Peer-to-Peer Federated Learning with Trusted Data Sharing for Non-IID Mitigation

Mahran Jazi[a] and Irad Ben-Gal[b]
Department of Industrial Engineering, Tel Aviv University, Tel Aviv, Israel

Keywords: Federated Learning, Data Sharing, Non-IID Data, Decentralized Machine Learning, Edge Intelligence, Distributed Optimization.

Abstract: Collaboration between edge devices without a central server defines the foundation of Peer-to-Peer Federated Learning (P2P FL), a decentralized approach to machine learning that preserves user privacy. However, P2P FL faces significant challenges when data distributions across clients are non-independent and identically distributed (non-IID), which can severely degrade learning performance. In this work, we propose an enhancement to P2P FL through direct data sharing between trusted peers, such as friends, colleagues, or collaborators, where each client shares a small, controlled portion of its local dataset with a selected set of neighbors. While this data-sharing mechanism enhances consistency in learning and improves model performance across the decentralized network, it introduces a trade-off between privacy and performance, as limited data sharing may increase privacy risks. To mitigate these risks, our approach assumes a trusted peer-to-peer network and avoids reliance on any central authority. We evaluate our approach using standard datasets (MNIST, CIFAR-10, and CIFAR-100) and models, including logistic regression, multilayer perceptron, convolutional neural networks (CNNs), and DenseNet-121. The results demonstrate that even modest amounts of peer data sharing significantly improve performance in non-identically distributed (non-IID) settings, offering a simple yet effective strategy to address the challenges of decentralized learning in peer-to-peer federated learning (P2P FL) systems.

1 INTRODUCTION

Edge devices such as smartphones, IoT sensors, and embedded systems increasingly serve as the primary source of private user data. These devices collect and process sensitive information, from health metrics to personal media, and support applications powered by machine learning (ML). Although traditional machine learning (ML) pipelines rely on aggregating data on centralized servers, this architecture raises significant concerns about user privacy, communication overhead, and system scalability (McMahan et al., 2017; Konečnỳ et al., 2015).

Federated Learning (FL) is recognized as a privacy-preserving alternative to traditional centralized ML, enabling distributed training of models while clients retain their data locally and only share model updates (McMahan et al., 2017). This approach mitigates privacy concerns and reduces the need for transferring raw data to centralized servers. This architecture introduces challenges such as communication bottlenecks, system scalability issues, and a single point of failure, which can hinder the robustness and efficiency of FL systems.

To overcome these limitations, *Peer-to-Peer Federated Learning (P2P FL)* has gained traction. In P2P FL, clients collaborate over a decentralized network without a central aggregator(Lalitha et al., 2019; Hegedüs et al., 2022; Tang et al., 2018). Clients share and update their models through local interactions with neighbors, forming communication graphs such as rings, meshes, or random networks. This decentralized setup enhances fault tolerance and eliminates reliance on a central authority, making it suitable for dynamic or large-scale systems, such as ad hoc networks or IoT environments.

However, a core challenge remains: *heterogeneity of the data*. In real-world scenarios, clients typically possess non-independent and identically distributed (non-IID) data due to their unique usage patterns, local contexts, or environments (Zhao et al., 2018; Kairouz et al., 2019). This heterogeneity can

[a] https://orcid.org/0000-0001-6432-3800
[b] https://orcid.org/0000-0003-2411-5518

Jazi, M. and Ben-Gal, I.
Peer-to-Peer Federated Learning with Trusted Data Sharing for Non-IID Mitigation.
DOI: 10.5220/0013685100004000

In *Proceedings of the 17th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2025) - Volume 1: KDIR*, pages 248-256

lead to divergent local model updates, degraded convergence, and suboptimal global performance.

To address this, we propose a novel enhancement to P2P FL: *data sharing between trusted peers*. In many practical scenarios, such as those involving friends, colleagues, or family, privacy concerns are often minimized due to social trust, and data sharing is already common (e.g., through messaging apps, shared documents, or collaborative platforms). Inspired by this natural behavior, we enable peers to share a small, controlled portion of their private datasets with their neighbors. This peer-level data sharing introduces beneficial overlap in local training sets, smoothing the non-IID effects and improving model convergence.

While the core philosophy of Federated Learning (FL) avoids sharing raw data to preserve privacy, our work explores a controlled extension of Peer-to-Peer FL (P2P FL) for contexts where limited data sharing is acceptable. Specifically, we assume settings such as small research collaborations, corporate departments, or circles of peers with established confidentiality agreements, where participants are willing to share a *small, predefined subset* of their data with trusted neighbors.

We fully acknowledge that this assumption does not hold in all FL scenarios and that any sharing of raw data introduces privacy risks. Rather than claiming zero privacy cost, we position our approach as a **trade-off** between improved learning performance in highly non-IID environments and a consciously accepted level of privacy risk in domains with existing trust relationships. This approach is not intended as a general replacement for FL, but as a targeted strategy for specific, privacy-aligned networks.

