

# Multi-Modal Style Transfer-based Prompt Tuning for Efficient Federated Domain Generalization

Yuliang Chen<sup>1,2\*</sup>, Xi Lin<sup>1,2\*</sup>, Jun Wu<sup>1,2</sup>, Xiangrui Cai<sup>3†</sup>, Qiaolun Zhang<sup>4</sup>,  
Xichun Fan<sup>5</sup>, Jiapeng Xu<sup>1,2</sup>, Xiu Su<sup>6†</sup>

<sup>1</sup>School of Computer Science, Shanghai Jiao Tong University

<sup>2</sup>Shanghai Key Laboratory of Integrated Administration Technologies for Information Security

<sup>3</sup>College of Computer Science, Nankai University

<sup>4</sup>Department of Electronics, Information and Bioengineering, Polytechnic Institute of Milan

<sup>5</sup>New York University Shanghai

<sup>6</sup>Big Data Institute, Central South University

{chenyuliang, linxi234, junwuhn, xjp20021018}@sjtu.edu.cn, caixr@nankai.edu.cn, qiaolun.zhang@mail.polimi.it, xf731@nyu.edu, xiusu1994@csu.edu.cn

## Abstract

Federated Domain Generalization (FDG) aims to collaboratively train a global model across distributed clients that can generalize well on unseen domains. However, existing FDG methods typically struggle with cross-client data heterogeneity and incur significant communication and computation overhead. To address these challenges, this paper presents a new FDG framework, dubbed FaST-PT, which facilitates local feature augmentation and efficient unseen domain adaptation in a distributed manner. First, we propose a lightweight Multi-Modal Style Transfer (MST) method to transform image embedding under text supervision, which could expand the training data distribution and mitigate domain shift. We then design a dual-prompt module that decomposes the prompt into global and domain prompts. Specifically, global prompts capture general knowledge from augmented embedding across clients, while domain prompts capture domain-specific knowledge from local data. Besides, Domain-aware Prompt Generation (DPG) is introduced to adaptively generate suitable prompts for each sample, which facilitates unseen domain adaptation through knowledge fusion. Extensive experiments on four cross-domain benchmark datasets, e.g., PACS and DomainNet, demonstrate the superior performance of FaST-PT over SOTA FDG methods such as FedDG-GA and DiPrompt. Ablation studies further validate the effectiveness and efficiency of FaST-PT.

## Introduction

Federated Learning (FL) (McMahan et al. 2017) is a widely adopted distributed machine learning paradigm, which facilitates collaborative model training while protecting privacy. Instead of relying on a single centralized server, FL distributes the training process across multiple clients (e.g., mobile devices or institutions) to collaboratively train a global model without sharing their private data. It is specifically

well-suited for real-world applications where sensitive data is distributed across multiple sources (Guan et al. 2024).

Although FL has shown promising performance in various tasks, existing studies predominantly focus on enhancing model performance for the participating clients. However, they often overlook the generalization capability to unseen domains beyond the participating clients. Data encountered in these unseen environments can exhibit substantial distributional differences from training clients, due to variations in data sources, collection methods, or domain-specific characteristics. Effectively adapting FL to such domain shifts (Muandet, Balduzzi, and Schölkopf 2013; Li et al. 2018) remains a technically challenging problem.

Federated Domain Generalization is defined as training a global model on multiple distributed source domains, aiming to achieve robust generalization to completely unseen target domains. Early attempts have been made to address the FDG problem. For example, ELCFS (Liu et al. 2021) allows clients to share the amplitude spectrum in frequency space for federated medical image segmentation. Similarly, CCST (Chen et al. 2023) extracts and exchanges style information to perform style transfer. Besides, FPL (Huang et al. 2023) facilitates the exchange of prototypes among clients. A common limitation of these methods is their reliance on sharing domain-related information among clients, which compromises the data privacy and communication constraints inherent in FL. Additionally, extracting domain-specific information often requires auxiliary modules, leading to increased computational overhead.

