

# pFedGPA: Diffusion-based Generative Parameter Aggregation for Personalized Federated Learning

Jiahao Lai<sup>1</sup>, Jiaqi Li<sup>1</sup>, Jian Xu<sup>1</sup>, Yanru Wu<sup>1</sup>, Boshi Tang<sup>1</sup>, Siqi Chen<sup>1</sup>,  
Yongfeng Huang<sup>2</sup>, Wenbo Ding<sup>1</sup>, Yang Li<sup>1\*</sup>

<sup>1</sup> Tsinghua Shenzhen International Graduate School, Tsinghua University<sup>†</sup>

<sup>2</sup> Department of Computer Science and Engineering, The Chinese University of Hong Kong  
laihj22@mails.tsinghua.edu.cn, yangli@sz.tsinghua.edu.cn

## Abstract

Federated Learning (FL) offers a decentralized approach to model training, where data remains local and only model parameters are shared between the clients and the central server. Traditional methods, such as Federated Averaging (FedAvg), linearly aggregate these parameters which are usually trained on heterogeneous data distributions, potentially overlooking the complex, high-dimensional nature of the parameter space. This can result in degraded performance of the aggregated model. While personalized FL approaches can mitigate the heterogeneous data issue to some extent, the limitation of linear aggregation remains unresolved. To alleviate this issue, we investigate the generative approach of diffusion model and propose a novel generative parameter aggregation framework for personalized FL, pFedGPA. In this framework, we deploy a diffusion model on the server to integrate the diverse parameter distributions and propose a parameter inversion method to efficiently generate a set of personalized parameters for each client. This inversion method transforms the uploaded parameters into a latent code, which is then aggregated through denoising sampling to produce the final personalized parameters. By encoding the dependence of a client’s model parameters on the specific data distribution using the high-capacity diffusion model, pFedGPA can effectively decouple the complexity of the overall distribution of all clients’ model parameters from the complexity of each individual client’s parameter distribution. Our experimental results consistently demonstrate the superior performance of the proposed method across multiple datasets, surpassing baseline approaches.

## Introduction

To meet with the increasing needs of privacy protection, federated learning (FL) has been a popular machine learning paradigm and research topic for many years (McMahan et al. 2017a). Typically, multiple devices such as smartphones, sensors, and IoTs (Internet of Things) collaboratively train a global model under the coordination of a central server. FL

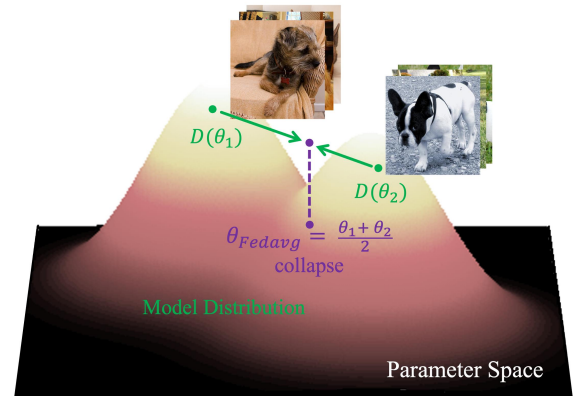


Figure 1: Parameter collapse can occur when linearly averaging the parameters from different clients. Bright colors indicate high-probability regions of the parameter space, where the parameters located at the peaks of the model distribution are well-optimized for specific tasks.

has been extended to various application domains, including healthcare and finance (Nguyen et al. 2022). In healthcare, FL trains models on distributed patient data, enabling cooperative research without compromising privacy (Yang et al. 2019). In finance, FL helps developing fraud detection models using data from multiple financial institutions while ensuring confidentiality (Li et al. 2019).

One challenge in FL is handling the heterogeneity of data distributions across various devices (Zhao et al. 2018; Kairouz et al. 2021). For example, in healthcare applications, patient data collected from different hospitals can vary greatly due to differences in patient demographics, medical equipment, and local practices. In linear aggregation methods such as FedAvg (McMahan et al. 2017a), data in different clients are assumed to share the same distribution. Non-IID data in local clients can lead to unstable results (Zhao et al. 2018; Karimireddy et al. 2020b), as illustrated in Fig. 1. In this case, the parameters  $\theta_1$  and  $\theta_2$  from non-IID datasets are aggregated to form  $\theta_{Fedavg} = (\theta_1 + \theta_2)/2$ . While  $\theta_1$  and  $\theta_2$  are located in the high-probability regions of the parameter space and are well-optimized for their respective tasks, their average  $\theta_{Fedavg}$  falls into a low-probability region, causing a collapse. This example high-

\*Corresponding author.

<sup>†</sup>Jiahao Lai, Jian Xu, Yanru Wu, Siqi Chen, Wenbo Ding and Yang Li are from the Shenzhen Key Laboratory of Ubiquitous Data Enabling, Tsinghua University.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

lights the limitations of linear aggregators and the need of more ingenious aggregation methods.

