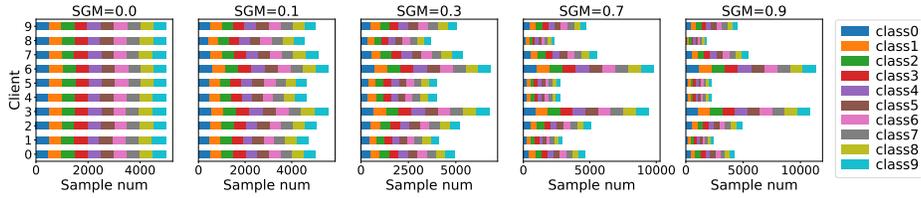


Fig. 6: Data partitioning with IID labels across different levels of SGM

For imbalanced datasets, we explore two aggregation strategies: **weighted aggregation**, used in FedAvg, where the contribution of each client is proportional to its dataset size, and **naive aggregation**, where updates from all clients are averaged equally, regardless of dataset size.

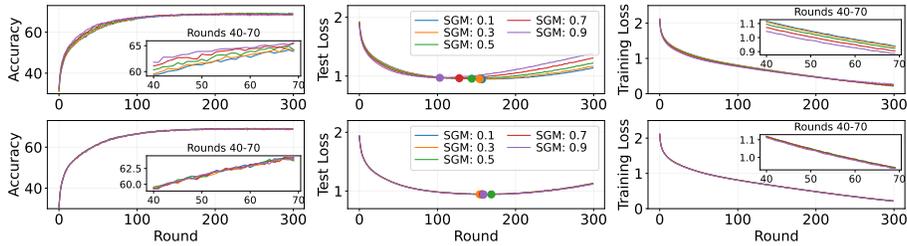


Fig. 7: Performance under varying levels of data imbalance. The first row shows results using weighted aggregation, while the second row uses naive averaging.

Contrary to intuitive expectations that increased imbalance would degrade performance, our experiments reveal that **weighted aggregation** actually accelerates convergence in IID settings. As shown in the first row of Fig. 7, models converge faster as imbalance increases, especially for SGM 0.9 (purple), which shows higher accuracy, earlier overfitting, and lower training loss between rounds 40-70. This is because, in weighted aggregation, clients with larger datasets exert more influence, effectively increasing the total number of updates. For instance, when Client A has 10 data points and Client B has 90, the weighted average of updates is $\frac{10 \times 10 + 90 \times 90}{100} = 82$. In contrast, with balanced datasets (50 data points each), the average number of updates is $\frac{50 \times 50 + 50 \times 50}{100} = 50$. Thus, in imbalanced scenarios, the total number of updates per round is higher, leading to faster convergence. Furthermore, when data is IID, the imbalance does not negatively impact performance—in fact, it can be advantageous.

We also conducted experiments using **naive aggregation**, which showed minimal performance differences across varying levels of imbalance, as shown in the second row of Fig. 7. This outcome can be understood by considering the number of updates: whether the dataset sizes are imbalanced (10 and 90) or balanced (50 and 50), naive averaging still results in an equal average of 50 updates per round. Interpreting updates as vectors, we can infer that when data is IID, the update vectors are aligned in direction, meaning that as long

as the magnitudes are adjusted appropriately, the combined update remains the same. This suggests that in IID settings, updates from individual clients are consistently aligned in direction.

4.3 Training Dynamics Under Non-IID Setting

Non-IID with Dirichlet Distribution Next, we experiment non-IID label distribution with a Dirichlet distribution for data partitioning across clients, which introduces varying degrees of non-IID behavior. In this setup, the degree of non-IID distribution is controlled by the Dirichlet α : a lower α leads to more distinct distributions across clients, while a higher α results in more similar distributions. The resulting distributions are visualized in Fig. 8.

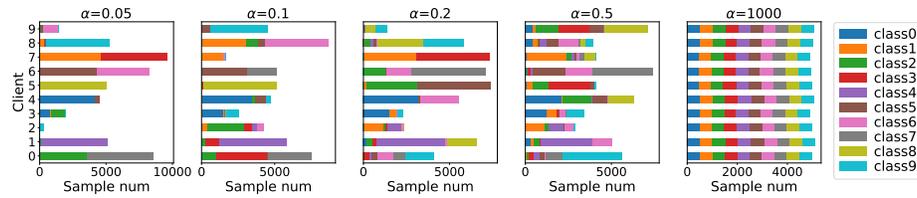


Fig. 8: Data partitioning with varying Dirichlet parameter α .

