

Federated Unsupervised Domain Generalization Using Global and Local Alignment of Gradients

Farhad Pourpanah*, Mahdiyar Molahasani*, Milad Soltany*, Michael Greenspan, Ali Etemad

Queen’s University, Canada

{f.pourpanahnavan, m.molahasani, milad.soltany, michael.greenspan, ali.etemad}@queensu.ca

Abstract

We address the problem of federated domain generalization in an unsupervised setting for the first time. We first theoretically establish a connection between domain shift and alignment of gradients in unsupervised federated learning and show that aligning the gradients at both client and server levels can facilitate the generalization of the model to new (target) domains. Building on this insight, we propose a novel method named FedGaLA, which performs gradient alignment at the client level to encourage clients to learn domain-invariant features, as well as global gradient alignment at the server to obtain a more generalized aggregated model. To empirically evaluate our method, we perform various experiments on four commonly used multi-domain datasets, PACS, OfficeHome, DomainNet, and TerraInc. The results demonstrate the effectiveness of our method which outperforms comparable baselines. Ablation and sensitivity studies demonstrate the impact of different components and parameters in our approach.

Code — <https://github.com/MahdiyarMM/FedGaLA>

1 Introduction

Federated learning (McMahan et al. 2017; Zhang et al. 2021a) has emerged as a promising framework for training machine learning models across multiple decentralized clients while preserving data privacy. It allows clients to collaboratively train a global model without the need to exchange their sensitive and local data. Each client trains a local model using its data and a server aggregates these models at a certain frequency (Ghosh et al. 2020; Charles et al. 2021). However, given that each client collects a different set of local training data, two issues arise. First, the data collected by each client is often recorded under unique conditions that may result in mutual domain shifts (Liu et al. 2021). Second, labeling training data is inherently challenging and resource-intensive; this issue is even more pronounced in the context of federated settings. A typical example of this scenario is a network of wearable activity monitors where variations in user conditions such as demographics or ambient factors can lead to significant domain shifts across devices, meanwhile, the

*These authors contributed equally.

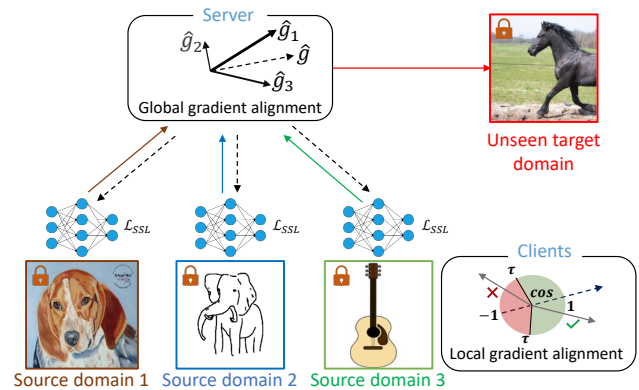


Figure 1: Overview of FedGaLA.

users are generally not asked to provide ground-truth labels for their performed activities.

In prior works, each of these issues has been addressed as a separate problem statement: (i) *federated domain generalization* (Zhang et al. 2021b; Nguyen, Torr, and Lim 2022; Bai, Bagchi, and Inouye 2024), and (ii) *federated unsupervised learning* (Zhuang et al. 2021; Han et al. 2022). Despite the effectiveness of both problem definitions, each ignores the fundamental assumptions of the other regarding the data in terms of distributions and availability of labels. To further approach federated learning in a more practical scenario, we propose to merge these two under a new umbrella called *federated unsupervised domain generalization*, which we define as Definition 1. To our knowledge, prior works have not studied federated learning under such constraints.

Definition 1. *Federated unsupervised domain generalization is the problem of learning general representations from various decentralized **unlabeled** datasets, each belonging to a **different domain**, in a federated setup where data sharing is restricted due to privacy concerns.*

To address this new problem, we propose a novel method called *Federated Unsupervised Domain Generalization using Global and Local Alignment of Gradients (FedGaLA)*. To learn more generalized representations from multiple domains, FedGaLA relies on gradient alignment at both client (local) and server (global) levels. At the client level, local

models are trained with unlabeled local data available to each client using self-supervised learning (SSL). To learn domain-invariant representations, gradients that are not aligned with the reference gradient, i.e., the global learning direction, are discarded. At the server level, to achieve better generalization, the local models are aggregated based on their alignment with each other (see Figure 1). Specifically, the local gradients that are more closely aligned with the average gradient are given greater weight in the aggregation stage. Moreover, we provide a detailed theoretical framework establishing a connection between the alignment of gradients from different clients and the similarity between their data distributions. Since FedGaLA does not require the sharing of data between clients to perform alignment at either the server or client sides, it effectively preserves data privacy. We verify the performance of our approach using four public datasets, PACS (Yu et al. 2022), OfficeHome (Venkateswara et al. 2017), DomainNet (Peng et al. 2019), and TerraInc (Beery, Van Horn, and Perona 2018) and demonstrate strong performance in federated unsupervised domain generalization in comparison to various baselines. We also conduct different ablation and sensitivity studies to understand the impact of different parameters and the choice of SSL frameworks.

In summary, the main contributions of this study are as follows. (1) To our knowledge, this is the first work to address federated unsupervised domain generalization, introducing a new problem class in this area. (2) We propose a novel technique that aligns gradients at both the local and global levels. In doing so, our solution effectively extracts domain-invariant information in local training and aligns each client’s contributions during the aggregation process, boosting the model’s generalization across various target domains. (3) We conduct extensive experiments and ablations to demonstrate the effectiveness of our proposed method on various benchmarks. To enable fast reproducibility and contribute to the area, we make our code public at: <https://github.com/MahdiyarMM/FedGaLA>.

2 Related Works

Federated Learning. Federating learning is a technique of distributed training that enables learning from decentralized clients without the need to transfer raw data to a central server (McMahan et al. 2017; Zhang et al. 2021a). In the first work in the area of federated learning, FedAVG (McMahan et al. 2017) aggregates the weights of clients via averaging to form the global model. Various federated learning techniques have been since developed for different purposes. FedProx (Li et al. 2020) adds a proximal term to the objective to tackle the data heterogeneity problem. FedBuf (Nguyen et al. 2022) proposes to use buffered asynchronous aggregation to address scalability and privacy issues. FedALA (Zhang et al. 2023b) dynamically aggregates the downloaded global model with the local model on each client, ensuring alignment with the local objective. Moon (Li, He, and Song 2021) leverages the similarity between model representations to improve the local training of individual parties through contrastive learning. FLTrust (Cao et al. 2020) and SignGuard (Xu et al. 2022) leverage gradient filtering to detect and eliminate malicious gradients with the aim of enhancing the robustness of

federated learning systems against model poisoning attacks. **Federated Domain Generalization.** This category of problems aims to learn a general model in a privacy-preserving manner that generalizes well to new target domains with distribution shifts (Liu et al. 2021; Li et al. 2023). There have been several studies to explore this direction, mainly focusing on learning domain-invariant features (Nguyen, Torr, and Lim 2022; Wu and Gong 2021) or identifying common features across multiple domains (Zhang et al. 2021b). FedDG (Liu et al. 2021) allows clients to share their data in the frequency space for medical image segmentation. FedSR (Nguyen, Torr, and Lim 2022) proposes to learn a simple representation of data to better generalize to target domains through two regularizers, i.e., L2-norm and conditional Mutual Information. FedADG (Zhang et al. 2021b) leverages adversarial training to align source domain distributions by matching each distribution with a reference distribution. FedKA (Sun, Chong, and Ochiai 2023) learns domain-invariant features by employing feature distribution matching in a universal workspace. CCST (Chen et al. 2023) aligns various client distributions and mitigates model biases by adapting local models to diverse sample styles via cross-client style transfer. FedDG-GA (Zhang et al. 2023c) uses domain divergence and a moment mechanism to enhance generalization through dynamic domain weight adjustment.