In our design, data exchange is limited to such trusted peers, where the benefits of improved performance are considered to outweigh the controlled risks. We argue that this assumption is both realistic and practical in modern edge-computing environments involving social, collaborative, or co-located devices. Furthermore, we investigate how P2P FL performs when users generate non-IID data but share a small portion with peers, an environment where adversarial threats are less prevalent and privacy or security concerns are comparatively minimal (Lyu et al., 2022; Chen et al., 2022; Liu et al., 2022; Jazi and Ben-Gal, 2024; Zang et al., 2024).

The remainder of this paper is organized as follows: Section 2 discusses the contributions of this work. Section 3 provides a detailed review of the related work in the domain. Section 4 presents the problem formulation and its theoretical framework. Section 5 introduces and explains the proposed peer-level data sharing mechanism and the P2P algorithm. Section 6 highlights the experimental results and their analysis, section 7 delves into non-IID partitioning in P2P-FL, exploring its implications on the proposed methodology. Finally, Section 8 concludes the paper, summarizing the contributions and providing directions for future research.

2 CONTRIBUTIONS

We propose trusted peer-level data sharing as a simple and efficient technique to boost the performance of peer-to-peer federated learning (P2P FL), as illustrated in Figure 1. Our findings apply to two distinct scenarios: active and passive data sharing among decentralized clients.

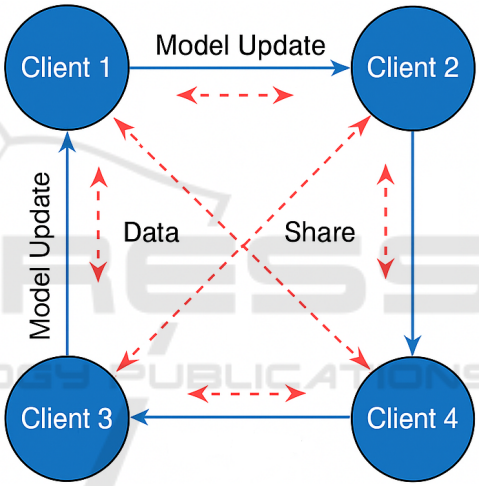

Figure 1: Illustration of Peer-to-Peer Federated Learning with Data Sharing. Clients exchange model updates with neighbors in a decentralized fashion (solid blue arrows) while optionally sharing a portion of their data with trusted peers (dashed red arrows) to mitigate non-IID effects.

In *active data sharing*, clients are explicitly encouraged to share a portion of their data with a selected set of trusted neighbors within the communication graph. While this involves a partial relaxation of local data privacy, it is essential to note that no data is transmitted to any centralized server; data instead is shared only between peers with pre-established trust relationships (e.g., friends, family, or colleagues). In practical implementations, this can be supported by device-level applications that enable users to share data with specific peers selectively. The dynamics of such interactions can be further modeled using incentive mechanisms or game-theoretic frameworks (Buratto et al., 2024).

In *passive data sharing*, we assume that data ex-

change naturally exists within certain peer groups due to ongoing digital interactions, such as shared cloud folders, group chats, or collaborative devices. When P2P FL is executed over such socially connected clients, the implicit overlap in their datasets results in improved data diversity across the network. Our results demonstrate that natural groupings can significantly enhance model performance, even when the total number of data samples per client remains constant.

Importantly, our experiments abstract away from the specific sharing mechanism (active or passive) and instead evaluate the effect of sharing varying percentages of data among clients. We demonstrate that peer-level data sharing improves model accuracy and convergence under non-IID data distributions. Moreover, we find that a relatively small fraction of shared data (e.g., 20,40%) is sufficient to yield near-optimal improvements.

The principal finding of this study is, therefore, that *Running decentralized FL on socially connected peers who share data either actively or passively outperforms training on arbitrarily grouped clients, even under identical data volumes*.

3 RELATED WORK

Federated Learning (FL) has emerged as a leading approach to enable collaborative machine learning across multiple clients while preserving data privacy (McMahan et al., 2017). In its canonical form, FL relies on a centralized server that orchestrates communication rounds and aggregates model updates from clients, as exemplified by the widely-used FedAvg algorithm (McMahan et al., 2016). This architecture, however, introduces potential bottlenecks and single points of failure, raising concerns over scalability, robustness, and trust in large-scale and dynamic environments.