To address the above challenges, we propose a new framework dubbed Federated Multi-Modal Style Transfer with Prompt Tuning (FaST-PT). In FaST-PT, we introduce CLIP into FL and utilize its text-image alignment and generalization ability for handling FDG problems. On the one hand, we leverage CLIP’s powerful image-text alignment capabilities and propose Multi-Modal Style Transfer (MST), which enables client-side external domain style transfer to expand the training data distribution. Specifically, MST introduces a

\*These authors contributed equally.

†Corresponding authors

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

transform network that transforms image embeddings using the text representation of the external domain, facilitating the style transfer while retaining the class-specific information of the original image. On the other hand, we employ the prompt tuning paradigm of CLIP to facilitate efficient unseen domain adaptation. We first design a dual-prompt module that incorporates both global and domain prompts. In detail, global prompts are trained on augmented embeddings and aggregated across clients to capture general knowledge, while local prompts are trained solely on local data to preserve domain-specific knowledge. Additionally, we propose a domain-aware prompt generation mechanism to further incorporate domain knowledge based on domain similarity, thereby enhancing the adaptability of prompts to unseen domains. Our contribution can be summarized as follows:

- We present FaST-PT, a new FL framework that tackles the FDG problem through both client-side distribution alignment and efficient unseen domain adaptation by leveraging the text-image alignment and generalization capabilities of the pre-trained CLIP model.
- We propose a multi-modal style transfer method that allows clients to transform the image embedding to the style of the external domains while retaining the class information under the supervision of text description to mitigate the distribution discrepancy among clients.
- We design a dual-prompt module that leverages prompt partition to simplify feature disentanglement in traditional domain generalization, and employs domain-aware prompt generation to integrate domain knowledge for efficient unseen domain adaptation.
- Extensive experiments are conducted with FaST-PT on four benchmark datasets, consistently demonstrating its superior performance over baseline methods. Comprehensive ablation studies further validate the effectiveness and efficiency of FaST-PT.

## Related Work

### Federated Learning

Federated learning (FL) distributes the training progress across multiple devices or compute nodes instead of a single central server to facilitate global training of the model while protecting data privacy. The most typical FL framework is FedAvg (McMahan et al. 2017; Konečný et al. 2016), which decomposes the distributed learning process from the perspectives of local training and global averaging. However, traditional FL performs poorly under domain shift. Existing methods mainly focus on regularizing the local training on the client side to attain a well-performed global model. FedProx (Li et al. 2020) added an extra term in the loss function to control the  $L_2$  distance of model weights. FedBN (Li et al. 2021) retains the parameters of the batch normalization layers locally and excludes them from the aggregation process.

### Federated Domain Generalization

Federated domain generalization is a variant of DG tailored to the FL setting, aiming to mitigate domain shifts across

clients and improve the model’s generalization to unseen domains (Liu et al. 2021; Nguyen, Torr, and Lim 2022; Wu and Gong 2021; Chen et al. 2023; Zhang et al. 2021a; Huang et al. 2023; Le et al. 2024). ELCFS (Liu et al. 2021) is the first to solve the FDG problem, which exchanges the amplitude spectrum in the frequency domain and employs episodic learning to further enhance generalization performance. CCST (Chen et al. 2023) allows clients to share their style information and achieve local style transfer, leading to more uniform distributions of source clients. Apart from domain information sharing methods, the remaining approaches focus on extracting domain-invariant features. For example, FedSR (Nguyen, Torr, and Lim 2022) leverages L2-norm and conditional mutual information regularizations to learn simple yet generalizable data representations. COPA (Wu and Gong 2021) simultaneously incorporates a domain-invariant representation extractor along with an ensemble of domain-specific classifiers.