To handle Non-IID client data, various personalized FL methods have been proposed. They aim to learn a customized model for each individual client to fit its local data distribution. Although many efforts have been devoted to develop advanced model adaptation methods to obtain better local models, the limitation of linear aggregation remains since most of the advanced personalized FL methods still highly rely on the linear aggregation (Collins et al. 2021; Sattler, Müller, and Samek 2020; Yang et al. 2023; Xu, Tong, and Huang 2023). For instance, a simple but tough-to-beat baseline is combining FedAvg with local fine-tuning, where the quality of the global model by FedAvg may directly affect the resulted local models.

To build better FL models, a fundamental challenge is to explicitly model the relationship between data distributions and model parameters within specific optimization constraints. However, when considering the parameter space where the model distribution resides, we recognize that the intrinsic difficulty arises from modeling this distribution as a low-dimensional manifold within a high-dimensional space (Zhou et al. 2022, 2023). Simple network architectures like multi-layer perceptrons (MLPs) or basic Transformers struggle to effectively learn such complex distributions and scale robustly. Moreover, they lack effective optimization objectives and training algorithms.

Inspired by the remarkable success of diffusion models in achieving state-of-the-art results in image generation (Dhariwal and Nichol 2021) and recent attempts in model parameter generation (Wang et al. 2024), we propose to leverage the diffusion models to capture the distribution of model parameters in the parameter space. Diffusion models, which convert the complex data distribution to an analytically tractable prior, estimate the score function at each time step and iteratively refine samples through stochastic processes (Song et al. 2020). Although initially applied to image data, this versatile approach shows promise for modeling high-dimensional parameter distributions (Zhang et al. 2024; Wang et al. 2024).

Building on above idea, we introduce our framework, pFedGPA, which employs a diffusion model on the server to handle the parameters of clients, enabling the server to learn the distributions of all clients’ model parameters using a powerful generative model. Subsequently, we investigate the mechanisms of generative parameter aggregation and personalized model updates within this framework. We consider the generation process as an alternative to achieve the better model aggregation. Our method effectively guides the optimization direction for newly joined clients and significantly accelerates the initialization. Furthermore, we develop a novel parameter inversion method inspired by image editing techniques applied within diffusion models, which transforms the original parameters into a latent representation and then refines them to generate new parameters that retain the implicit semantics of the original data while incorporating the learned patterns from the diverse parameter distributions. Furthermore, we explore the configuration of the diffusion model architecture tailored for processing model

parameter data. Our experimental results demonstrate the effectiveness of the proposed method, consistently achieving superior performance across multiple datasets compared to baseline approaches. The main contributions of our work are concluded as follows:

- We present a novel FL framework that employs diffusion models to generate personalized parameters with the global guidance in heterogeneous settings. To the best of our knowledge, we are the first to improve model aggregation by applying diffusion models in FL.
- We conducted experiments to evaluate the generation quality by producing partial and full parameters for small models, and partial parameters for large models.
- We empirically verify the superiority of our proposed method in terms of generalization performance on several benchmark datasets under practical scenarios.

## Related Work

### FL with Data Heterogeneity

To address the non-IID data challenge, several strategies for improving the global models have been proposed. FedProx (Li et al. 2020) introduced a local regularization term to optimize each client’s model. SCAFFOLD (Karimireddy et al. 2020a) introduces control variates to correct the drift in local updates. Clustering-based methods (Sattler, Müller, and Samek 2020; Li et al. 2020) cluster similar local models into groups and assign each group a global model. Moreover, self-supervised methods have been incorporated to define the similarity between model representations for correcting local training (Li, He, and Song 2021). Personalized FL (Smith et al. 2017) focuses on tailoring models for individual clients by combining information across clients. Decentralized MAML (Fallah, Mokhtari, and Ozdaglar 2020) adapts the model-agnostic meta-learning framework to a federated setting, allowing clients to learn models through local adaptation (Finn, Abbeel, and Levine 2017; Nichol, Achiam, and Schulman 2018). Another category includes model mixing and layer adaptation strategies, where clients learn a mixture of global and local models (Hanzely and Richtárik 2020; Zhang et al. 2023). Decoupling model layers is also a popular and effective approach (Arivazhagan et al. 2019; Collins et al. 2021; Xu, Yan, and Huang 2022). For instance, FedRep (Collins et al. 2021) and FedBABU (Oh, Kim, and Yun 2022) train base layers globally using FedAvg, with personalized layers fine-tuned locally. Personalized aggregation based on model similarity has also been investigated (Zhang et al. 2021; Ye et al. 2023).