To address these limitations, decentralized FL architectures have been proposed, where clients interact directly with each other in a peer-to-peer (P2P) manner, thereby removing the dependency on a central aggregator (Lalitha et al., 2019; Hegedüs et al., 2022; Tang et al., 2018). A popular and effective communication paradigm in this domain is the use of gossip-based protocols, where clients iteratively exchange and average model parameters with randomly selected neighbors (Boyd et al., 2006). Gossip algorithms enhance scalability, fault tolerance, and privacy by leveraging localized communication and avoiding central coordination (Mishchenko et al., 2021). Recent works have demonstrated the theoretical convergence and practical viability of gossip-based federated learning (FL), especially in decentralized and infrastructure-less networks (Hegedüs et al., 2022; Mishchenko et al., 2021). However, these approaches primarily focus on model parameter aggregation without explicitly addressing the heterogeneity of client data distributions, and they do not consider non-convex problems such as those encountered in deep neural networks (DNNs).

Data heterogeneity, or non-independent and identically distributed (non-IID) data across clients, remains a fundamental challenge in both centralized and decentralized federated learning (FL) (Zhao et al., 2018; Kairouz et al., 2019; Hsieh et al., 2020; Sery et al., 2021). Non-IID data can cause local models to diverge significantly, leading to slower convergence and reduced global model accuracy (Zhu et al., 2021). Various strategies have been proposed to mitigate these issues, including personalized federated learning (FL) (Smith et al., 2017), client clustering (Sattler et al., 2020; Ouyang et al., 2021; Yang et al., 2023), and adaptive aggregation techniques (Li et al., 2020c). While these approaches improve performance at the model or algorithmic level, they do not directly modify the underlying data distributions.

To address this limitation, (Zhao et al., 2018) proposed using a shared synthetic dataset, uniformly distributed over the data space and generated by a central server. This dataset is then distributed to all clients to make their local data more independent and identically distributed (IID). Although the approach is conceptually simple and effective, it has significant drawbacks: it requires the central server to be aware of the global data distribution and imposes a storage burden on clients. Alternatively, (Yoshida et al., 2020) proposed a hybrid scheme in which only a small portion of private data is shared with the server to enhance privacy. While this reduces the amount of shared sensitive information, it still poses a risk of privacy leakage.

Limited and controlled data sharing between clients has been explored as a complementary approach to enhance learning under non-identical and independent (non-IID) conditions. Although conventional federated learning (FL) aims to avoid raw data exchange to protect privacy in specific trusted environments, such as among social peers, collaborative organizations, or co-located devices, small-scale data sharing is practical and beneficial (Li et al., 2020a). Hybrid frameworks that integrate model distillation with selective data sharing have been proposed to enrich client datasets and improve global learning (Jeong et al., 2019; Li et al., 2020a). Nevertheless, these frameworks typically operate within centralized

federated learning (FL) paradigms and do not leverage data sharing in fully decentralized peer-to-peer (P2P) networks.

Despite these prior advancements, the relationship between data-sharing mechanisms and decentralized federated learning in P2P networks remains underexplored. Existing gossip-based frameworks focus heavily on scalability and fault tolerance but neglect the complexities introduced by data heterogeneity (Boyd et al., 2006; Hegedüs et al., 2022). Similarly, hybrid approaches for data sharing (Jeong et al., 2019; Li et al., 2020a) are limited to centralized settings and do not address the unique challenges of decentralized, infrastructure-less networks. Our work builds on these foundations by combining structured peer-level data-sharing mechanisms with gossip-based communication to directly address non-IID data challenges in decentralized P2P FL environments. This integration positions our framework as a practical and scalable solution for real-world deployments, advancing beyond the state-of-the-art in both centralized and decentralized FL paradigms.

Our work advances the state-of-the-art by introducing a novel P2P FL framework that explicitly incorporates a structured data-sharing mechanism among trusted peers. This approach leverages the advantages of gossip-based decentralized communication while addressing data heterogeneity through peer-level data exchange. By doing so, it mitigates non-IID challenges without sacrificing privacy and decentralization. Our theoretical analysis and empirical results across diverse datasets and models demonstrate that even modest data sharing significantly enhances convergence speed and model accuracy, offering a practical and scalable solution for real-world P2P FL deployments.

4 PROBLEM FORMULATION

We propose a peer-to-peer federated learning (P2P-FL) framework enhanced with selective data sharing, where N clients collaboratively train a shared model by directly communicating model updates with their peers, eliminating the need for a centralized server. To counter the challenges posed by non-IID data distributions, each client is permitted to share a portion of its local data with neighboring clients. This targeted data sharing enhances convergence and alignment of distribution.

Inspired by prior work on decentralized training algorithms (He et al., 2020; Hegedüs et al., 2019), we extend the Federated Averaging (FedAvg) algorithm to function in a fully peer-to-peer (P2P) environment. Our model synchronization process is based on gossip-based communication protocols(Hegedüs et al., 2019), where each client periodically exchanges model parameters with a subset of its neighboring peers. After averaging these parameters, clients update their local models accordingly and proceed with further local training. This decentralized adaptation eliminates the need for a central server, enhancing scalability and robustness while preserving the collaborative benefits of federated learning.