In Fig. 9, each line represents the average of three independent runs, with the shaded region indicating the min-max range across these runs. This approach accounts for the randomness introduced by the non-IID distribution of the dataset. Additionally, in the case of PP, the selection of clients in each round can have a significant impact on the results.

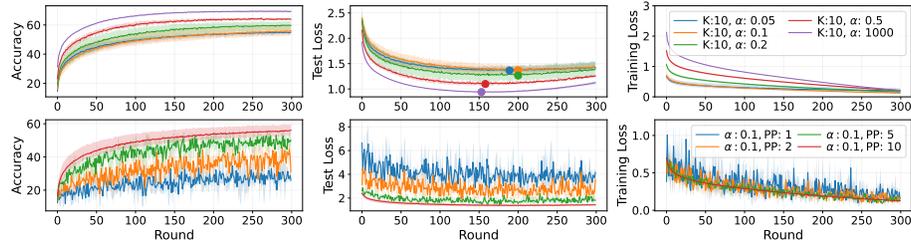


Fig. 9: Performance across varying α (Top) and PP (Bottom).

We observe that as the data distribution becomes more non-IID (i.e., as α decreases), model performance degrades (from 69% to 55%) and convergence slows. While lower α values do eventually converge within the limited communication rounds, the resulting solutions are highly suboptimal. Furthermore, unlike in IID settings, the level of PP plays a critical role in FL performance. When α is fixed at 0.1, lower levels of PP lead to significantly lower final accuracy, higher test loss, and increased fluctuations compared to full participation. This indicates that updates from different clients are not valid for others.

Impact of Hyperparameters under Non-IID Condition We investigate the effects of varying hyperparameters, as shown in Fig. 10.

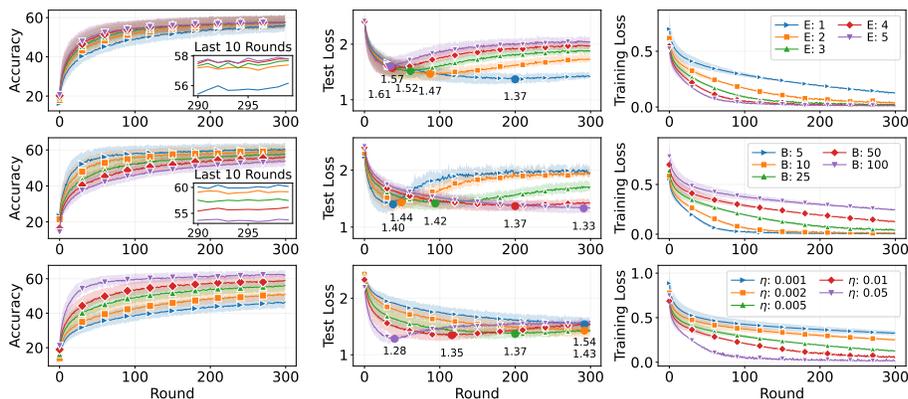


Fig. 10: Test accuracy, test loss, and training loss for $K = 10$. The first row shows results for varying E , the second for different B , and the third for various η .

First, increasing E results in faster convergence, and unlike the IID setting, a higher number of epochs leads to improved final performance. The difference in accuracy between the smallest and largest E is +1.7% (from 56.16% to 57.86%), whereas in the IID setting, this difference is -1.55% (from 68.75% to 67.2%). Second, lower values of B result in both higher final accuracy and faster convergence. The accuracy gap between the smallest and largest B is +6.45% (from 53.74% to 60.19%), compared to +3.57% (from 67.77% to 71.34%) in the IID case. Finally, η plays a crucial role in the non-IID scenario. The accuracy difference between the smallest and largest η is +15.48% (from 46.43% to 61.91%), compared to +7.79% (from 63.22% to 71.01%) in the IID setting. Notably, lowest test loss decreases as η increases.