### Diffusion Models

Diffusion models have emerged as a powerful technique of generative AI, particularly excelling in producing high-quality images. Their superior properties have led to widespread applications in various vision (Lugmayr et al. 2022) and multi-modal tasks (Rombach et al. 2022; Kawar et al. 2023; Kumari et al. 2023). The foundational work by (Sohl-Dickstein et al. 2015), studied non-equilibrium thermodynamics, highlighting its potential in generative model-

ing. (Ho, Jain, and Abbeel 2020) advanced this with Denoising Diffusion Probabilistic Models (DDPMs), improving the denoising process for better image synthesis. (Song and Ermon 2019) proposed Score-Based Generative Modeling, achieving remarkable results in diverse applications (Meng et al. 2021; Xu et al. 2022). (Song et al. 2020) further unified diffusion probabilistic modeling and score-based generative modeling using stochastic differential equations (SDEs), demonstrating their equivalence.

Guided diffusion models (Dhariwal and Nichol 2021) introduced mechanisms to steer the generative process. Notable examples include DALL·E 2 (Ramesh et al. 2022). In contrast, classifier-free guidance (Ho and Salimans 2022) integrates guidance directly into the model by conditioning on input prompts, thus avoiding the need for a separate classifier. Examples of this approach include Imagen (Saharia et al. 2022) and Stable Diffusion (Rombach et al. 2022).

## Model Parameter Generation

Model parameter generation has progressed from gradient optimization to advanced meta-learning and Bayesian methods. Early efforts primarily focused on optimizing parameters using gradient-based techniques like Stochastic Gradient Descent (SGD) and its variants (Amari 1993). Meta-learning emerged to enable models to adapt quickly to new tasks with minimal data, exemplified by methods that use hypernetworks as metaknowledge to dynamically generate model parameters from input data (Zhmoginov, Sandler, and Vladymyrov 2022). Bayesian deep learnig used Variational Bayes to infer the distributions of model parameters, which are assumed to be Gaussian (Wilson and Izmailov 2020).

Recently, diffusion models have emerged as a powerful paradigm for model parameter generation. G.pt (Peebles et al. 2022) trains conditional diffusion models using model checkpoints as datasets. MetaDiff (Zhang et al. 2024) introduces a diffusion-based meta-learning method for few-shot tasks. HyperDiffusion (Erkoç et al. 2023) uses diffusion models to generate new neural implicit fields. Additionally, p-diff (Wang et al. 2024) examines the quality of model parameters generated through diffusion. These works provide preliminary experiments and analyses on using diffusion models for model parameter generation.

## Problem Definition and Preliminaries

### Federated Learning Setting

FL aims to collectively train a centralized model for  $n$  edge-distributed clients. Each client  $i$  has access to  $m_i$  data samples  $D_i = (x_j^{(i)}, y_j^{(i)})_{j=1}^{m_i}$  from its own private data distribution  $P_i$  on  $\mathcal{X} \times \mathcal{Y}$  and  $N$  is the total number of data samples overall clients. Generally, the data distributions for each client are different, i.e.,  $P_i \neq P_j$  for any pair  $i, j \in 1, \dots, n$ . Let  $\ell_i : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{R}_+$  denote the loss function corresponding to client  $i$ , and  $F_i : \Theta \times \mathcal{X} \rightarrow \mathcal{Y}$  denote the local model parameterized by  $\theta^i$ . The goal of conventional FL is to optimize the following objectives:

$$\theta^* = \arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(x,y) \sim P_i} [\ell_i(F_i(\theta, x), y)], \quad (1)$$

where  $\theta^*$  represents the globally optimal parameters. The representative method for solving Eq. (1) is FedAvg (McMahan et al. 2017b). In each round, clients perform several epochs of SGD on their local loss functions and send the updated models to the server. The server then aggregates the models from all clients linearly by  $\bar{\theta} = \sum_{i=1}^n \frac{m_i}{N} \theta^i$  and broadcasts the averaged model  $\bar{\theta}$  back to the clients.

Moreover, due to data heterogeneity, the unified parameters may not be locally optimal for each client. Therefore, personalized FL adjusts the optimization objective to:

$$\Theta^* = \arg \min_{\Theta = \{\theta^i\}_{i=1}^n} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(x,y) \sim P_i} [\ell_i(F_i(\theta^i; x), y)]. \quad (2)$$

where  $\Theta^*$  denotes the collection of locally optimal parameters. The challenge lies in how to aggregate model parameters from heterogeneous clients on the server and produce parameters that incorporate insights from all clients while still adapting to specific data distributions. In this work, we propose to train a diffusion model at the server to address these challenges.

### Diffusion Probabilistic Models

Here, we focus on the Denoising Diffusion Probabilistic Models formulation, primarily because of its prevalence and consistency with the Denoising Score Matching using Langevin Dynamics. DDPMs gradually add noise to data, transforming it into standard Gaussian noise, and then learn to denoise step-by-step, generating new data. It involves a forward process and a reverse process as described below.