Our primary objective is to investigate how the inclusion of a shared data component affects model performance in decentralized, non-i.i.d. settings. Rather than enforcing a fixed method for sharing (e.g., active vs. passive), we focus on the statistical effect of data overlap. Our findings demonstrate that even modest data sharing can significantly enhance performance, outperforming purely decentralized setups with equivalent total data volumes. This suggests that distribution alignment, not just increased data quantity, plays a critical role.

We consider a peer-to-peer federated learning (P2P-FL) setting involving N clients collaboratively engaged in a classification task, where the objective is to learn a model that maps each input to one of K possible classes. Each client n holds a private local dataset denoted by $\mathcal{D}_n = \{x_i^n\}_{i=1}^{M_n}$, consisting of M_n data samples and their corresponding labels $\{y_i\}$. Each client maintains its model, parameterized by $\omega^n \in \mathbb{R}^p$, where p is the total number of trainable parameters. All clients share an identical model architecture.

Let $l(\omega, x, y)$ represent the loss incurred on a single data point (x, y). The local loss function for client n is given by:

$$\mathcal{L}_n(\omega^n) \triangleq \sum_{x \in \mathcal{D}_n} l(\omega^n, x, y(x)). \tag{1}$$

The global learning objective in this decentralized setting is to train models that collectively minimize the total loss across all clients:

$$\mathcal{L}(\omega) = \sum_{n=1}^{N} \mathcal{L}_n(\omega). \tag{2}$$

Unlike centralized FL, where a central server coordinates aggregation, in P2P-FL, each client independently exchanges model updates with a randomly selected subset of peers at each round t. Let $S_t^n \subseteq \mathcal{N}_n$ denote the set of m peers that client n communicates with at round t. The client computes a weighted average of the model parameters received from its peers, along with its locally updated model:

$$\overline{\omega}_t^n = \frac{M_n}{M_t^n} \omega_{\text{local}}^n + \sum_{j \in S_t^n} \frac{M_j}{M_t^n} \omega_t^j, \tag{3}$$

where $M_t^n = M_n + \sum_{j \in S_t^n} M_j$ denotes the total number of samples considered in the local peer aggregation. This averaged model $\overline{\omega}_t^n$ serves as the reference model for the next local update.

Each client then updates its model using a local stochastic gradient descent (SGD) step with momentum $0 \leq \beta < 1$:

$$\omega_{t+1}^n = \overline{\omega}_t^n - \eta_t v_{t+1}^n, \tag{4}$$

where η_t is the learning rate at round t, and v_{t+1}^n is the momentum-augmented gradient defined as:

$$v_{t+1}^n = \beta v_t^n + g_n(\overline{\omega}_t^n). \tag{5}$$

The stochastic gradient $g_n(\overline{\omega}_t^n)$ is computed over a mini-batch $\mathcal{S}_n \subset \mathcal{D}_n$ of size B:

$$g_n(\overline{\omega}_t^n) = \sum_{x \in \mathcal{S}_n} \nabla l(\overline{\omega}_t^n, x, y(x)). \tag{6}$$

This decentralized protocol enables each client to iteratively refine its model by leveraging knowledge from a dynamically selected neighborhood, thereby promoting robustness and scalability in the absence of a central server.

5 PEER-LEVEL DATA SHARING AND P2P ALGORITHM

To tackle the limitations introduced by non-IID data, we incorporate a data exchange protocol among clients. Each client contributes a fraction $\Delta \in [0,1]$ of its local dataset to selected neighbors. The updated dataset at client n becomes:

$$\tilde{\mathcal{D}}_n = \mathcal{D}_n \cup \bigcup_{m \in \mathcal{S}_n} \mathcal{D}_{m \to n}, \tag{7}$$

where $\mathcal{S}_n \subseteq \mathcal{N}_n$ represents the set of contributing peers, and $\mathcal{D}_{m \to n} \subseteq \mathcal{D}_m$ is the subset of data points (with $|\mathcal{D}_{m \to n}| = \Delta \cdot M_m$) transferred from client m.

This setting mirrors practical environments where users already share data in socially trusted relationships (e.g., family, friends, coworkers), allowing us to assume reduced privacy concerns. Importantly, no data is transferred to any central authority, preserving the decentralized and privacy-conscious nature of the system.

Through this lens, we evaluate how data sharing impacts learning quality under diverse model structures and datasets, focusing on realistic non-IID scenarios.

We investigate the impact of data sharing on the performance of Federated Learning (FL). Once data sharing is complete, the FL process proceeds using the previously outlined decentralized algorithm. A critical feature of our approach is that client data is never transmitted to a central server. Instead, data is exchanged exclusively among socially connected peers, maintaining the privacy of individual clients. This form of data sharing can occur organically without the need for centralized orchestration.