5 Understanding FL from a Loss Landscape Perspective

We have performed a series of experiments across datasets ranging from IID to non-IID, incorporating both partial participation and imbalanced data. As expected, the performance of the global model declines as the level of non-IIDness increases. The most commonly cited explanation for this behavior is client drift. *But what drives this phenomenon?*

To uncover the underlying causes, we examine the problem through the lens of the loss landscape [11]. The loss landscape is shaped by the interaction between the loss function, dataset, and model parameters. It provides valuable insights into how the loss changes as model parameters are updated for a fixed dataset.

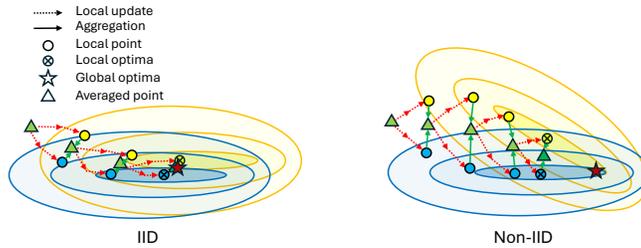


Fig. 11: Toy example illustrating the loss landscape and update paths for two clients with IID and Non-IID datasets.

In FL, each client operates on a distinct dataset, meaning each experiences a unique loss landscape [2,21]. Even though all clients use the same model and loss function, the variations in datasets lead to different optimization paths. As a result, updates from different clients are often misaligned, causing client drift. As each client optimizes locally, their models may converge to different local optima. When these locally optimized models are aggregated, the result may fail to reach the global optimum, as illustrated in Fig. 11.

In contrast, under IID conditions, these differences are minimized as the clients’ loss landscapes are more similar, leading to better-aligned update directions. Since directly visualizing the loss landscape in FL is difficult due to the high dimensionality and variability in client data, we instead performed two additional experiments to substantiate this observation.

First, we selected 5 clients from both IID and non-IID settings when K is 10 and traced their individual training loss. As shown in Fig. 12, under non-IID conditions, individual loss values decrease at different rates, while in the IID setting, the loss values for each client decrease more uniformly. This suggests that the local optimization processes in non-IID settings follow different loss landscapes, leading to varying loss values across clients.

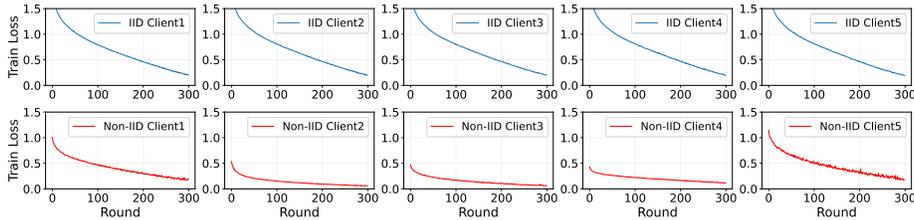


Fig. 12: Client training loss: Top row - IID, Bottom row - Non-IID.

Secondly, in Fig. 13, we measure the layer-wise (conv1, conv2, fc1, and fc2) cosine similarity of local updates ($\Delta\theta$) between pairs of clients, averaged across all combinations from 10 clients. As expected, in the IID setting, the updates across layers are well aligned (with cosine values close to 0), while in the non-IID setting, there is consistent misalignment between updates.

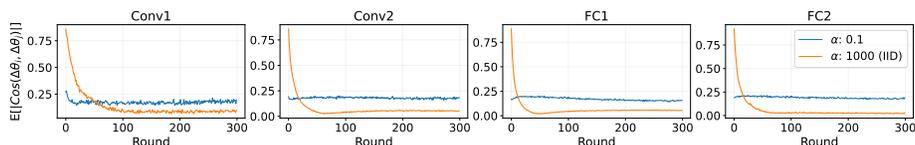


Fig. 13: Cosine similarity of local updates from different clients for each layer

These findings show that non-IID conditions lead to more diverse loss landscapes and less aligned client updates, while IID settings promote better alignment through consistent loss landscapes.

To guide the global model toward optimal performance in non-IID settings, two main strategies emerge:

1. **Adjusting the update path**, where optimization steps are tailored to direct the model toward the global optimum without altering the underlying loss landscape.
2. **Modifying the loss landscape**, which aims to increase the overlap between local optima of individual clients, thus promoting more consistent convergence across clients.

For example, methods like SCAFFOLD [7], which employs control variates, and FedOpt [16], which uses server-side optimizers (e.g., Adam, Yogi) or weight perturbation [15,9], focus on refining update directions. In contrast, strategies that modify the loss landscape add terms to the objective function [13,12,10]. Hybrid approaches like FedDyn [5] address both aspects simultaneously. However, a deeper analysis of these methods is necessary to fully understand their impact.