FedSB (Soltany et al. 2024) enhances domain generalization by introducing controlled uncertainty to local clients and ensuring balanced contributions from each client. StableFDG (Park et al. 2024) uses style and attention-based strategies to address the federated domain generalization problem. hFedF (Bartholet et al. 2024) utilizes hypernetworks for non-linear aggregation to facilitate generalization to unseen domains. PerAda (Xie et al. 2024) uses knowledge distillation to align each client’s adapter with a global adapter, blending their knowledge. This approach helps the model handle shifting distributions across diverse domains while keeping computational costs low. gPerXAN (Le et al. 2024) combines a personalized normalization scheme with a guiding regularizer to address domain shifts in federated learning, reducing privacy risks and communication costs. MCGDM (Wei and Han 2024) uses intra- and inter-domain gradient matching to minimize domain-specific overfitting, promoting robust generalization on unseen domains. Finally, FGPP (Wan, Huang, and Ye 2024) captures domain information using clustered prototypes and contrastive learning.

Federated Unsupervised Learning. This category attempts to learn representations from unlabeled data distributed across clients while preserving data privacy (Jin et al. 2020). The study by (van Berlo, Saeed, and Ozcelebi 2020) is the first to propose federated unsupervised learning using an encoder-decoder structure. FedU (Zhuang et al. 2021) employs a contrastive learning approach with online and target networks to enable each client to learn representations from unlabeled data independently. It also introduces a dynamic aggregation mechanism to update the predictor, either locally or globally. Similarly, FedEMA (Zhuang, Wen, and Zhang 2022) uses the exponential moving average of the global model to update local ones. FedX (Han et al. 2022) proposes to use knowledge distillation to learn representation

from local data and refine the central server’s knowledge. FedCA (Zhang et al. 2023a) consists of two components: (i) a dictionary module to gather the representations of samples from each client to maintain consistency in the representation space across all clients, and (ii) an alignment module to adjust each client’s representation to match a base model trained on public data.

3 Approach

3.1 Problem Formulation

To formalize Definition 1, assume K clients, C_i , in a federated setup, each with its own *unlabeled* data $\mathcal{D}_i = \{\mathbf{x}_i^{(n)}\}_{n=1}^{N_i}$. Each dataset consists of N_i data points sampled from a distinct data distribution $p(\mathbf{x}_i)$, where \mathbf{x}_i is a vector of F features, i.e., $\mathbf{x}_i = [x_i^1, x_i^2, \dots, x_i^F]^T$. The data distributions are assumed to be different among the clients with each distribution $p(\mathbf{x}_i)$ sampled from a family of distributions \mathcal{P} . Privacy constraints prevent the transfer of data between clients or to the server S . The objective is to learn generalized representations from \mathbf{x}_i that perform well across unseen distributions $p(\mathbf{x}_i) \sim \mathcal{P}$, where $p(\mathbf{x}_i) \neq p(\mathbf{x}_i)$. This is formulated as minimizing the expected loss over the unseen distributions:

$$\min_{\theta} \mathbb{E}_{p(\mathbf{x}_i) \sim \mathcal{P}} [\mathbb{E}_{p(\mathbf{x}_i)} [\mathcal{L}(\theta; \mathbf{x})]], \quad (1)$$

where \mathcal{L} is the unsupervised loss function, and θ is the set of global model parameters. Each client contributes to this goal by computing a local objective function approximating the expected loss with respect to its own data distribution:

$$\min_{\theta_i} \mathbb{E}_{p(\mathbf{x}_i)} [l_i(\theta_i; \mathbf{x})] \approx \frac{1}{N_i} \sum_{n=1}^{N_i} \mathcal{L}(\theta_i; \mathbf{x}_i^{(n)}), \quad (2)$$

where θ_i indicates the local parameters of client C_i , and θ is the global aggregation of all local θ_i .

3.2 Gradient Alignment and Domain Shift

Under a federated learning framework, privacy constraints prevent clients and servers from accessing each other’s data, including distribution information such as data means and variances. They can, however, observe individual client gradients at the server level and the average aggregated gradient across clients at the client level. We motivate our work on the fact that alignment of gradients may infer characteristics of the client domain distributions, thus facilitating improved model generalization. While empirically the utility of gradients has been demonstrated in the area of domain generalization (Mansilla et al. 2021), no theoretical basis has been proposed for this approach under federated constraints. Accordingly, we aim to establish a link between gradient alignment and domain shifts within the proposed problem formulation (Definition 1). Our theoretical findings provide the basis for effective local parameter updates and global model aggregation to address federated unsupervised domain generalization.

Assumption 1. Let each \mathbf{x}_i^f be a random variable drawn from a Normal distribution $p(\mathbf{x}_i^f) \sim \mathcal{P}$ (Bar-Hillel et al. 2003). Within a single domain, following (Štrumbelj and

Kononenko 2014, 2011), features are assumed to be independent ($\text{Cov}(x_i^{f_1}, x_i^{f_2}) = 0$). Across different domains, corresponding features of \mathbf{x}_i^f and \mathbf{x}_j^f are bivariate with a covariance of $\sigma_{x_i^f, x_j^f}$. In line with contemporary practices in deep learning, the features are normalized with $\mu = 0$ and $\sigma^2 = 1$ (Yu and Spiliopoulos 2022; Qi, Zhou, and Wang 2022; Huang et al. 2023). We also assume that each client is a one-layer encoder with the sigmoid activation function. Following (Chen et al. 2020), local training is performed using contrastive loss where a random augmentation of a data sample is used as the positive pair and all other samples are utilized as the negative pairs. The random augmentation is performed with the same random Affine transformation ($Ax_i + B$) (Wang et al. 2023) broadcast to all clients. After each epoch, the local models are aggregated in the server and sent back to all clients. The gradient \mathbf{g}_i of the model is assumed to be differentiable.

Theorem 1 (Gradient Misalignment in Federated Self-supervised Learning Dependent upon Domain Shift). *Given Assumption 1, under the problem proposed in Definition 1, for two distinct domains characterized by random variables \mathbf{x}_i and \mathbf{x}_j belonging to two different clients C_i and C_j , an increase in domain shift across the clients results in a decrease in covariance $\text{Cov}(\mathbf{g}_i, \mathbf{g}_j)$ of the corresponding gradients $\mathbf{g}_i, \mathbf{g}_j$ across C_i and C_j ’s respective local models.*

Proof of Theorem 1. Motivated by previous works where the similarity between the representations of different domains under domain shift is measured by Mutual Information (Gao et al. 2020; Menapace, Lathuilière, and Ricci 2020; Wang et al. 2021), we use this concept for modeling the similarity between different domains drawn from a family of distributions. To calculate the mutual information we introduce the following Lemma.

Lemma 1. *Given Assumption 1, the mutual information between two random variables \mathbf{x}_i and \mathbf{x}_j can be calculated as:*

$$I(\mathbf{x}_i; \mathbf{x}_j) = -\frac{1}{2} \sum_{f=1}^F \log(1 - \sigma_{x_i^f, x_j^f}^2). \quad (3)$$

The full proof of this lemma is presented in Appendix A.1.

Since $\sigma_{x_i^f, x_i^f} = \sigma_{x_i^f}^2$, given identical and standardized domains we have $\sigma_{x_i^f, x_j^f} = 1$. Therefore, we observe from Eq. (3) that as the shift between the domain approaches zero, the Mutual Information approaches infinity. On the other hand, as the two domains shift apart, $\sigma_{x_i^f, x_j^f}$ approaches zero, and consequently the Mutual information monotonically decreases toward zero. Accordingly, $I(\mathbf{x}_i; \mathbf{x}_j)$ and $\sigma_{x_i^f, x_j^f}$ are positively and monotonically correlated. To establish the link between the covariance of the features and the variance of the gradient and demonstrate their relationship, the following lemma and claim are introduced.