**Forward process.** Given training data  $z_0$  from a target distribution  $q(z_0)$ , the forward process gradually adds Gaussian noise to diffuse  $z_0$  into  $z_1, z_2, \dots, z_T$ , where  $z_T$  is approximately sampled from the standard Gaussian distribution  $q(z_T) \approx \mathcal{N}(z_T; \mathbf{0}, \mathbf{I})$ . This process can be formulated as follows:

$$q(z_{1:T}|z_0) = \prod_{t=1}^T q(z_t|z_{t-1}) \quad (3)$$

$$q(z_t|z_{t-1}) = \mathcal{N}(z_t; \sqrt{1 - \beta_t} z_{t-1}, \beta_t \mathbf{I}) \quad (4)$$

where  $\beta_1, \dots, \beta_T \in (0, 1)$  is a variance schedule. The diffusion model aims to approximate the noise added at each time step  $t$  using a neural estimator  $\epsilon_\phi(z_t, t)$ . The training objective  $L_{ddpm}$  is given by the Eq. (5) below:

$$L_{ddpm} = \mathbb{E}_{t \sim [1, T], z_0 \sim q_0, \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\epsilon_t - \epsilon_\phi(z_t, t)\|_2^2], \quad (5)$$

where  $z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t$  and  $\bar{\alpha}_t = \prod_{j=1}^t (1 - \beta_j)$ .

**Reverse process.** During the reverse process, we iteratively reconstruct the target data  $z_0$  from the random noise  $z_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . At each time step  $t$ , we sample  $z_{t-1}$  from the estimated reverse Markov chain  $p_\phi(z_t|z_{t-1})$  parameterized by the noise estimator  $\epsilon_\phi(z_t, t)$ . That is,

$$z_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left( z_t - \frac{\beta_t}{\sqrt{(1 - \bar{\alpha}_t)}} \epsilon_\phi(z_t, t) \right) + \sigma_t z_t. \quad (6)$$

where  $z_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $\sigma_t^2 = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$  represents the derived variance.

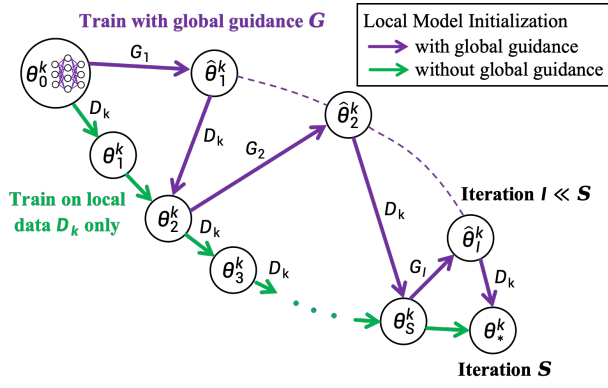


Figure 2: Illustration of the training process for a new client  $k$  with and without global guidance. The green arrows represent training starting from the initial parameters  $\theta_0^k$  solely on local data  $D_k$ , gradually converging to the final optimized parameters  $\theta_*^k$  within  $S$  iterations. The purple arrows indicate training with global guidance  $G$  alternated with local data training, which accelerates initialization and converges within  $I \ll S$  iterations.

**Analysis.** As demonstrated in (Song et al. 2020), DDPMs can be viewed as discretizations to SDEs, coherently bridging diffusion probabilistic modeling and noise conditional score networks into a unified framework. Thus, we can reformulate  $L_{ddpm}$  equivalently as denoising score matching:

$$\mathbb{E}_{q(z_0)} \mathbb{E}_{q(z_t|z_0)} [\|s_\phi(z_t, t) - \nabla_{z_t} \log q(z_t|z_0)\|_2^2], \quad (7)$$

where  $s_\phi(z_t, t) = -\frac{\epsilon_\phi(z_t, t)}{\sqrt{(1-\alpha_t)}}$ . This allows us to use DDPMs to effectively estimate the scores of data distributions.

## Methodology

In this section, we present our framework pFedGPA in detail. Given a collection of clients with their own specific data distributions in a FL system, our goal is to integrate the diverse parameter distributions captured by the diffusion model on the server, which reflects the underlying data distributions of the clients. Meanwhile, the diffusion model aims to generate new, personalized model parameters for each client by leveraging these learned parameter distributions. In the following subsections, we first illustrate the parameter aggregation using the diffusion model. Next, we explain the personalized parameter generation. Finally, we delve into the architecture design of the diffusion model. More details about the training procedure are presented in Algorithm 1.

### Generative Parameter Aggregation

In our framework, we deploy a diffusion model on the server, where both input and output are the weights of the clients' models. The core of this approach is to enable the server to learn the distribution of all clients' model parameters using a powerful generative model, a distribution which resides on a low-dimensional manifold in the parameter space.