6 Discussion and Conclusion

Our experiments consistently show that a higher η and lower B improve performance in both IID and non-IID settings. According to our categorization, these hyperparameters are part of **adjusting the update path**. A larger η leads to bigger steps in gradient descent, which can skip sharp minima and converge on flatter ones, although too large of a step risks divergence. Similarly, a smaller B aids in escaping sharp minima and encourages convergence to flatter minima [11], improving generalization [8] and potentially increasing the likelihood of finding overlapping local optima, a topic that requires further investigation.

In conclusion, our exploration of optimization algorithms, from GD to FL, reveals that client drift stems from inconsistencies in loss landscapes. Based on these findings, we suggest that methods addressing performance degradation in non-IID settings can be categorized into two main strategies: adjusting the update path or modifying the loss landscape. While the motivations behind these methods may vary, they generally align with these categories. We hope this analysis provides valuable insights for new FL researchers and contributes to ongoing discussions in the field.

Acknowledgments. The authors acknowledge the Research Council of Norway and the industry partners of NCS2030 – RCN project number 331644 – for their support.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Acar, D.A.E., Zhao, Y., Matas, R., Mattina, M., Whatmough, P., Saligrama, V.: Federated learning based on dynamic regularization. In: ICLR (2021)
2. Al-Shedivat, M., Gillenwater, J., Xing, E., Rostamizadeh, A.: Federated learning via posterior averaging: A new perspective and practical algorithms. In: ICLR (2021)
3. Chu, C.T., Kim, S., Lin, Y.A., Yu, Y., Bradski, G., Olukotun, K., Ng, A.: Map-reduce for machine learning on multicore. *NeurIPS* (2006)
4. 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. *NeurIPS* (2012)
5. Durmus, A.E., Yue, Z., Ramon, M., Matthew, M., Paul, W., Venkatesh, S.: Federated learning based on dynamic regularization. In: ICLR (2021)
6. 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. *Foundations and trends in machine learning* (2021)
7. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S., Stich, S., Suresh, A.T.: Scaffold: Stochastic controlled averaging for federated learning. In: ICML. PMLR (2020)
8. Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On large-batch training for deep learning: Generalization gap and sharp minima. In: ICLR (2017)
9. Lee, G., Jeong, M., Kim, S., Oh, J., Yun, S.Y.: Fedsol: Stabilized orthogonal learning with proximal restrictions in federated learning. In: CVPR (2024)
10. Lee, G., Jeong, M., Shin, Y., Bae, S., Yun, S.Y.: Preservation of the global knowledge by not-true distillation in federated learning. *NeurIPS* (2022)
11. Li, H., Xu, Z., Taylor, G., Studer, C., Goldstein, T.: Visualizing the loss landscape of neural nets. *NeurIPS* (2018)
12. Li, Q., He, B., Song, D.: Model-contrastive federated learning. In: CVPR (2021)
13. Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., Smith, V.: Federated optimization in heterogeneous networks. *MLSys* (2020)
14. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: AISTATS. PMLR (2017)
15. Qu, Z., Li, X., Duan, R., Liu, Y., Tang, B., Lu, Z.: Generalized federated learning via sharpness aware minimization. In: ICML. PMLR (2022)
16. Reddi, S.J., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., McMahan, H.B.: Adaptive federated optimization. In: ICLR (2021)
17. Ruder, S.: An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747* (2016)
18. Sergeev, A., Del Balso, M.: Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799* (2018)
19. Stich, S.U.: Local SGD converges fast and communicates little. In: ICLR (2019)
20. Sun, Y., Shen, L., Chen, S., Ding, L., Tao, D.: Dynamic regularized sharpness aware minimization in federated learning: Approaching global consistency and smooth landscape. In: ICML. PMLR (2023)
21. Zhang, J., Li, Z., Li, B., Xu, J., Wu, S., Ding, S., Wu, C.: Federated learning with label distribution skew via logits calibration. In: ICML. PMLR (2022)
22. Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., Chandra, V.: Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582* (2018)
23. Zinkevich, M., Weimer, M., Li, L., Smola, A.: Parallelized stochastic gradient descent. *NeurIPS* (2010)