### Prompt Tuning

Prompt tuning (Lester, Al-Rfou, and Constant 2021) is a parameter-efficient fine-tuning technique that introduces a small set of learnable prompt tokens as additional inputs, while keeping the parameters of foundation models frozen. As a pioneering work, CoOp (Zhou et al. 2022b) adapts CLIP (Radford et al. 2021) to downstream image recognition tasks by converting the input context tokens in the text encoder into learnable prompts. CoCoOp (Zhou et al. 2022a) extends CoOp by generating input-dependent prompts through a lightweight network. MaPLe (Khattak et al. 2023) further enhances the alignment between textual and visual representations in CLIP by introducing visual prompts alongside textual ones. In recent years, prompt tuning has been introduced into federated learning due to its inherent suitability for efficient communication. For instance, PromptFL (Guo et al. 2023) and FedPrompt (Zhao et al. 2023) independently learn text prompts on each local client and aggregate the optimized prompts on the server.

## Methods

We assume  $K$  clients, where each client  $i$  possesses a dataset from domain  $\mathcal{D}_i$ . The primary objective is to jointly train a global prompt module that generalizes effectively to the unseen target domain  $\mathcal{D}_u$ . A key assumption in our framework is that each client has access to simple text descriptions of external training domains  $t_i$  (e.g., “a clipart of a”), which are considered non-sensitive and do not compromise client privacy. The framework of FaST-PT is shown in Figure 1.

### Multi-Modal Style Transfer

Existing FDG approaches improve the generalization capability of local models by incorporating knowledge from diverse domains, which necessitates sharing part of the domain knowledge among clients. While sharing such sensitive information between clients may bring the risk of data leakage, we propose a lightweight and secure local style transfer scheme, MST, which leverages CLIP’s extensive domain knowledge and strong text-image alignment capabili-

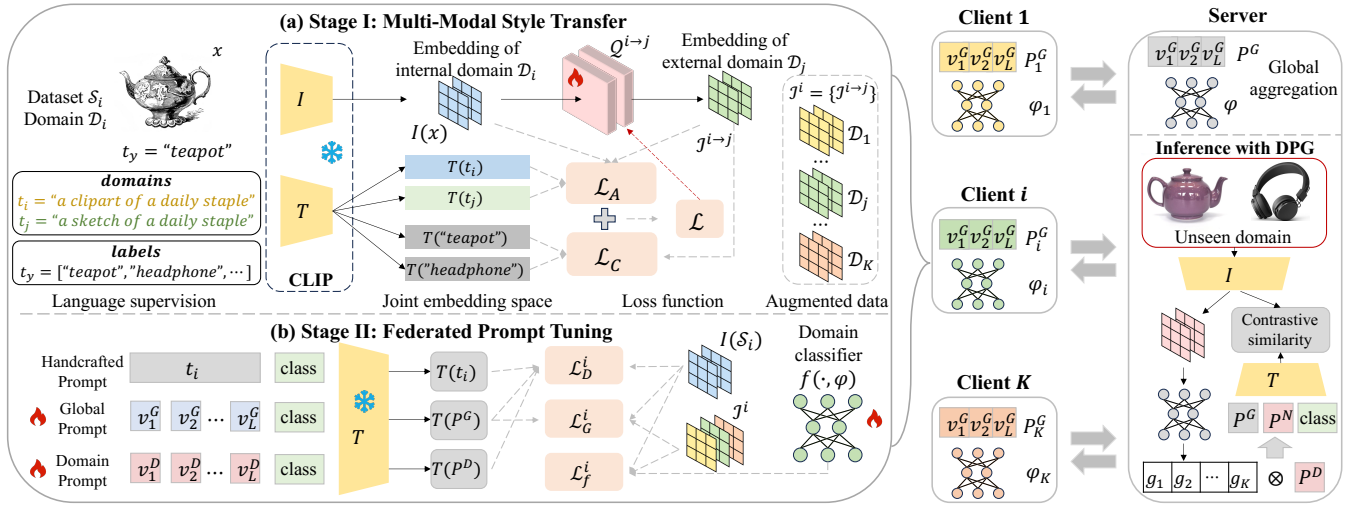


Figure 1: Framework of our proposed FaST-PT. The left side illustrates the local training process, while the right side depicts the communication between client and server. The client-side training is divided into two phases: (a) MST and (b) FPT. (a) In the MST phase, the client employs a frozen CLIP model along with text descriptions from various domains to train a transform network, enabling style transfer at the embedding level. (b) In the FPT phase, the client trains global prompts, domain prompts, and domain classifiers using the augmented data. The inference process is also depicted on the server side.

ties to facilitate latent feature augmentation at the embedding level. It allows client  $k$  with domain  $\mathcal{D}_k$  to generate image embeddings for external domains  $\mathcal{D}_{\{1,2,\dots,K\}\setminus k}$ .