Algorithm 1 describes the procedure for assigning data to each client. It is designed to ensure that every client ends up with a fixed number of data points, regardless of the degree of data sharing. While sharing data naturally leads to an increase in the total number of samples available to a client, we aim to isolate the effect of data distribution rather than data volume. To achieve this, the baseline FL setup with $\Delta = 0\%$ (no data sharing) is constructed to match the final dataset size per client (M_n). However, it contains more unique data samples than the setup involving shared data. This approach allows for a fair comparison that focuses on distributional benefits rather than data size advantages.

We adopt the simplifying assumption that all client pairs share an equal amount of data. This ensures the sharing process can be captured using a single, interpretable parameter Δ, representing the proportion of shared data per client.

6 EXPERIMENTAL RESULTS

In this section, we evaluate the performance of our proposed Peer-to-Peer Federated Learning (P2P-FL) framework with data sharing. Data sharing in P2P-FL facilitates improved learning outcomes by guiding the optimization process toward better stationary points, particularly under non-convex loss landscapes that are common in real-world machine learning models.

The experiments were conducted using Google Colab Pro, which provides access to high-performance resources, including an A100 GPU and extended RAM, to support the execution of computationally intensive models and large-scale datasets in a Python-based environment.

The evaluation metric used is classification accuracy, defined as the percentage of correctly classified samples in the test dataset after the training phase using the P2P-FL model.

To ensure a fair and comprehensive assessment, we employed the following widely used benchmark datasets:

- **MNIST:** A dataset of 70,000 grayscale images of handwritten digits (0–9), with 60,000 samples used for training and 10,000 for testing (LeCun et al., 1998).

Algorithm 1: P2P Data Sharing and Training Procedure.

Require: Number of clients N, initial dataset size M_0, final dataset size M, data sharing ratio Δ, communication graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

Initialize each client $n \in \mathcal{V}$ with M_0 private data samples $\mathcal{D}_n$
for all client pairs (n,m) such that $(n,m) \in \mathcal{E}$ **do**
  Client n shares $\Delta \cdot M_0$ randomly selected samples from $\mathcal{D}_n$ with client m
  Client m augments its dataset: $\mathcal{D}_m \leftarrow \mathcal{D}_m \cup \mathcal{D}_{n \to m}$
end for
for all clients $n \in \mathcal{V}$ **do**
  Compute $|\mathcal{D}_n|$ after sharing
  Add new (non-overlapping) samples from a global pool to reach final size $|\mathcal{D}_n| = M$
end for
Initialize local model parameters ω^n for each client n
for each communication round $t = 1$ to T **do**
  for all clients $n \in \mathcal{V}$ **do**
    Perform local SGD on $\mathcal{D}_n$ to compute $g_n(\omega_t^n)$
    Receive ω_t^j from neighbors $j \in \mathcal{N}_n$
    Compute weighted average:

$$\overline{\omega}_t^n = \frac{M_n}{M_t^n}\omega_t^n + \sum_{j \in \mathcal{N}_n} \frac{M_j}{M_t^n}\omega_t^j$$

    Update model with momentum:

$$v_{t+1}^n = \beta v_t^n + g_n(\overline{\omega}_t^n)$$

$$\omega_{t+1}^n = \overline{\omega}_t^n - \eta_t v_{t+1}^n$$

  end for
end for
return Final model parameters ω^n for all clients

- **CIFAR-10:** A dataset consisting of 60,000 32×32 color images across 10 object classes. It is split into 50,000 training images and 10,000 test images (Krizhevsky and Hinton, 2009).
- **CIFAR-100:** Similar to CIFAR-10, but with 100 classes, using the same 50,000/10,000 split for training and testing, respectively (Krizhevsky and Hinton, 2009).

To establish the general applicability of P2P-FL, we tested it using four standard machine-learning models with varying levels of complexity:

- **Logistic Regression (LR):** A baseline linear classifier.
- **2NN:** A simple multilayer perceptron with two hidden layers, each containing 200 ReLU-activated units (McMahan et al., 2017; Zhao et al., 2018).
- **LeNet-5:** A convolutional neural network (CNN) for image recognition (LeCun et al., 1998).
- **DenseNet-121:** A deeper CNN model representing a high-complexity architecture (Huang et al., 2017).

Each model was trained using stochastic gradient descent (SGD) over 100 communication rounds. The hyperparameters used were batch size $B = 32$, local epoch $E = 1$ (i.e., one pass over the local dataset), learning rate $\eta = 0.01$, and momentum parameter $\beta = 0.9$. Each experiment was repeated 10 times independently, and we report the average results. The error bars represent one standard deviation above and below the average.