Lemma 2. *Given Assumption 1, the covariance between the differentiable function \mathbf{g} with inputs \mathbf{x}_i and \mathbf{x}_j can be*

estimated as:

$$\text{Cov}(\mathbf{g}_i, \mathbf{g}_j)_{mn} \approx \sum_{f=1}^F \sigma_{x_i^f, x_j^f} \left(\left. \frac{\partial g_{i(m)}}{\partial x_i^f} \right|_{\mu_i^f} \right) \left(\left. \frac{\partial g_{j(n)}}{\partial x_j^f} \right|_{\mu_j^f} \right), \quad (4)$$

where $\sigma_{x_i^f, x_j^f}$ is the covariance between the f^{th} feature of x_i and x_j and $g_{i(m)}$ is the m^{th} dimension of g_i . The proof for this Lemma is provided through the Taylor expansion derived in Appendix A.2.

Claim 1. For all clients trained under the federated unsupervised domain generalization setup using self-supervised learning described in Assumption 1, for two clients C_i and C_j , for both positive and negative contrastive data pairs, the sign of $(\left. \frac{\partial \mathbf{g}_i}{\partial x_i^f} \right|_{x_i^f = \mu_i^f}) (\left. \frac{\partial \mathbf{g}_j}{\partial x_j^f} \right|_{x_j^f = \mu_j^f})$ is always positive in all dimensions. See proof in Appendix A.3.

According to Eq. (4) introduced in Lemma 2 and Claim 1, we conclude that $\text{Cov}(\mathbf{g}_i, \mathbf{g}_j)$ and $\sigma_{x_i^f, x_j^f}$ are positively and monotonically correlated. Therefore, $\text{Cov}(\mathbf{g}_i, \mathbf{g}_j)$ and $I(\mathbf{x}_i; \mathbf{x}_j)$ are also positively and monotonically correlated, which completes the proof. \square

Our theoretical findings may also be applicable to Federated Supervised Domain Generalization, which could form the subject of future research as expressed in Appendix A.4 and (Molahasani et al. 2024). We further investigate the relationship between the distribution of the domains and their corresponding gradients through the following Corollary.

Corollary 1. Given the assumptions stated, for two distinct domains characterized by random variables \mathbf{x}_i and \mathbf{x}_j from the distribution family \mathcal{P} , as $I(\mathbf{x}_i, \mathbf{x}_j)$ decreases, then the variance of the difference of the corresponding gradients, $\text{Var}(\mathbf{g}_i - \mathbf{g}_j)$, increases. The full proof is presented in Appendix A.6.

3.3 FedGaLA

The analysis above establishes a link between gradient alignment and domain shift, forming the basis for our proposed FedGaLA framework. Figure 1 illustrates an overview of our framework for addressing the newly proposed problem setup (Definition 1). FedGaLA is remotely inspired by prior works that demonstrate improved generalization in centralized (non-federated) learning through gradient alignment (Mansilla et al. 2021; Shi et al. 2022; Rame, Dancette, and Cord 2022). However, our framework extends this notion by integrating gradient alignment into the field of unsupervised and federated learning, and assumes, unlike (Mansilla et al. 2021; Shi et al. 2022; Rame, Dancette, and Cord 2022), that the distribution of data from different clients is not known. Our core idea includes (i) enabling clients to learn domain-invariant representations at the client level through local gradient alignment, and (ii) adjusting the aggregation weights at the server level using global gradient alignment.

Algorithm 1: FedGaLA

```

1: Input: data  $\mathcal{D}_i$ , initialization  $\Theta^0$ 
2: Output:  $\Theta^T$ 
3: for  $t$  from 1 to  $T$  do
4:   Server:
5:     Calculate client updates:  $\mathbf{g}_i^{(t)} \leftarrow \Theta_i^{(t)} - \Theta^{(t-1)}$ 
6:     Initialize global update:  $\mathbf{g}^{(t+1)} \leftarrow \text{FedAVG}(\mathbf{g}_i^{(t)})$ 
7:     for  $j$  from 1 to  $iter$  do
8:       Calculate weights:  $w_i \leftarrow \frac{\text{Cosine}(\mathbf{g}_i^{(t)}, \mathbf{g}^{(t+1)})+1}{2}$ 
9:       Normalize weights:  $w_i \leftarrow \frac{w_i}{\sum w_i}$ 
10:      Aggregate updates:  $\hat{\mathbf{g}}^{(t+1)} \leftarrow \sum_{i=1}^K w_i \mathbf{g}_i^{(t)}$ 
11:    end for
12:    Update global model:  $\Theta^{(t+1)} \leftarrow \Theta^{(t)} + \hat{\mathbf{g}}^{(t+1)}$ 
13:    Communicate:  $\Theta_i \leftarrow \Theta^{(t+1)}$ 
14:   Client:
15:     for  $j$  from 1 to  $N_{\text{batch}}$  do
16:       Compute batch gradient:  $\mathbf{g}_{i,j}^{(t)} \leftarrow \nabla l_i(x_{i,j}, \theta_i^{(t)})$ 
17:       for  $l$  from 1 to  $L$  do
18:         Estimate reference:  $\hat{\mathbf{g}}_{est}^{(l,t)} = \theta^{(l,t)} - \theta^{(l,t-1)}$ 
19:         if  $\text{Cosine}(\mathbf{g}_{i,j}^{(l,t)}, \hat{\mathbf{g}}_{est}^{(l,t)}) > \tau$  then
20:           Update weights:
21:              $\theta_i^{(l,t)} \leftarrow \theta_i^{(l,t-1)} - \eta \cdot \nabla l_i^{(l,t)}(x_{i,j}, \theta_i^{(l,t)})$ 
22:         end if
23:       end for
24:     end for
25:     Communicate:  $S \leftarrow \Theta_i$ 
26:   end for

```

At each communication round, clients are initialized with the global model. Subsequently, each client updates its parameters using SSL for E epochs based on the local data through local gradient alignment and sends these updates back to the server. Finally, the server employs the global gradient alignment technique to perform aggregation. This procedure is repeated for T communications rounds to determine the global model. The remaining part of this section provides details on local and global gradient alignment, and the complete framework is outlined in Algorithm 1.

Local Gradient Alignment. Our method performs layer-wise local gradient alignment using a reference gradient. This reference is derived from the l^{th} layer of the global model’s parameter updates between the current and the previous communication round. Suppose $\Theta = \{\theta^{(l)}\}_{l=1}^L$ indicates the parameters of the global model, where $\theta^{(l)}$ represents the parameters of the l^{th} layer. The reference gradient for the l^{th} layer is computed as $\hat{\mathbf{g}}_{est}^{(l,t)} = \theta^{(l,t)} - \theta^{(l,t-1)}$, where $\theta^{(l,t)}$ and $\theta^{(l,t-1)}$ are the parameters of the l^{th} layer of the global model at rounds t and $t-1$, respectively. Then, $\hat{\mathbf{g}}_{est}^{(l,t)}$ is locally used to determine whether the gradient of each layer obtained during training (e.g., via SGD) is aligned with the reference. The cosine similarity between the batch gradient and the reference for each layer l at round t is computed as $\text{cos}(\phi)^{(l,t)} = \frac{\langle \mathbf{g}_{i,k}^{(l,t)}, \hat{\mathbf{g}}_{est}^{(l,t)} \rangle}{\|\mathbf{g}_{i,k}^{(l,t)}\| \cdot \|\hat{\mathbf{g}}_{est}^{(l,t)}\|}$, where $\mathbf{g}_{i,k}^{(l,t)}$ is the gradient of k^{th} batch of the i^{th} client at layer l and round t . Finally, batch gradients whose similarity with $\hat{\mathbf{g}}_{est}^{(l,t)}$ are less than a user-defined threshold τ are discarded during the update pro-

cess. This prevents clients from learning domain-specific features by disregarding local gradients that are not aligned with the global model. The rationale for discarding unaligned gradients instead of applying soft weighting is that when a gradient vector is unaligned with $\hat{\mathbf{g}}_{est}$, scaling does not change this alignment, as the cosine of the angle between them is independent of scale. To establish a theoretical basis for the proposed local alignment, we introduce the following.