MST mainly learns a transform network  $\mathcal{Q}^{i \rightarrow j}$  that transforms the image embedding of the input image from  $\mathcal{D}_i$  to  $\mathcal{D}_j$ , with the goal formulated as the following two aspects:

- The alignment between  $\mathcal{D}_i$  and  $\mathcal{D}_j$
- The preservation of class-consistent features

Let  $x$  denote the input image and  $I(x)$  denote the image embedding of  $x$ . Besides, let  $t_k$  denote the domain-related description of domain  $\mathcal{D}_k$  and  $t_y$  denote the class label. We compose these two parts to obtain the text input as  $t = [t_k, t_y]$ , which is further encoded to get the text embedding  $T(t)$ . For instance, if  $t_i = \text{“a clipart of a \{ \}”}$ ,  $t_j = \text{“a sketch of a \{ \}”}$ , and  $t_y$  could be “tiger”, then the composed text  $t = \text{“a clipart of a tiger”}$ .

To achieve the embedding transform from  $\mathcal{D}_i$  to  $\mathcal{D}_j$  while preserving class-consistent features, we design two loss functions to train  $\mathcal{Q}^{i \rightarrow j}$  for each domain  $\mathcal{D}_j$  distinct from the local domain  $\mathcal{D}_i$ : alignment loss and consistency loss.

**Alignment Loss** The alignment loss  $\mathcal{L}_A(\mathcal{Q}^{i \rightarrow j})$  aims to ensure that the embeddings of augmented images more faithfully reflect the distinctive properties of their respective domains.  $\mathcal{L}_A(\mathcal{Q}^{i \rightarrow j})$  is supervised by textual descriptions associated with different source domains. While CLIP effectively captures relationships between text and image embeddings, directly obtaining their precise correspondence remains challenging. To address this, we align the change direction in the image and text embedding spaces as the optimization target. Specifically, this change direction is defined as the normalized difference between embeddings from  $\mathcal{D}_i$  and  $\mathcal{D}_j$ . The alignment loss is thus formulated as:

$$\mathcal{L}_A(\mathcal{Q}^{i \rightarrow j}) = \sum_{(x,y) \in \mathcal{S}_i} \left( 1 - \frac{\mathcal{Q}^{i \rightarrow j}(I(x)) - I(x)}{\|\mathcal{Q}^{i \rightarrow j}(I(x)) - I(x)\|} \cdot \frac{T(t_j, t_y) - T(t_i, t_y)}{\|T(t_j, t_y) - T(t_i, t_y)\|} \right). \quad (1)$$

**Consistency Loss** The consistency loss  $\mathcal{L}_C$  is designed to preserve the class knowledge in transformed embeddings. The intuition behind designing  $\mathcal{L}_C$  is that using only the alignment loss may ensure style transfer of the embedding, but it can also erase some class information. We leverage CLIP’s zero-shot inference capabilities to assess the class information contained in the enhanced image embeddings. Specifically, we perform classification by comparing these enhanced embeddings with the text embeddings of class names. To this end, we adopt CLIP’s standard contrastive loss, formulated as follows:

$$\mathcal{L}_C(\mathcal{Q}^{i \rightarrow j}) = \mathbb{E}_{(x,y) \in \mathcal{S}_i} \ell_{ce}([\mathcal{Q}^{i \rightarrow j}(I(x)) \cdot T(t_y)], y), \quad (2)$$

where  $[I \cdot T]$  denote the Softmax-based contrastive similarity computation and Top-1 prediction process.