Let  $H : \Phi \times \Theta \rightarrow \Theta$  denotes the diffusion model parametrized by  $\phi$  and we modify the optimization objective in Eq. (2) based on the parameter generation to derive:

$$\phi^*, \Theta^* = \arg \min_{\phi, \Theta := \{\theta^i\}_{i=1}^n} \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\phi, \theta^i), \quad (8)$$

where  $\mathcal{L}_i(\phi, \theta^i) = \mathbb{E}_{(x,y) \sim P_i} [\ell_i(F_i(H(\phi; \theta^i); x), y)]$  and  $\{\theta_*^i\}_{i=1}^n$  are the stationary points of the system  $H$  under well-fitted state  $\phi^*$  that represent the optimal model parameters of each client. In each round, the server receives the client parameters and treats them as training data for the diffusion model, training it over a predefined number of epochs. By performing parameter inversion, which will be introduced in the next section, the server then generates personalized model parameters for each client.

**Global Guidance.** To quickly initialize a newly joined client, we propose a global guidance approach to help the client model adapt its parameters to the local data in just a few iterations. In (Shamsian et al. 2021), a hypernetwork is used to take a trainable vector descriptor as input to differentiate each client. However, this approach has scalability limitations. Specifically, when new clients join the network, it often requires training new embedding vectors from scratch, which becomes increasingly inefficient as the network grows. Moreover, fine-tuning all clients' embedding vectors to achieve consistent representation is necessary but complex. In diffusion models, if the input also includes client descriptors  $e_{client}$ , it can be implemented as a form of classifier-free guidance, utilizing a linear combination of conditional and unconditional score estimates to provide client-specific guidance:

$$\tilde{\epsilon}_\phi(\theta, e_{client}) = (1 + \omega)\epsilon_\phi(\theta, e_{client}) - \omega\epsilon_\phi(\theta), \quad (9)$$

where  $\omega$  controls the strength of the guidance. The advantage is that there is no need to train a separate classifier. However, in our approach, we adopt classifier-based guidance because, in the FL setting, each client inherently serves as a classifier relative to the diffusion model. In turn, the diffusion model offers global guidance to the local client models. That is:

$$\tilde{\epsilon}_\phi(\theta) = \underbrace{\epsilon_\phi(\theta)}_{\text{global guidance}} - \underbrace{(1 + \omega)\nabla_\theta \log P_y(x|\theta)}_{\text{local update}}, \quad (10)$$

where  $\log P_y(x|\theta)$  represents the training loss of the client model in probabilistic form (e.g., cross-entropy loss). This setup allows the diffusion model to share information across clients while preserving the adaptability of personalized models. Instead of fitting each client's distribution individually, a unified model is used to estimate the overall score of model distributions, thereby facilitating efficient knowledge sharing.

For newly joined clients, this approach enables rapid initialization within just a few iterations. In each iteration, the client sends its fine-tuned model parameters to the server, where the diffusion model provides global guidance to update these parameters. This process is illustrated in Fig. 2.



Method	EMNIST			Fashion-MNIST			CIFAR-10		
	10 clients	20 clients	100 clients	10 clients	20 clients	100 clients	10 clients	20 clients	100 clients
Local-only	70.85	70.45	73.73	85.11	84.89	85.48	63.83	64.38	65.41
FedAvg	71.10	72.89	74.60	81.90	82.96	83.93	64.13	68.51	69.87
FedAvg-FT	81.91	<b>85.09</b>	86.52	88.45	<b>89.39</b>	89.28	71.76	76.46	76.55
FedPer	74.56	76.61	77.08	87.13	87.20	88.65	63.20	63.68	69.43
FedRep	74.22	74.37	76.54	88.69	88.27	88.42	64.71	65.85	67.63
LG-FedAvg	71.23	70.90	76.25	85.28	85.25	85.96	64.20	64.94	65.41
FedBABU	80.28	83.92	84.39	87.44	88.86	89.43	69.42	74.25	75.32
pFedHN	77.33	81.79	77.23	87.69	86.84	86.25	71.43	75.41	77.12
<b>pFedGPA</b>	<b>82.53</b>	<u>84.70</u>	<b>88.22</b>	<b>88.90</b>	<u>89.18</u>	<b>90.08</b>	<b>72.95</b>	<b>78.54</b>	<b>78.33</b>

Table 1: The comparison of final round average test accuracy (%) across different datasets: full participation with 10 and 20 clients, and client sampling at 30% with 100 clients in the FL system. The best performance is in **bold** font and the second best is marked with underline.

Additionally, we introduce noise augmentation into the input and latent representations to improve the robustness and generalization of the generated results. Experiments in p-diff show that this augmentation is crucial for producing stable models with high performance. Unlike generating model parameters trained on IID data in p-diff, our diffusion model is designed to fit model parameters from different data distributions in FL. For smaller models, we generate the entire model, while for larger models, we generate specific layers due to hardware constraints. The complete details of the model are provided in the Appendix.