Proposition 1. *Given two sets of gradient vectors \mathbf{g}_i and \mathbf{g}_j , by removing the K^{th} vector in \mathbf{g}_j where $\cos(\mathbf{g}_{j,K}, \mathbf{g}_{est}) < 0$, the covariance of two sets increases. For more details on this Proposition and its proof, see Appendix A.8.*

Theorem 1 highlighted how an increase in domain shift between clients C_i and C_j correlates with a decrease in the covariance of their gradient vectors, \mathbf{g}_i and \mathbf{g}_j . Therefore, gradient covariance can be used as an indicator of domain shift. Proposition 1 complements this by showing that by selectively removing gradients from \mathbf{g}_j that have a negative cosine similarity with an estimated target direction, \mathbf{g}_{est} , we can effectively increase the covariance of the gradient sets, thus potentially counteracting the effects of domain shift.

Global Gradient Alignment. To aggregate the local (client) models at the server, the locally measured gradients that closely match the average gradient across all clients are assigned greater weights. This soft weighting process operates as follows. Once the server receives the local models, it first obtains the local update $\hat{\mathbf{g}}_i^{(t)} = \Theta_i^{(t)} - \Theta^{(t-1)}$. Then, the initial global update $\hat{\mathbf{g}}^{(t+1)}$ is calculated by averaging all $\hat{\mathbf{g}}_i^{(t)}$. Subsequently, the weight of each client is computed using $w_i = \frac{\cos(\hat{\mathbf{g}}_i^{(t)}, \hat{\mathbf{g}}^{(t+1)}) + 1}{2}$, where w_i is the weight for the i^{th} client and $\hat{\mathbf{g}}_i^{(t)}$ is the gradient of the i^{th} client at round t . This weight reflects the degree of alignment between each client’s update and the global model’s update direction. To ensure these weights are properly normalized using $w_i = w_i / \sum_{k=1}^K w_k$. The normalization of weights allows for the proportional contribution of each client’s update.

Finally, the normalized weights are used to perform aggregation. Each client’s model update is scaled by its respective weight, and these weighted updates are then aggregated to compute the weighted average update. The global model at round $t + 1$ is updated based on $\hat{\mathbf{g}}^{(t+1)} = \sum_{i=1}^K w_i \hat{\mathbf{g}}_i^{(t)}$. This aggregation step is repeated three times, refining the weights with each iteration. This ensures the global model update is significantly influenced by clients whose updates align closely with the global learning objective. After completing the aggregation process, the global model is further updated based on $\Theta^{(t+1)} \leftarrow \Theta^{(t)} + \hat{\mathbf{g}}^{(t+1)}$. The rationale for using different alignment strategies at the client and server levels stems from the inherent differences between local training and global aggregation. At the client level, we manage batch gradients, allowing us to specifically discard the unaligned ones without significant information loss. However, discarding gradients at the server-level corresponds to the deletion of entire clients, which can adversely affect the outcome.

4 Experiment Setup

Datasets. To evaluate the effectiveness of our proposed method, we conduct experiments across four commonly used benchmarks for domain generalization. They include: **PACS** (Li et al. 2017), which consists of 9,991 images from four domains: ‘Photo’, ‘Art-painting’, ‘Cartoon’, and ‘Sketch’, across seven classes; **Office-Home** (Venkateswara et al. 2017), which consists of 15,588 images from four domains: ‘Art’, ‘Clipart’, ‘Product’, and ‘Real-world’, across 65 classes; **TerraInc** (Beery, Van Horn, and Perona 2018), which includes 24,788 images from four domains ‘Location 38’, ‘Location 43’, ‘Location 46’, and ‘Location 100’, across nine classes; **DomainNet** (Peng et al. 2019), which consists of 569,010 images in six domains: ‘Clipart’, ‘Infograph’, ‘Painting’, ‘Quickdraw’, ‘Real’, and ‘Sketch’, covering 345 classes. Following (Zhang et al. 2022), for the DomainNet dataset, we select the following classes: zigzag, tiger, tornado, flower, giraffe, toaster, hexagon, watermelon, grass, hamburger, blueberry, violin, fish, sun, broccoli, Eiffel tower, horse, train, bird, and bee (Zhang et al. 2022), resulting in a total of 38556 samples. In all experiments using this dataset, three domains are used for training (‘Painting’, ‘Real’, and ‘Sketch’) and the other three domains (‘Clipart’, ‘Infographics’, and ‘Quickdraw’) are used for testing.

Evaluation. We use the leave-one-domain-out setting used in prior works (Zhang et al. 2023a, 2022; Gulrajani and Lopez-Paz 2021). This involves selecting one domain as the target, training the model on the rest of the domains, and then testing the model’s performance on the selected target domain. Linear evaluation, a common feature evaluation approach, is utilized to evaluate the quality of learned representations (Feng, Xu, and Tao 2019; Zhang, Isola, and Efros 2017; Kolesnikov, Zhai, and Beyer 2019). For linear evaluation, following (van Berlo, Saeed, and Ozcelebi 2020; Zhuang, Wen, and Zhang 2022), we utilize 10% and 30% of the target data to train the linear classifier and evaluate the remaining 90% and 70% of the data, respectively.

Baselines. To evaluate our method, we take a two-pronged approach: (1) We adapt several popular SSL approaches to the federated domain generalization task, denoting them as FedSimCLR, FedMoCo, FedBYOL, and FedSimSiam. To this end, we employed SimCLR (Chen et al. 2020), MoCo (He et al. 2020), BYOL (Grill et al. 2020), and SimSiam (Chen and He 2021) in a federated setup. Each of these models consists of two encoders. BYOL and MoCo utilize exponential moving averages to update one of the encoders (target network) using the other encoder (online network). In contrast, SimCLR and SimSiam share weights between their two encoders. Additionally, SimCLR and MoCo, being contrastive-based SSL models, leverage negative samples in the learning process. Conversely, BYOL and SimSiam, which do not rely on negative samples, include a predictor on top of the online encoder to enhance learning from unlabeled samples. We then train each client locally using the respective SSL method. Next, we aggregate the trained encoders at the server using FedAVG (McMahan et al. 2017). For BYOL and SimSiam, we follow the procedure in (Zhuang et al. 2021) and apply FedAVG on the online encoder and projector. We also adapt FedEMA (Zhuang, Wen, and Zhang 2022),

Model	Labeled ratio	PACS					DomainNet				
		P	A	C	S	Ave.	C	I	Q	Ave.	
FedEMA	10%	50.0(0.7)	29.5(1.9)	42.4(2.3)	45.6(0.16)	41.9	38.6(1.1)	13.7(0.5)	45.5(1.6)	32.4	
FedBYOL		52.1(1.1)	31.8(1.1)	45.4(2.2)	47.4(1.9)	44.2	38.1(0.5)	14.1(0.4)	53.6(2.9)	31.8	
FedMoCo		58.5(2.2)	35.7(9.6)	37.7(12.9)	36.6(8.2)	42.1	30.5(0.6)	10.9(2.1)	46.4(0.8)	27.2	
FedSimSiam		46.2(1.1)	28.6(1.0)	46.7(0.6)	37.6(1.3)	39.8	44.8(1.5)	12.2(0.3)	40.3(2.4)	36.9	
FedSimCLR		64.2(1.2)	41.9(1.5)	58.4(1.3)	70.1(1.2)	58.6	45.2(0.4)	13.7(0.3)	59.7(0.7)	39.5	
FedGaLA (ours)		64.7(1.9)	44.2(1.2)	60.5(2.2)	70.5(1.3)	60.0	47.6(0.9)	14.2(0.5)	61.4(0.3)	41.1	
FedEMA	30%	53.5(0.3)	33.9(2.8)	48.3(2.4)	45.3(2.3)	45.2	43.5(0.8)	20.0(1.5)	49.0(2.8)	37.5	
FedBYOL		55.1(1.2)	35.3(1.5)	48.3(1.5)	48.9(0.5)	46.9	44.6(1.2)	20.5(1.1)	50.4(0.9)	38.5	
FedSimSiam		58.5(2.2)	35.7(9.6)	37.7(12.9)	36.6(8.2)	42.1	35.7(1.1)	14.9(3.9)	44.6(3.5)	31.8	
FedMoCo		47.3(1.6)	30.2(3.4)	47.4(0.7)	34.4(1.7)	39.9	51.5(1.8)	18.2(0.1)	59.9(3.6)	43.2	
FedSimCLR		69.8(1.1)	46.4(2.1)	63.9(1.6)	73.4(3.0)	63.3	51.7(1.0)	16.3(0.2)	66.5(0.9)	44.8	
FedGaLA (ours)		71.1(2.0)	46.8(2.1)	65.7(1.6)	74.5(1.1)	64.6	52.4(0.7)	16.2(0.6)	68.8(0.6)	45.8	