The general objective  $\mathcal{L}_{\text{MST}}(\mathcal{Q}^{i \rightarrow j})$  of training the transform network  $\mathcal{Q}^{i \rightarrow j}$  is the composition of the above two loss functions. To balance the effectiveness of alignment and consistency, we additionally introduce a hyperparameter  $\lambda$  to achieve the linear composition:

$$\mathcal{L}_{\text{MST}}(\mathcal{Q}^{i \rightarrow j}) = \lambda \mathcal{L}_A(\mathcal{Q}^{i \rightarrow j}) + (1 - \lambda) \mathcal{L}_C(\mathcal{Q}^{i \rightarrow j}). \quad (3)$$

Client  $i$  with domain  $\mathcal{D}_i$  locally learn  $(K - 1)$  transform networks for the other  $(K - 1)$  clients and generate the corresponding augmentation embeddings  $\mathcal{I}^{i \rightarrow j} = \mathcal{Q}^{i \rightarrow j}(x), x \in \mathcal{S}_i$ . Thus, each client maintains a set of augmentation embeddings as  $\mathcal{I}_i = \{\mathcal{I}^{i \rightarrow j}, j \in \{1, 2, \dots, K\} \setminus i\}$ .

## Federated Prompt Tuning

Yet, conventional DG methods often involve complex feature extraction techniques. How to leverage augmented embeddings to achieve domain generalization with minimal resource consumption remains an open challenge. Building on the powerful feature extraction capabilities of CLIP, we introduce prompt tuning of CLIP for efficient domain generalization. Motivated by the idea of feature disentanglement (Bui et al. 2021; Zhang et al. 2022) in traditional DG, we propose a dual-prompt module that divides the prompts into global and domain prompts.

**Global Prompt**  $P^G$  is designed to capture the general knowledge of the downstream task, which should be shared across clients.  $P^G$  is a set of  $L$  learnable vectors  $P^G = \{v_1^G, v_2^G, \dots, v_L^G\}$ , with the same dimension as text embedding  $d$ . We denote the class label as  $t_y$ . For the  $j$ -th class, we define the input of the text encoder with global prompt as  $t = [P^G, t_y^j]$ . We denote the feature map of image  $x$  as  $I(x)$  and the text feature map as  $T([P^G, t_y^j])$ . Thus, the prediction probability of the  $j$ -th class computed with contrastive similarity can be formulated as:

$$p_G(y = j|x) = \frac{\exp(\mathcal{H}(T([P^G, t_y^j]), I(x))/\tau)}{\sum_{i=1}^C \exp(\mathcal{H}(T([P^G, t_y^i]), I(x))/\tau)}, \quad (4)$$

where  $\mathcal{H}$  is the cosine similarity,  $C$  is the number of class and  $\tau$  is the temperature of Softmax function. We optimize the global prompts by minimizing the cross-entropy loss computed over both the augmented embeddings and the original dataset. Specifically, the training loss on client  $i$  is defined as follows:

$$\mathcal{L}_G^i = \mathbb{E}_{(x,y) \in \mathcal{S}_i \cup \mathcal{I}_i} [\ell_{ce}(y, p_G(y|x))]. \quad (5)$$

**Domain Prompt**  $P_i^D$  is designed to capture the domain-specific knowledge of client  $i$ . Similarly,  $P_i^D$  is a set of vectors  $P_i^D = \{v_1^D, v_2^D, \dots, v_L^D\}$  with length  $L$  and dimension  $d$ . We concatenate  $P^G$  with the domain prompt to form a new composite prompt  $P^C = [P^G, P_i^D]$ . Specifically, the predicted probability is updated accordingly after introducing the domain prompt, by replacing  $P^G$  with  $P^C$ . The core idea of the domain-specific prompt is to enhance the performance of  $P_i^D$  on local data  $\mathcal{S}_i$ . We similarly optimize it using the cross-entropy loss  $\mathcal{L}_{\text{Cla}}^i = \mathbb{E}_{(x,y) \in \mathcal{S}_i} [\ell_{ce}(y, p_C(y|x))]$ . Besides, the training of  $P_i^D$  should also enforce orthogonality with  $P^G$  to guarantee the uniqueness of the extracted knowledge. To this end, we additionally introduce a contrastive loss term as follows:

$$\begin{aligned} \mathcal{L}_D^i &= \mathcal{L}_{\text{Cla}}^i + \mathcal{L}_{\text{Con}}^i = \mathbb{E}_{(x,y) \in \mathcal{S}_i} [\ell_{ce}(y, p_C(y|x))] \\ & - \log \frac{\exp(\text{sim}(P_i^D, t_i))}{\exp(\text{sim}(P_i^D, t_i)) + \exp(\text{sim}(P_i^D, P^G))}, \end{aligned} \quad (6)$$

where  $t_i$  is the handcrafted prompt for Domain  $\mathcal{D}_i$  and  $p_C(y|x)$  is the probability predicted with prompt  $P^C$ .

Notably, clients share global prompts for aggregation during training, while domain-specific prompts are retained locally and only shared after training. The final domain prompt list is formed as  $P^D = [P_1^D, P_2^D, \dots, P_K^D]$ .

---

## Algorithm 1: FaST-PT

---

- 1: **Input:** Datasets  $\mathcal{S}_i$ , client number  $K$ , training round  $R$ , epoch number  $E_G$  and  $E_L$
  - 2: **Output:** Global Prompt  $P^G$ , Domain prompt list  $P^D$ , domain classifier  $f(\cdot, \varphi)$
  - 3: **Stage 1: Multi-Modal Style Transfer**
  - 4: **for** client  $i = 1, 2, \dots, K$  **do**
  - 5:   **for** Domain  $\mathcal{D}_j \in \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\} \setminus \mathcal{D}_i$  **do**
  - 6:     Train  $\mathcal{Q}^{i \rightarrow j}$  with  $\mathcal{L}_{\text{MST}}(\mathcal{Q}^{i \rightarrow j})$  (Eq. 3)
  - 7:   **end for**
  - 8: **end for**
  - 9: Clients locally augments image embedding  $\mathcal{I}^{i \rightarrow j}$
  - 10: **Stage 2: Federated Prompt Tuning**
  - 11: **for** round  $r = 1, 2, \dots, R$  **do**
  - 12:   **for** client  $i = 1, 2, \dots, K$  **do**
  - 13:     **for** epoch  $e = 1, 2, \dots, E_G$  **do**
  - 14:       Training  $P_i^G$  and  $\varphi_i$  with  $\mathcal{L}_G^i$  and  $\mathcal{L}_F^i$
  - 15:     **end for**
  - 16:   **end for**
  - 17:   Upload  $P_i^G$  and  $\varphi_i$
  - 18:   Server aggregates  $P^G$  and  $\varphi$  and broadcasts
  - 19:   **for** client  $i = 1, 2, \dots, K$  **do**
  - 20:     **for** epoch  $e = 1, 2, \dots, E_D$  **do**
  - 21:       Training  $P_i^D$  with  $\mathcal{L}_D^i$
  - 22:     **end for**
  - 23:   **end for**
  - 24: **end for**
  - 25: Server collects  $P^D$  and broadcasts
- 

**Domain-aware Prompt Generation** Considering the latent similarity between training and unseen domains, we further design a domain-aware prompt generation scheme to integrate knowledge from diverse domains to enhance generalization performance, which enables better adaptation of the target domain  $\mathcal{D}_u$ . We utilize a domain classifier  $f(\cdot, \varphi)$  to facilitate the domain similarity analysis.  $f(\cdot, \varphi)$  takes image feature as the input and outputs the one-hot format result with  $K$  dimension as:  $\mathbf{G} = [g_1, g_2 \dots g_K]$ .