## Experiments

### Setup

**Datasets and Local Models.** We conduct image classification tasks and evaluate our method on three widely-used benchmark datasets: Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), EMNIST (Cohen et al. 2017), and CIFAR-10 (Krizhevsky and Hinton 2009). A small CNN model for Fashion-MNIST and two larger CNN models for EMNIST and CIFAR-10 are constructed, respectively. Details of the datasets and model architectures are provided in the Appendix.

**Data Partitioning.** For the heterogeneous data distribution setting, we adopt the approach proposed in (Karimireddy et al. 2020b; Zhang et al. 2021), ensuring that all clients have equal data sizes. A portion of the data ( $s\%$ , default 20%) is uniformly sampled from all classes, while the remaining  $(100 - s)\%$  is sampled from a set of dominant classes specific to each client. Clients are grouped based on their dominant classes, though this grouping is unknown to the server. Additionally, the size of the local training data is kept small, specifically at 600 samples per client, to emphasize the necessity of FL. The testing data on each client is drawn from the same distribution as the training data.

**Compared Baselines.** We evaluate the performance of pFedGPA against the following baselines: Local-only, where each client trains its model locally without communication; FedAvg (McMahan et al. 2017a) and its locally fine-tuned variant (FedAvg-FT); pFedHN (Shamsian et al. 2021),

which leverages a server-side hypernetwork; and other personalized FL methods, including FedPer (Arivazhagan et al. 2019), FedRep (Collins et al. 2021), FedBABU (Oh, Kim, and Yun 2022), and LG-FedAvg (Liang et al. 2020). Hyperparameters for all baselines were tuned to their optimal values based on the settings reported in the original papers or determined through grid search.

**Training Details.** All local models are trained using mini-batch SGD as the optimizer, with 2 epochs per round and a batch size of 50. The number of global communication rounds is set to 200 for Fashion-MNIST and EMNIST, and 300 for CIFAR-10. For Fashion-MNIST, the entire model is generated, while for EMNIST and CIFAR-10, only the final two fully connected layers are generated. We report the average test accuracy across clients. For pFedGPA, we trained the diffusion models using parameters collected from the last 20 rounds, generating new parameters in the final round.

### Experimental Results

**Performance Comparison.** The results of average test accuracy across clients are reported in Table 1, where our pFedGPA is able to outperform other baselines in most cases. The improvement by our method increases as the learning task becomes harder. For example, our method improve the accuracy by  $\sim 2\%$  with 100 clients on the CIFAR-10. The performance gaps between FedPer/FedRep and FedAvg-FT/FedBABU on various cases indicate that only sharing the feature extractor during the FL process could have inverse affect, as diverse local classifier head could make the feature learning sub-optimal. By contrast, our framework leverages the powerful generation ability of diffusion models to facilitate the personalized model aggregation.

**Parameter Initialization.** The generative approach captures the distribution of model parameters across different clients and can provide parameter initialization for subsequent networks. The diffusion model trained within one FL network can subsequently serve as a pre-trained model for a newly initialized FL network, thus avoiding a cold start and accelerating the initialization process. To validate this idea,

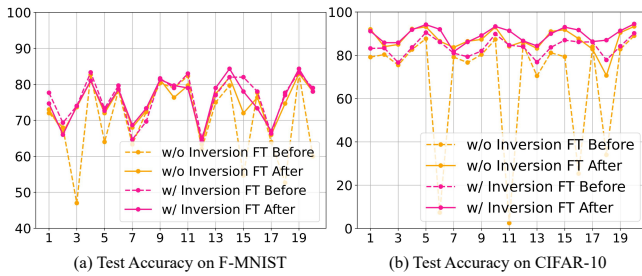


Figure 4: Comparison of test accuracies across 20 clients on Fashion-MNIST and CIFAR-10 using the pFedGPA method, before and after the fine-tuning (FT) operation, with and without parameter inversion. Dashed lines indicate results before FT, and solid lines indicate results after FT.

Dataset	w/o inversion		w/ inversion	
	before FT	after FT	before FT	after FT
EMNIST (%)	69.38	81.09	83.54	<b>84.70</b>
Fashion-MNIST (%)	68.48	87.24	83.82	<b>89.37</b>
CIFAR-10 (%)	68.9	74.77	<b>76.28</b>	75.72

Table 2: Ablation study on the effect of the parameter inversion mechanism in pFedGPA. **before FT** and **after FT** refer to the accuracy before and after fine-tuning on local data in the final round. **w/o inversion** denotes without the inversion mechanism, and **w/ inversion** denotes with the inversion mechanism applied.

we employed a diffusion model trained within one FL network to generate parameters for a newly initialized network. We then compared the number of communication rounds required by each method to reach 95% of its respective peak performance, measured as the average test accuracy. Experimental results, detailed in the appendix, demonstrate that our approach significantly reduces the number of communication rounds required to achieve high performance, compared to baseline methods.