Model	Labeled ratio	Office-Home					TerraInc				
		A	C	P	R	Ave.	L38	L43	L46	L100	Ave.
FedEMA	10%	6.7(0.5)	12.5(0.4)	20.8(0.9)	14.1(0.8)	13.5	62.1(0.4)	43.5(3.0)	43.4(0.4)	63.9(0.7)	53.2
FedBYOL		7.4(0.3)	12.9(0.5)	20.9(1.1)	13.8(0.1)	13.8	63.4(0.2)	43.3(0.5)	44.3(0.3)	66.3(2.4)	54.3
FedSimSiam		8.5(0.5)	19.8(0.8)	28.2(0.8)	16.2(0.6)	18.9	54.0(6.8)	36.0(1.5)	40.9(3.7)	60.8(1.6)	47.9
FedMoCo		10.8(0.4)	9.4(0.3)	12.4(0.9)	10.1(0.7)	10.7	50.0(2.4)	33.8(0.3)	29.7(2.3)	60.2(0.1)	45.7
FedSimCLR		8.9(0.4)	24.3(0.3)	35.2(1.2)	20.0(0.2)	22.0	62.8(0.2)	45.8(1.1)	42.9(1.8)	68.8(0.3)	55.1
FedGaLA (ours)		8.9(0.4)	25.3(0.6)	36.6(0.3)	21.2(0.5)	23.0	63.6(0.1)	47.6(1.4)	43.9(2.9)	71.7(0.9)	56.7
FedEMA	30%	10.0(0.3)	17.3(0.7)	28.2(1.1)	17.8(0.3)	18.2	62.7(0.6)	47.1(0.3)	46.1(0.4)	67.6(2.0)	55.9
FedBYOL		10.4(0.3)	17.4(1.2)	29.4(1.3)	17.6(0.6)	18.7	63.6(0.9)	46.7(1.3)	46.1(0.2)	68.8(2.5)	56.3
FedSimSiam		12.7(1.1)	24.1(1.9)	36.5(0.7)	21.6(1.1)	23.7	56.8(4.6)	35.7(2.0)	38.9(4.6)	61.4(1.4)	48.2
FedMoCo		15.2(0.6)	11.2(0.9)	15.8(0.7)	11.5(1.6)	13.4	60.7(0.1)	28.3(3.3)	31.7(2.3)	60.0(0.6)	45.2
FedSimCLR		13.6(1.3)	35.3(0.7)	47.6(0.9)	26.6(1.0)	30.7	65.8(0.3)	54.5(1.2)	49.0(0.4)	71.0(1.6)	60.1
FedGaLA (ours)		13.6(0.6)	35.1(0.2)	48.0(1.3)	27.0(0.7)	30.9	66.2(0.5)	52.4(1.7)	51.7(0.8)	74.6(1.3)	61.3

Table 1: Results of linear eval. on PACS, DomainNet, Office-Home, and TerraInc datasets.

which is a commonly used method originally developed for federated unsupervised learning. FedEMA integrates BYOL as an SSL technique into its structure.

It is important to note that we avoid direct comparisons with FedDG (Liu et al. 2021), FedSR (Nguyen, Torr, and Lim 2022), FedADG (Zhang et al. 2021b), FedIIR (Guo et al. 2023), and FedDG-GA (Zhang et al. 2023c) since unlike our method, they employ the label information in their solutions. (2) Given the absence of prior research specifically addressing the problem of *federated* unsupervised domain generalization, we also compare FedGaLA to established *centralized unsupervised domain generalization* methods on the PACS and DomainNet datasets. We could not identify unsupervised domain generalization methods for Office-Home and TerraInc datasets. This comparison includes the following solutions: SimCLR (Chen et al. 2020), MoCo (He et al. 2020), BYOL (Grill et al. 2020), AdCo (Hu et al. 2021), and DARLING (Zhang et al. 2022). It is important to note that these methods are implemented in a non-federated environment and do not incorporate any data privacy constraints. All the results for these models are reported from (Zhang et al. 2022).

Implementation Details. We use SimCLR as the SSL module in FedGaLA due to its performance on domain generalization problems as previously shown (Zhang et al. 2022). Following (Zhang et al. 2022), ResNet-18 (He et al. 2016) is employed as the encoder network architecture for all experiments, which we train from scratch. We present details regarding data augmentations, projector architecture, and en-

coder hyperparameters in Appendix B. Following (Feng, Xu, and Tao 2019; Zhang, Isola, and Efros 2017), we first learn a representation by FedGaLA and the baseline models for 100 communication rounds with 7 local epochs. Next, we freeze the backbone model and train a linear classifier for 100 epochs to perform prediction on the target domain. Appendix B presents all the hyperparameters used for FedGaLA. For FedEMA, we use the hyperparameters reported in (Zhuang, Wen, and Zhang 2022). All experiments were implemented using PyTorch and trained on 8 NVIDIA GeForce RTX 3090 GPUs. For each experiment, we train the models three times with random initialization seeds and report the average.

Cross-Silo vs. Cross-Device Federated Learning. Two frameworks can generally be considered in federated learning: cross-silo and cross-device. In cross-silo, the number of clients K is small, and each client is capable of retaining its local model and state across training rounds. In contrast, cross-device federated learning works under the assumption of having millions of stateless clients, where only a limited number of clients participate in training during each communication round. However, due to limitations in experimental settings, the majority of studies typically conduct experiments involving at most hundreds of clients (Wang et al. 2020; Zhuang, Wen, and Zhang 2022). In this study, given the limited number of domains FedGaLA works within the cross-silo framework. Moreover, each client is equipped with the necessary resources to participate in the collaborative learning process. Additionally, we ensure that every client updates its local model with a sufficient number of epochs.




Figure 2: (a) Covariance under domain shift across various domains of the PACS dataset. (b) Average ratio of discarded local gradients over 100 communication rounds ($\tau = 0$ and $E = 1$). (c) Impact of various batch sizes. (d) Impact of τ on performance across different number of local epochs. Results in (b-d) are reported for the domain P of the PACS dataset, with a 10% label ratio.

Model	PACS					DomainNet			
	P	A	C	S	Ave.	C	I	Q	Ave.
BYOL	27.0	25.9	21.0	19.7	23.4	14.6	8.7	5.9	9.7
MoCo	44.2	25.9	33.5	25.0	32.1	32.5	18.5	8.1	19.7
SimCLR	54.7	37.7	46.0	28.3	41.6	37.1	19.9	12.3	23.1
AdCo	46.5	30.2	31.5	22.9	32.8	32.3	17.9	11.6	20.6
DARLING	53.4	39.9	46.4	30.2	42.5	35.2	20.9	15.7	23.9
FedGaLA (ours)	64.7	44.2	60.5	70.5	60.0	47.6	14.2	61.4	41.1

Table 2: Comparison of linear eval. (10%) results with non-federated (centralized) domain generalization methods on PACS and DomainNet datasets.

5 Results

In this section, we first present the performance evaluation of FedGaLA across various datasets, comparing it against baseline methods to demonstrate its effectiveness. This is followed by ablation studies and sensitivity analyses, which provide detailed insights into the contributions of individual components and the robustness of our proposed framework.