To assess the benefits of our data-sharing mechanism in the P2P-FL setting, we compared it against the classical Federated Averaging (FedAvg) algorithm (McMahan et al., 2017), which does not involve peer-to-peer data sharing. Additionally, we included the Federated Proximal (FedProx) algorithm (Li et al., 2020b) as a benchmark, as it is designed to address heterogeneity in federated learning environments. We benchmark our results against three scenarios:

- **No Data Sharing (FedAvg):** Standard FL with no overlap in data between clients.
- **Federated Proximal (FedProx):** A variant of FL that incorporates a proximal term to better handle statistical heterogeneity across clients, serving as a benchmark for heterogeneous federated learning settings.
- **Full Data Sharing (Centralized Baseline):** All data is shared among all clients, effectively equivalent to a centralized model with access to the full dataset. Training proceeds by averaging the gradients from all N clients as if the data were fully centralized.

These comparisons highlight the potential of P2P data sharing to bridge the performance gap between decentralized and centralized learning settings, while preserving privacy and avoiding dependency on a central server.

7 NON-IID PARTITIONING IN P2P-FL

The results for our proposed P2P-FL approach under non-IID data distribution are presented in Figure 2. In this setting, ten clients participate in training, each starting with local data that only includes

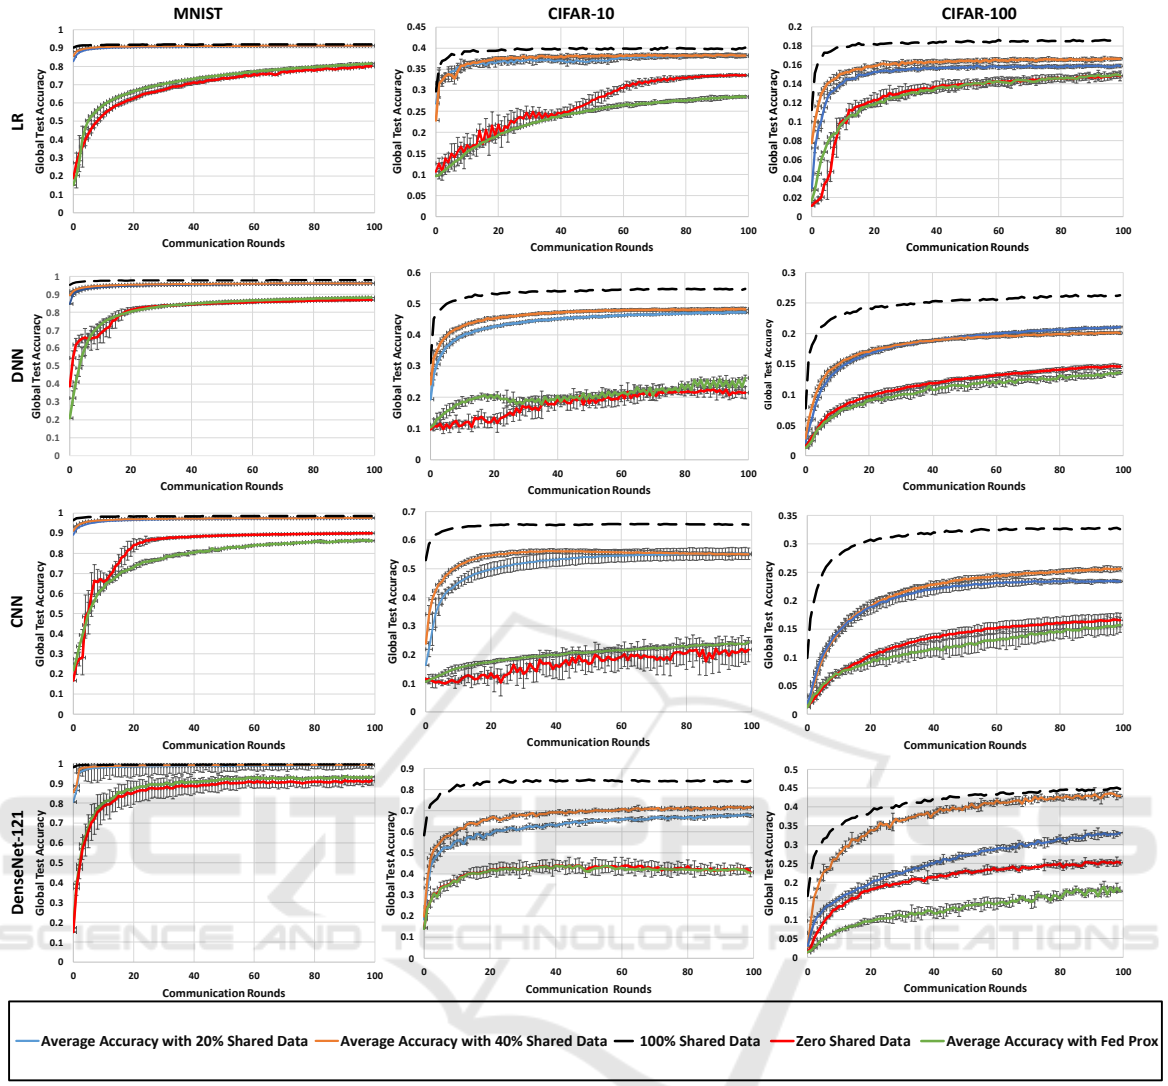

Figure 2: Data sharing with non-IID data distributions. The columns correspond to the MNIST, CIFAR-10, and CIFAR-100 datasets. The rows correspond to the models LR, 2NN, LeNet-5, and DenseNet-121. The comparison includes FedAvg (red), FedProx (green), and SFL with various data sharing levels.

exclusive classes: for MNIST or CIFAR-10, client i contains data solely from class i, for $i = 1, \ldots, 10$, and for CIFAR-100, client i contains classes indexed as $10(i-1), 10(i-1)+1, \ldots, 10(i-1)+9$. Thus, no class overlap occurs across clients, representing an extreme non-IID scenario.

Each client initially holds $M_0 = 1000$ data samples. After applying the peer-to-peer data-sharing mechanism, the final number of local data points per client becomes $M = 6000$. In our framework, '20% data sharing' refers to each client randomly selecting 20% of its local dataset and independently sharing that subset with its connected peers, rather than broadcasting to a central server. This behavior aligns with the decentralized P2P communication paradigm (see Algorithm 1).

Our empirical results demonstrate that a fully decentralized P2P-FL system with no data sharing (red line in Figure 2) performs poorly in the presence of extreme non-IID distributions, consistent with observations in prior work (Zhao et al., 2018). In addition, FedProx ($\mu = 0.01$, represented by the green line), which incorporates a proximal term to mitigate client drift, demonstrates marginal improvements over FedAvg in specific datasets and models. However, it continues to face significant challenges when operating under highly non-IID data distributions. The modest peer-level data sharing (e.g., 20%) significantly improves learning outcomes across all tested models and datasets.