**Ablation Study.** We conduct ablation studies to verify the efficacy of the inversion mechanism in pFedGPA. Specifically, we compare the performance of using inversion versus direct generation on the three datasets. We report the final round’s average accuracy on 20 clients, both before and after fine-tuning. As demonstrated in Table 2, the inversion mechanism significantly improve the average test accuracy before and after fine-tuning. Beyond the advantage of inversion in generating personalized parameters, we found that direct generation is less stable, often resulting in failed parameter generation. We compared the test accuracies across 20 clients on Fashion-MNIST and CIFAR-10 using the pFedGPA method, evaluating the performance both before and after the fine-tuning (FT) operation, with and without parameter inversion. In the results shown in Fig. 4, dashed lines represent the accuracies before fine-tuning, while solid lines show the accuracies after fine-tuning. In about 3 out of 20 cases, directly generated parameters had a pre-finetune accuracy below 60%. While increasing the size of the diffusion model reduced this failure rate, inversion re-

mained the more robust design, with failures occurring in only 1 out of 40 cases.

## Discussion

**Practical Considerations.** In our experiments, training a round of the diffusion model takes about an hour on a single Nvidia 4090 24GB GPU, with the entire FL process completing in four hours. However, the diffusion model’s training time significantly exceeds the time between communication rounds, causing a bottleneck in the overall process. In real-world applications, it is expected that the server would have more hardware resources, thereby enabling the training of larger diffusion models, which can in turn generate larger local models. Additionally, employing larger batch sizes could further accelerate the training process. In contrast, the inference time remains relatively short, enabling the trained diffusion model to quickly generate parameters for clients.

**Communication and Privacy Concerns.** The trained parameters of the diffusion model, which are much larger than those of the clients, are never transmitted. Consequently, we do not introduce any additional communication costs compared to other methods. Moreover, in our framework, only the local model parameters are exchanged between clients and the server, without transmitting additional local feature statistics, thereby maintaining the same privacy levels as other methods that exchange only parameters.

**Future Work.** Beyond the advantages demonstrated in our experiments, generative aggregation methods have the potential to be integrated with other FL approaches, such as prototype-enhanced algorithms (Tan et al. 2022; Xu, Tong, and Huang 2023). These methods can complement each other effectively: generative models learn the distribution of model parameters, while prototype learning focuses on the distribution of data representations. This integration could enable the development of novel inversion strategies that may lead to the generation of more accurate personalized parameters in the future.

## Conclusion

In this paper, we introduce a novel diffusion-based parameter aggregation method for personalized FL. In our framework, the server deploys a diffusion model to consolidate the uploaded parameters and generates personalized parameters for each client with the global guidance by using a newly developed inversion mechanism. Experimental results on three datasets verify the effectiveness and further advantages of our method.

## Acknowledgments

This work is supported in part by the Natural Science Foundation of China (Grant 62371270) and Shenzhen Key Laboratory of Ubiquitous Data Enabling (No.ZDSYS20220527171406015).

## References

- Amari, S.-i. 1993. Backpropagation and stochastic gradient descent method. *Neurocomputing*, 5(4-5): 185–196.
- Arivazhagan, M. G.; Aggarwal, V.; Singh, A. K.; and Choudhary, S. 2019. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*.
- Cohen, G.; Afshar, S.; Tapson, J.; and van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks, IJCNN 2017, Anchorage, AK, USA, May 14-19, 2017*, 2921–2926. IEEE.
- Collins, L.; Hassani, H.; Mokhtari, A.; and Shakkottai, S. 2021. Exploiting Shared Representations for Personalized Federated Learning. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 2089–2099. PMLR.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Erkoç, Z.; Ma, F.; Shan, Q.; Nießner, M.; and Dai, A. 2023. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14300–14310.
- Fallah, A.; Mokhtari, A.; and Ozdaglar, A. 2020. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in neural information processing systems*, 33: 3557–3568.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Hanzely, F.; and Richtárik, P. 2020. Federated learning of a mixture of global and local models. *arXiv preprint arXiv:2002.05516*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2): 1–210.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020a. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, 5132–5143. PMLR.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S. J.; Stich, S. U.; and Suresh, A. T. 2020b. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 5132–5143. PMLR.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Krizhevsky, A.; and Hinton, G. 2009. Learning Multiple Layers of Features from Tiny Images. *University of Toronto*.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1931–1941.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.
- Li, W.; Milletari, F.; Xu, D.; Rieke, N.; Hancox, J.; Zhu, W.; Baust, M.; Cheng, Y.; Ourselin, S.; Cardoso, M. J.; et al. 2019. Privacy-preserving federated brain tumour segmentation. In *Machine Learning in Medical Imaging: 10th International Workshop, MLMI 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13, 2019, Proceedings 10*, 133–141. Springer.
- Liang, P. P.; Liu, T.; Ziyin, L.; Allen, N. B.; Auerbach, R. P.; Brent, D.; Salakhutdinov, R.; and Morency, L.-P. 2020. Think locally, act globally: Federated learning with local and global representations. *arXiv preprint arXiv:2001.01523*.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11461–11471.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017a. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- McMahan, H. B.; Moore, E.; Ramage, D.; et al. 2017b. Communication-efficient learning of deep networks from decentralized data. *AISTATS*.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.
- Nguyen, D. C.; Pham, Q.-V.; Pathirana, P. N.; Ding, M.; Seneviratne, A.; Lin, Z.; Dobre, O.; and Hwang, W.-J. 2022. Federated learning for smart healthcare: A survey. *ACM Computing Surveys (Csur)*, 55(3): 1–37.
- Nichol, A.; Achiam, J.; and Schulman, J. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.
- Oh, J.; Kim, S.; and Yun, S. 2022. FedBABU: Toward Enhanced Representation for Federated Image Classification. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.