5.1 Performance

Federated Unsupervised Domain Generalization. We report the accuracy rates of FedGaLA and baseline models on the four datasets. As shown in Table 1, FedGaLA consistently outperforms all baselines across all four datasets, with the exception of the ‘Art-painting’ domain in Office-Home, for the 10% data regime. When 30% of the data are used, our method still generally outperforms the baseline models, although, for some of the domains, the baseline solutions produce slightly better results. This observation is expected given that with the introduction of more domain-specific training data, the need for domain generalization declines, and thus, methods that are not explicitly designed for domain generalization can produce competitive results. Across the four datasets, we observe that among the baseline models, FedSimCLR achieves better results compared to FedSimSiam, FedBYOL, and FedMoCo. This finding is consistent with (Zhang et al. 2022) where it was demonstrated that SimCLR provides a

Model	P	A	C	S	Ave.
FedMoCo	46.2(1.1)	28.6(1.0)	46.7(0.6)	37.6(1.3)	39.8
FedGaLA w/ MoCo	46.5(0.3)	28.9(0.4)	47.9(0.3)	39.6(2.4)	40.7
FedBYOL	52.1(1.1)	31.8(1.1)	45.4(2.2)	47.4(1.9)	44.2
FedGaLA w/ BYOL	52.8(0.6)	31.7(0.5)	46.3(3.1)	47.6(1.2)	44.6
FedSimCLR	64.2(1.2)	41.9(1.5)	58.4(1.3)	70.1(1.2)	58.6
FedGaLA w/ SimCLR	64.7(1.9)	44.2(1.2)	60.5(2.2)	70.5(1.3)	60.0

Table 3: The performance of FedGaLA when employing different SSL methods.

Model	P	A	C	S	Ave.
FedGaLA	64.7(1.9)	44.2(1.2)	60.5(2.2)	70.5(1.3)	60.0
w/o GA	64.7(0.4)	42.6(1.2)	58.1(0.6)	69.9(1.1)	58.8
w/o LA	63.7(1.5)	41.6(1.4)	59.8(1.5)	68.9(1.1)	58.5
w/o GA & LA	64.2(1.2)	41.9(1.5)	58.4(1.3)	70.1(1.2)	58.6

Table 4: Ablation study on PACS (GA: global alignment; LA: local alignment).

better foundation for domain generalization versus other SSL methods, albeit in a non-federated setup.

Centralized Unsupervised Domain Generalization. We compare the performance of FedGaLA with centralized (non-federated) methods for PACS and DomainNet datasets, where the baselines are trained on the entire dataset consisting of all the domains. The results presented in Table 2 show that FedGaLA outperforms centralized methods by large margins. This is an expected observation as prior works have shown that federation can boost domain generalization (Nguyen, Torr, and Lim 2022; Arasteh et al. 2023).

FedGaLA With Other SSL Techniques. Recall that we used SimCLR in our federated framework. To further explore whether FedGaLA can improve the performance of other federated SSL methods, we perform additional experiments in which we employ MoCo and BYOL in FedGaLA instead of SimCLR. We evaluate the performance on the PACS dataset, as presented in Table 3. The results show that FedGaLA consistently improves the performance of all federated SSL




Figure 3: (a) Impact of the number of iterations for global alignment. (b) Stability of the training. (c) Effect of different labeled data ratios on linear evaluation performance. (d) Performance of the model in different communication rounds. Results are reported for domain P of PACS dataset with 10% label ratio, except for (c) where the label ratio is changing.

Model	Communication Frequency (E)			
	$E = 1$	$E = 5$	$E = 7$	$E = 9$
FedGaLA	67.8	65.2	64.8	64.7

Table 5: Impact of communication frequency (E) on model accuracy over 900 total local epochs.

frameworks, demonstrating its ability to adapt to different SSL methodologies while boosting domain generalization.

5.2 Empirical Verification of Theorem 1

Gradient Misalignment Due to Domain Shift. In Figure 2a we illustrate the amount of measured covariance for different domains in PACS versus the amount of domain shift between each pair measured through Mutual Information. We observe that, except for a single outlier, the trend follows our prediction based on Theorem 1.

5.3 Ablation Studies and Component Analysis

Ablation Studies. Here we examine the effectiveness of the local and global gradient alignment components individually on the final performance of FedGaLA. To this end, we systematically remove each of these components. As shown in Table 4, each component plays an important role in the overall performance. It is noteworthy to mention that FedGaLA essentially becomes the FedSimCLR baseline by removing both global and local alignments.

Ratio of Discarded Local Gradients. Figure 2b demonstrates the ratio of local gradients discarded due to local gradient alignment during training versus communication rounds. For this experiment, the local models are trained for 1 epoch at each communication round. As observed, the ratio of discarded local gradients decreases from approximately 68% in the early communication rounds to 37% by round 100. This trend indicates that as the number of communication rounds increases, the local gradient directions become more aligned with the global model, suggesting that with increasing the number of communications, local models learn more domain-invariant features.

Models	P	A	C	S	Ave.
RF: 0.001	64.3(1.9)	41.8(1.2)	58.1(2.3)	68.3(1.1)	58.1
RF: 0.01	63.8(1.4)	42.0(1.7)	58.1(0.4)	68.3(1.6)	58.6
RF: 0.1	64.1(1.7)	40.9(2.1)	58.5(1.4)	70.4(0.4)	58.5
FedGaLA (ours)	64.7(1.9)	44.2(1.2)	60.5(2.2)	70.5(1.3)	60.0

Table 6: Results on PACS dataset for different reweight-factors and our FedGaLA method. ‘RF’ indicates the reweight factor.

5.4 Hyperparameter Sensitivity and Stability

Batch Size. To identify the optimal training hyperparameters, we conduct a grid search on various batch sizes, testing values of 32, 64, 128, and 256, as illustrated in Figure 2c. The results indicate that FedGaLA achieves its best performance with a batch size of 128, which balances efficient gradient updates and stability during training.

Local Threshold vs. the Number of Local Epochs. In Figure 2d we study the impact of the threshold for the similarity between gradients and the reference (τ) and the number of local epochs E on performance. In this experiment, we use three different values for τ : -0.1 , 0 , and 0.1 , and three different values for E : 5 , 7 or 9 (local epochs per communication round). We observe that our method produces the best results when $\tau = 0$, i.e., when we keep all gradients with positive cosine similarity with the reference. Expectedly, even discarding gradients with small amounts of alignment ($\tau = 0.1$) degrades the results, while keeping gradients that are not aligned with the reference ($\tau = -0.1$) also hurts performance. Moreover, we see that setting $E = 7$ yields the best performance as increasing the number of epochs beyond 7 does not have a positive impact, and only increases computational time.

Global Gradient Alignment Iterations. Figure 3a shows the impact of different gradient alignment iterations (line 7 in Algorithm 1) on performance. FedGaLA produces the best results when the iteration number is 3. When the number of iterations increases, the model places too much emphasis on clients aligned with the global average and decreases the

Model	PACS					DomainNet				
	P	A	C	S	Ave.	C	I	Q	Ave.	
FedGaLA+L2 ($\lambda = 0.001$)	64.8(0.5)	42.8(1.1)	58.7(1.6)	67.3(0.6)	58.4	46.3(0.8)	16.4(0.9)	62.8(0.2)	41.6	
FedGaLA+L2 ($\lambda = 0.01$)	66.5(0.9)	43.4(0.9)	56.9(2.7)	65.2(0.6)	57.9	38.9(1.2)	14.6(0.7)	52.5(0.8)	35.3	
FedGaLA+FedProx ($\mu = 0.001$)	62.1(0.2)	41.4(0.7)	58.9(3.1)	68.8(0.1)	57.8	44.6(0.2)	13.4(1.0)	60.9(1.1)	39.6	
FedGaLA+FedProx ($\mu = 0.01$)	63.1(0.2)	41.6(1.2)	58.5(2.0)	68.9(1.9)	58.0	44.7(1.2)	13.2(0.4)	61.7(0.4)	39.9	
FedGaLA (ours)	64.7(1.9)	44.2(1.2)	60.5(2.2)	70.5(1.3)	60.0	47.6(0.9)	14.2(0.5)	61.4(0.3)	41.1	