The figure illustrates a diminishing return effect with respect to the percentage of shared data and the number of communication rounds. We evaluate sharing levels of 0%, 20%, 40%, and 100%, with each client maintaining a consistent private dataset size of 6,000 points. This ensures that performance gains stem from better statistical diversity rather than increased data quantity. Notably, increasing the data-sharing level beyond 20% yields only marginal benefits, especially after 40 communication rounds.

The performance boost is more pronounced in complex scenarios, such as DenseNet-121 on CIFAR-100, where the interplay between model capacity and data heterogeneity becomes increasingly critical. These findings suggest that P2P-FL with partial data sharing not only mitigates statistical heterogeneity but also aligns better with task complexity when deeper models or more nuanced datasets are used.

8 CONCLUSION

In this paper, we propose a novel peer-to-peer federated learning framework enhanced with a structured data-sharing mechanism among trusted peers. Our approach addresses the critical challenge of non-IID data distributions in fully decentralized FL systems by enabling clients to exchange small subsets of their local data directly with neighboring clients. This strategy enhances dataset diversity and significantly mitigates the adverse effects of data heterogeneity, resulting in faster convergence and improved model accuracy.

Through extensive experiments on diverse datasets and models, we demonstrated that even modest data sharing substantially enhances learning performance compared to P2P FL without data sharing. Importantly, our method preserves the core principles of privacy and decentralization in federated learning by eliminating reliance on a central server and restricting data exchange to trusted peers.

Future work includes exploring adaptive data-sharing strategies to optimize the trade-off between privacy and performance, as well as extending the framework to dynamic and large-scale P2P networks with varying trust relationships. To further address privacy concerns, future work can focus on integrating privacy-preserving techniques such as differential privacy (Dwork et al., 2006), secure multiparty computation (Bonawitz et al., 2017), or homomorphic encryption (Aono et al., 2017) with our proposed approach. These extensions could reduce the risks associated with data sharing while maintaining the performance benefits demonstrated in this study. Our findings open up promising avenues for the practical deployment of decentralized federated learning (FL) in real-world applications where privacy, scalability, and data heterogeneity coexist.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge funding from the Koret Foundation Grant for Smart Cities and Digital Living 2030 and the Neubauer Family Foundation.

REFERENCES

Aono, Y., Hayashi, T., Abadi, M., Kono, K., and Han, S. (2017). Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5):1333–1345.

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. (2017). Practical secure aggregation for federated learning on user-held data. In *Proceedings of the 24th ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191. ACM.

Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. (2006). Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530.