- Peebles, W.; Radosavovic, I.; Brooks, T.; Efros, A. A.; and Malik, J. 2022. Learning to learn with generative models of neural network checkpoints. *arXiv preprint arXiv:2209.12892*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.
- 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.
- Shamsian, A.; Navon, A.; Fetaya, E.; and Chechik, G. 2021. Personalized federated learning using hypernetworks. In *International Conference on Machine Learning*, 9489–9502. PMLR.
- Smith, V.; Chiang, C.-K.; Sanjabi, M.; and Talwalkar, A. S. 2017. Federated multi-task learning. *Advances in neural information processing systems*, 30.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Tan, Y.; Long, G.; Liu, L.; Zhou, T.; Lu, Q.; Jiang, J.; and Zhang, C. 2022. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8432–8440.
- Wang, K.; Xu, Z.; Zhou, Y.; Zang, Z.; Darrell, T.; Liu, Z.; and You, Y. 2024. Neural network diffusion. *arXiv preprint arXiv:2402.13144*.
- Wilson, A. G.; and Izmailov, P. 2020. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33: 4697–4708.
- Wu, C. H.; and De la Torre, F. 2023. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7378–7387.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *CoRR*, abs/1708.07747.
- Xu, J.; Tong, X.; and Huang, S.-L. 2023. Personalized Federated Learning with Feature Alignment and Classifier Collaboration. In *The Eleventh International Conference on Learning Representations*.
- Xu, J.; Yan, Y.; and Huang, S. 2022. FedPer++: Toward Improved Personalized Federated Learning on Heterogeneous and Imbalanced Data. In *IJCNN*.
- Xu, M.; Yu, L.; Song, Y.; Shi, C.; Ermon, S.; and Tang, J. 2022. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*.
- Yang, M.; Xu, J.; Ding, W.; and Liu, Y. 2023. Federated Capsule Graph Neural Network with Personalization. In *Adjunct Proceedings of the 2022 ACM International Joint Conference on Pervasive and Ubiquitous Computing and the 2022 ACM International Symposium on Wearable Computers, UbiComp/ISWC '22 Adjunct*.
- Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2): 1–19.
- Ye, R.; Ni, Z.; Wu, F.; Chen, S.; and Wang, Y. 2023. Personalized Federated Learning with Inferred Collaboration Graphs. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 39801–39817. PMLR.
- Zhang, B.; Luo, C.; Yu, D.; Li, X.; Lin, H.; Ye, Y.; and Zhang, B. 2024. Metadiff: Meta-learning with conditional diffusion for few-shot learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16687–16695.
- Zhang, J.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Guan, H. 2023. Fedala: Adaptive local aggregation for personalized federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11237–11244.
- Zhang, M.; Sapra, K.; Fidler, S.; Yeung, S.; and Alvarez, J. M. 2021. Personalized Federated Learning with First Order Model Optimization. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; and Chandra, V. 2018. Federated learning with non-iid data. *arXiv:1806.00582*.
- Zhmoginov, A.; Sandler, M.; and Vladymyrov, M. 2022. Hypertransformer: Model generation for supervised and semi-supervised few-shot learning. In *International Conference on Machine Learning*, 27075–27098. PMLR.
- Zhou, H.; Qixuan, Z.; Luo, T.; Zhang, Y.; and Xu, Z.-Q. 2022. Towards understanding the condensation of neural networks at initial training. *Advances in Neural Information Processing Systems*, 35: 2184–2196.
- Zhou, Z.; Zhou, H.; Li, Y.; and Xu, Z.-Q. J. 2023. Understanding the initial condensation of convolutional neural networks. *arXiv preprint arXiv:2305.09947*.