Model	Office-Home					TerraInc				
	A	C	P	R	Ave.	L38	L43	L46	L100	Ave.
FedGaLA+L2 ($\lambda = 0.001$)	11.8(0.4)	23.9(0.5)	37.1(0.4)	21.9(0.7)	23.6	61.6(0.9)	45.8(0.6)	46.6(0.7)	70.8(0.7)	56.2
FedGaLA+L2 ($\lambda = 0.01$)	11.8(0.2)	21.2(0.1)	34.5(0.3)	22.1(0.2)	22.4	61.5(0.7)	38.1(2.9)	44.9(0.1)	57.8(3.3)	50.6
FedGaLA+FedProx ($\mu = 0.001$)	10.2(0.4)	24.2(0.4)	36.4(1.1)	19.7(0.32)	22.6	63.8(0.7)	49.3(1.5)	47.2(0.1)	71.5(0.9)	57.9
FedGaLA+FedProx ($\mu = 0.01$)	9.2(0.3)	23.6(0.7)	37.1(0.6)	20.4(0.4)	22.6	63.2(1.1)	47.8(2.1)	44.8(0.4)	71.4(0.8)	56.0
FedGaLA (ours)	8.9(0.4)	25.3(0.6)	36.6(0.3)	21.2(0.5)	23.0	63.6(0.1)	47.6(1.4)	43.9(2.9)	71.7(0.9)	56.7

Table 7: The effect of regularizers on the performance of FedGaLA across PACS, DomainNet, Office-Home, and TerraInc.

impact of unaligned clients harshly, leading to a decline in the overall performance of the model. On the other hand, with small iteration values, global gradient alignment is not performed effectively, and thus performance declines.

Training Stability. Figure 3b demonstrates that FedGaLA maintains robust training stability, rapidly improving accuracy in the initial iterations and stabilizing after approximately 40 iterations. Despite discarding dissimilar gradients, the method avoids divergence and ensures consistent optimization, effectively learning domain-invariant features while maintaining steady performance throughout training.

5.5 Data and Communication Efficiency

Ratio of Labeled Data. Figure 3c presents the results when evaluated with different ratios of labeled data, ranging from 10% to 60%. As can be seen, the accuracy increases significantly from 64% to 72% when the label ratio rises from 10% to 30%, followed by a steady climb to 74% as the label ratios increase to 60%. Overall, when more labeled data are used to train the linear classifier, the final performance expectedly improves.

Communication Rounds. We investigate the effect of the communication rounds T on the model’s performance. Following (Zhuang et al. 2021; Zhuang, Wen, and Zhang 2022) we set $E = 1$ and vary the number of communication rounds T from 100 to 900. We observe from Figure 3d that performance improves significantly from 1 to 200 communication rounds, with this trend slowing down from 200 to 900.

FedGaLA and Communication Cost. We conduct an experiment to analyze the trade-off between communication frequency and model performance. In this setup, local models are trained for a total of 900 epochs, with aggregation occurring after every E local epochs. As E increases, communication frequency decreases, thereby reducing the overhead. The performance of the model is then assessed for different values of E , as demonstrated in Table 5. The results indicate that increasing the communication frequency (lower E) improves accuracy and the model maintains consistent performance even with reduced communication, particularly for $E > 1$. This indicates the ability of the model to adapt to less frequent aggregation intervals.

5.6 Model Variations

Gradient Discarding vs. Gradient Re-Weighting. Gradient re-weighting is an alternative approach to discarding unaligned gradients. To further highlight the effectiveness of gradient discarding in the local alignment module of FedGaLA, we conducted an additional experiment replacing gradient discarding with gradient re-weighting in FedGaLA. Here, rather than discarding the unaligned gradients, they are multiplied by a re-weight factor reducing their impact. The results are presented in Table 6, demonstrating that while re-weighting is effective, our original method of discarding unaligned gradients yields the best overall performance.

Effects of Regularizers on FedGaLA. Prior works have demonstrated that adding regularizers can indeed improve generalization across domains or non-IID data (Li et al. 2020; Nguyen, Torr, and Lim 2022). To this end, we test the impact of regularizers on FedGaLA by applying two types of regularizers based on L2 norm (Li et al. 2020) and FedProx (Nguyen, Torr, and Lim 2022). Please refer to Appendix B.4 for more details regarding these two techniques. The results in Table 7 demonstrate that FedGaLA is highly compatible with regularizers and that the addition of such approaches can further boost the performance of our method.

6 Conclusion

In this work, we first introduced a new problem category, federated unsupervised domain generalization, and subsequently proposed FedGaLA to tackle it. Building the relationship between gradient alignment and domain shift, FedGaLA comprises two alignment strategies at the global and local levels designed to address the problem of domain generalization in an unsupervised and federated setup. We assessed the performance of FedGaLA through extensive experiments where our approach outperformed various baseline models. Detailed ablation studies and sensitivity analyses were also conducted to provide more insights into different aspects of our solution.

Acknowledgments

We would like to thank Geotab Inc., the City of Kingston, and NSERC for their support of this work.

References

- Arasteh, S. T.; Kuhl, C.; Saehn, M.-J.; Isfort, P.; Truhn, D.; and Nebelung, S. 2023. Mind the Gap: Federated Learning Broadens Domain Generalization in Diagnostic AI Models. *arXiv:2310.00757*.
- Bai, R.; Bagchi, S.; and Inouye, D. I. 2024. Benchmarking Algorithms for Federated Domain Generalization. In *ICLR*.
- Bar-Hillel, A.; Hertz, T.; Shental, N.; and Weinshall, D. 2003. Learning distance functions using equivalence relations. In *ICML*, 11–18.
- Bartholet, M.; Kim, T.; Beuret, A.; Yun, S.-Y.; and Buhmann, J. M. 2024. Non-linear Fusion in Federated Learning: A Hypernetwork Approach to Federated Domain Generalization. *arXiv:2402.06974*.
- Beery, S.; Van Horn, G.; and Perona, P. 2018. Recognition in terra incognita. In *ECCV*, 456–473.
- Cao, X.; Fang, M.; Liu, J.; and Gong, N. Z. 2020. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *arXiv:2012.13995*.
- Charles, Z.; Garrett, Z.; Huo, Z.; Shmulyian, S.; and Smith, V. 2021. On large-cohort training for federated learning. *NeurIPS*, 34: 20461–20475.
- Chen, J.; Jiang, M.; Dou, Q.; and Chen, Q. 2023. Federated domain generalization for image recognition via cross-client style transfer. In *WACV*, 361–370.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *ICLR*, 1597–1607.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *CVPR*, 15750–15758.
- Feng, Z.; Xu, C.; and Tao, D. 2019. Self-supervised representation learning from multi-domain data. In *ICCV*, 3245–3255.
- Gao, J.; Hua, Y.; Hu, G.; Wang, C.; and Robertson, N. M. 2020. Reducing distributional uncertainty by mutual information maximisation and transferable feature learning. In *ECCV*, 587–605.
- Ghosh, A.; Chung, J.; Yin, D.; and Ramchandran, K. 2020. An efficient framework for clustered federated learning. In *NeurIPS*, volume 33, 19586–19597.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *NeurIPS*, 33: 21271–21284.
- Gulrajani, I.; and Lopez-Paz, D. 2021. In Search of Lost Domain Generalization. In *ICLR*.
- Guo, Y.; Guo, K.; Cao, X.; Wu, T.; and Chang, Y. 2023. Out-of-distribution generalization of federated learning via implicit invariant relationships. In *ICML*, 11905–11933. PMLR.
- Han, S.; Park, S.; Wu, F.; Kim, S.; Wu, C.; Xie, X.; and Cha, M. 2022. FedX: Unsupervised federated learning with cross knowledge distillation. In *ECCV*, 691–707.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hu, Q.; Wang, X.; Hu, W.; and Qi, G.-J. 2021. Adco: Adversarial contrast for efficient learning of unsupervised representations from self-trained negative adversaries. In *CVPR*, 1074–1083.
- Huang, L.; Qin, J.; Zhou, Y.; Zhu, F.; Liu, L.; and Shao, L. 2023. Normalization techniques in training DNNs: Methodology, analysis and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Jin, Y.; Wei, X.; Liu, Y.; and Yang, Q. 2020. Towards utilizing unlabeled data in federated learning: A survey and prospective. *arXiv:2002.11545*.
- Kolesnikov, A.; Zhai, X.; and Beyer, L. 2019. Revisiting self-supervised visual representation learning. In *CVPR*, 1920–1929.
- Le, K.; Ho, L.; Do, C.; Le-Phuoc, D.; and Wong, K.-S. 2024. Efficiently Assemble Normalization Layers and Regularization for Federated Domain Generalization. In *CVPR*, 6027–6036.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *ICCV*, 5542–5550.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *CVPR*, 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, Y.; Wang, X.; Zeng, R.; Donta, P. K.; Murturi, I.; Huang, M.; and Dustdar, S. 2023. Federated Domain Generalization: A Survey. *arXiv:2306.01334*.
- Liu, Q.; Chen, C.; Qin, J.; Dou, Q.; and Heng, P.-A. 2021. FedDG: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *CVPR*, 1013–1023.
- Mansilla, L.; Echeveste, R.; Milone, D. H.; and Ferrante, E. 2021. Domain generalization via gradient surgery. In *ICCV*, 6630–6638.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 1273–1282.
- Menapace, W.; Lathuilière, S.; and Ricci, E. 2020. Learning to Cluster Under Domain Shift. In *ECCV*, 736–752.
- Molahasani, M.; Soltany, M.; Pourpanah, F.; Greenspan, M.; and Etemad, A. 2024. A Theoretical Framework for Federated Domain Generalization with Gradient Alignment. In *NeurIPS Workshop of Mathematics of Modern Machine Learning*.
- Nguyen, A. T.; Torr, P.; and Lim, S. N. 2022. FedSR: A simple and effective domain generalization method for federated learning. *NeurIPS*, 35: 38831–38843.
- Nguyen, J.; Malik, K.; Zhan, H.; Yousefpour, A.; Rabbat, M.; Malek, M.; and Huba, D. 2022. Federated learning with buffered asynchronous aggregation. In *AISTATS*, 3581–3607.