Buratto, A., Guerra, E., Miozzo, M., Dini, P., and Badia, L. (2024). Energy minimization for participatory federated learning in iot analyzed via game theory. In *2024 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*, pages 249–254. IEEE.

Chen, Y., Gui, Y., Lin, H., Gan, W., and Wu, Y. (2022). Federated learning attacks and defenses: A survey. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 4256–4265. IEEE.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284. Springer.

He, X. et al. (2020). Training linear models in a fully decentralized environment. In *Proceedings of the International Conference on Machine Learning*.

Hegedüs, I., Berta, D., and Jelasity, M. (2019). Gossip learning as a decentralized alternative to federated learning. In *Proceedings of the ACM International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Hegedüs, I., Danner, G., and Schad, J. (2022). Decentralized federated learning: A survey and perspective. *ACM Computing Surveys*, 55(11):1–37.

Hsieh, K., Phanishayee, A., Mutlu, O., and Gibbons, P. (2020). The non-iid data quagmire of decentralized machine learning. In *International Conference on Machine Learning*, pages 4387–4398. PMLR.

Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708.

Jazi, M. and Ben-Gal, I. (2024). Federated learning for xss detection: A privacy-preserving approach. In *Proceedings of the 16th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management - Volume 1: KDIR*, pages 283–293. INSTICC, SciTePress.

Jeong, E., Oh, S., Kim, S., Kang, J., Kim, J.-S., and Lee, S.-E. (2019). Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5713–5720.

Kairouz, P., McMahan, H. B., et al. (2019). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2):1–210.

Konečnỳ, J., McMahan, B., Ramage, D., et al. (2015). Federated optimization: Distributed optimization beyond the datacenter. *arXiv preprint arXiv:1511.03575*.

Krizhevsky, A. and Hinton, G. (2009). Learning multiple layers of features from tiny images. Technical Report, University of Toronto.

Lalitha, A., Javidi, T., and Koushanfar, F. (2019). Fully decentralized federated learning. In *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pages 582–586. IEEE.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.

Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020a). Fedmd: Heterogeneous federated learning via model distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3542–3549.

Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020b). Fedprox: A scalable federated learning framework with heterogeneity. *arXiv preprint arXiv:1812.06127*.

Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. (2020c). Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450.

Liu, P., Xu, X., and Wang, W. (2022). Threats, attacks and defenses to federated learning: Issues, taxonomy and perspectives. *Cybersecurity*, 5(1):1–19.

Lyu, L., Yu, H., Ma, X., Chen, C., Sun, L., Zhao, J., Yang, Q., and Yu, S. (2022). Privacy and robustness in federated learning: Attacks and defenses. *IEEE Transactions on Neural Networks and Learning Systems*.

McMahan, B., Moore, E., Ramage, D., Hampson, S., and Aguera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 1273–1282. PMLR.

McMahan, H. B., Moore, E., Ramage, D., et al. (2016). Federated averaging: Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*.

Mishchenko, K., Khaled, A., and Richtarik, P. (2021). Distributed stochastic gradient tracking methods. *Journal of Machine Learning Research*, 22(153):1–54.

Ouyang, X., Xie, Z., Zhou, J., Huang, J., and Xing, G. (2021). Clusterfl: A similarity-aware federated learning system for human activity recognition. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 54–66.

Sattler, F., Müller, K.-R., and Samek, W. (2020). Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3710–3722.

Sery, T., Shlezinger, N., Cohen, K., and Eldar, Y. C. (2021). Over-the-air federated learning from heterogeneous data. *IEEE Transactions on Signal Processing*, 69:3796–3811.

Smith, V., Chiang, C.-K., Sanjabi, M., and Talwalkar, A. (2017). Federated multi-task learning. *Advances in Neural Information Processing Systems*, 30.

Tang, H., Dube, X., Wang, S., Joshi, G., and Kar, S. (2018). D2: Decentralized training over decentralized data. In *International Conference on Learning Representations (ICLR)*.

Yang, L., Huang, J., Lin, W., and Cao, J. (2023). Personalized federated learning on non-iid data via group-based meta-learning. *ACM Transactions on Knowledge Discovery from Data*, 17(4):1–20.

Yoshida, N., Nishio, T., Morikura, M., Yamamoto, K., and Yonetani, R. (2020). Hybrid-fl for wireless networks: Cooperative learning mechanism using non-iid data. In *ICC 2020 - IEEE International Conference on Communications*, pages 1–7. IEEE.

Zang, M., Zheng, C., Koziak, T., Zilberman, N., and Dittmann, L. (2024). Federated in-network machine learning for privacy-preserving iot traffic analysis. *ACM Transactions on Internet Technology*, 24(4):1–24.

Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. (2018). Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.

Zhu, H., Xu, J., Liu, S., and Jin, Y. (2021). Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390.