- Park, J.; Han, D.-J.; Kim, J.; Wang, S.; Brinton, C.; and Moon, J. 2024. StableFDG: style and attention based learning for federated domain generalization. *NeurIPS*, 36.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment matching for multi-source domain adaptation. In *ICCV*, 1406–1415.
- Qi, H.; Zhou, J.; and Wang, H. 2022. A note on factor normalization for deep neural network models. *Scientific Reports*, 12(1): 5909.
- Rame, A.; Dancette, C.; and Cord, M. 2022. Fishr: Invariant Gradient Variances for Out-of-distribution Generalization. In *ICLR*.
- Shi, Y.; Seely, J.; Torr, P.; N, S.; Hannun, A.; Usunier, N.; and Synnaeve, G. 2022. Gradient Matching for Domain Generalization. In *ICLR*.
- Soltany, M.; Pourpanah, F.; Molahasani, M.; Greenspan, M.; and Etemad, A. 2024. Federated Domain Generalization with Label Smoothing and Balanced Decentralized Training. *arXiv preprint arXiv:2412.11408*.
- Štrumbelj, E.; and Kononenko, I. 2011. A general method for visualizing and explaining black-box regression models. In *International Conference on Adaptive and Natural Computing Algorithms*, 21–30.
- Štrumbelj, E.; and Kononenko, I. 2014. Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41: 647–665.
- Sun, Y.; Chong, N.; and Ochiai, H. 2023. Feature distribution matching for federated domain generalization. In *Asian Conference on Machine Learning*, 942–957.
- van Berlo, B.; Saeed, A.; and Ozcelebi, T. 2020. Towards federated unsupervised representation learning. In *Proceedings of the ACM International Workshop on Edge Systems, Analytics and Networking*, 31–36.
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 5018–5027.
- Wan, G.; Huang, W.; and Ye, M. 2024. Federated graph learning under domain shift with generalizable prototypes. In *AAAI*, 14, 15429–15437.
- Wang, H.; Yurochkin, M.; Sun, Y.; Papailiopoulos, D.; and Khazaeni, Y. 2020. Federated Learning with Matched Averaging. In *ICLR*.
- Wang, Z.; Luo, Y.; Qiu, R.; Huang, Z.; and Baktashmotlagh, M. 2021. Learning to diversify for single domain generalization. In *ICCV*, 834–843.
- Wang, Z.; Wu, C.; Yang, Y.; and Li, Z. 2023. Learning Transformation-Predictive Representations for Detection and Description of Local Features. In *CVPR*, 11464–11473.
- Wei, Y.; and Han, Y. 2024. Multi-Source Collaborative Gradient Discrepancy Minimization for Federated Domain Generalization. In *AAAI*, 14, 15805–15813.
- Wu, G.; and Gong, S. 2021. Collaborative optimization and aggregation for decentralized domain generalization and adaptation. In *ICCV*, 6484–6493.
- Xie, C.; Huang, D.-A.; Chu, W.; Xu, D.; Xiao, C.; Li, B.; and Anandkumar, A. 2024. PerAda: Parameter-Efficient Federated Learning Personalization with Generalization Guarantees. In *CVPR*, 23838–23848.
- Xu, J.; Huang, S.-L.; Song, L.; and Lan, T. 2022. Byzantine-robust federated learning through collaborative malicious gradient filtering. In *IEEE International Conference on Distributed Computing Systems*, 1223–1235.
- Yu, J.; and Spiliopoulos, K. 2022. Normalization effects on deep neural networks. *arXiv preprint arXiv:2209.01018*.
- Yu, S.; Wu, P.; Liang, P. P.; Salakhutdinov, R.; and Morency, L.-P. 2022. PACS: A dataset for physical audiovisual commonsense reasoning. In *ECCV*, 292–309.
- Zhang, C.; Xie, Y.; Bai, H.; Yu, B.; Li, W.; and Gao, Y. 2021a. A survey on federated learning. *Knowledge-Based Systems*, 216: 106775.
- Zhang, F.; Kuang, K.; Chen, L.; You, Z.; Shen, T.; Xiao, J.; Zhang, Y.; Wu, C.; Wu, F.; Zhuang, Y.; et al. 2023a. Federated unsupervised representation learning. *Frontiers of Information Technology & Electronic Engineering*, 24(8): 1181–1193.
- Zhang, J.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Guan, H. 2023b. FedALA: Adaptive local aggregation for personalized federated learning. In *AAAI*, volume 37, 11237–11244.
- Zhang, L.; Lei, X.; Shi, Y.; Huang, H.; and Chen, C. 2021b. Federated learning with domain generalization. *arXiv:2111.10487*.
- Zhang, R.; Isola, P.; and Efros, A. A. 2017. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 1058–1067.
- Zhang, R.; Xu, Q.; Yao, J.; Zhang, Y.; Tian, Q.; and Wang, Y. 2023c. Federated domain generalization with generalization adjustment. In *CVPR*, 3954–3963.
- Zhang, X.; Zhou, L.; Xu, R.; Cui, P.; Shen, Z.; and Liu, H. 2022. Towards unsupervised domain generalization. In *CVPR*, 4910–4920.
- Zhuang, W.; Gan, X.; Wen, Y.; Zhang, S.; and Yi, S. 2021. Collaborative unsupervised visual representation learning from decentralized data. In *ICCV*, 4912–4921.
- Zhuang, W.; Wen, Y.; and Zhang, S. 2022. Divergence-aware Federated Self-Supervised Learning. In *ICLR*.