We weight  $\otimes$  the domain prompt list  $P^D$  with the pseudo-domain label  $\mathbf{G}$  element-wise to generate the dynamic prompt  $P^N$  for arbitrary image  $x$  as follows:

$$P^N = f(I(x), \varphi) \otimes P^D. \quad (7)$$

$P^N$  integrates knowledge from multiple source domains to generate knowledge that approximates the target domain. During the inference stage,  $P^N$  is concatenated with  $P^G$  and fed into the text encoder. For the optimization of  $f(\cdot, \varphi)$ , we leverage the augmented dataset  $\mathcal{S}_i \cup \mathcal{I}_i$  for training and denote the domain index as the label. We utilize the cross-entropy loss and formulate the objective as follows:

$$\mathcal{L}_F^i = \mathbb{E}_{(x,y,j) \in \mathcal{S}_i \cup \mathcal{I}_i} [\ell_{ce}(j, f(x, \varphi_i))], \quad (8)$$

where  $j$  is the domain index.

During the training of the prompt module, we adopt a two-stage alternating update strategy. First, clients update the global parameters  $P^G$  and  $\varphi$ , and upload the gradient

Methods	PACS (%)					VLCS (%)				
	A	C	P	S	Avg	C	L	S	V	Avg
FedAvg	81.27	76.78	92.33	63.54	78.48	92.78	62.39	73.02	74.98	75.79
FedProx	80.22	77.17	96.08	64.28	79.44	93.22	62.80	74.64	74.96	76.41
ELCFS	85.89	74.02	92.14	69.16	80.30	93.28	61.64	73.42	76.06	76.10
FedSR	84.97	85.13	91.37	67.13	82.15	95.55	63.29	71.93	75.43	76.55
FedADG	80.16	81.92	91.28	68.43	80.45	95.96	65.63	70.29	76.70	77.15
FedDG-GA	94.03	83.19	76.85	82.93	84.25	95.14	62.82	74.41	77.49	77.47
PromptFL	93.19	94.72	99.10	83.76	92.70	99.26	66.08	79.38	78.56	80.82
FedAPT	93.78	95.52	99.63	85.31	93.56	99.01	67.91	78.66	81.53	81.78
DiPrompt	95.72	96.39	99.24	88.07	94.86	<b>99.70</b>	68.97	84.72	82.93	84.08
<i>FaST-PT</i>	<b>97.32</b>	<b>99.31</b>	<b>99.93</b>	<b>96.09</b>	<b>98.16</b>	99.17	<b>70.59</b>	<b>86.05</b>	<b>86.26</b>	<b>85.52</b>

Table 1: Performance comparison of different methods on the PACS dataset and the VLCS dataset.

updates, which are then aggregated by the server and broadcast back to all clients. Second, clients update the local parameters  $P_i^D$  and upload them after the final round. Since the training of  $P^G$  and  $f(\cdot, \varphi)$  does not interfere with each other, it is reasonable to train them simultaneously. The complete training process is illustrated in Alg.1.

## Discussion

FaST-PT achieves efficient local feature enhancement and robust domain generalization. Specifically, although FaST-PT leverages MST to mitigate domain shift and adopts prompt tuning for domain generalization, both components are built upon the CLIP model. This unified backbone design obviates the need for deploying additional models on local clients, thereby reducing memory consumption and computational overhead. Moreover, unlike traditional federated learning methods that require sharing full model parameters, FaST-PT only exchanges prompts and their associated lightweight generation modules across clients, substantially lowering communication costs (as further evidenced in our experimental results). Finally, in contrast to approaches that transmit domain-specific features, MST merely requires clients to share text descriptions of their respective domains—information that is inherently less sensitive, thus providing stronger privacy protection.

## Experiments

### Experimental Setup

**Datasets** We perform extensive experiments on four cross-domain datasets: **PACS** (Li et al. 2017) has four domains, *photo* (P), *art-painting* (A), *cartoon* (C) and *sketch* (S). **VLCS** (Fang, Xu, and Rockmore 2013) is a widely used benchmark for domain generalization, consisting of four domains, *VOC2007* (V), *LabelMe* (L), *Caltech101* (C), and *SUN09* (S). **DomainNet** (Peng et al. 2019) is a large-scale dataset across six domains, *clipart* (C), *infograph* (I), *painting* (P), *quickdraw* (Q), *real* (R), and *sketch* (S). **Office-Home** (Venkateswara et al. 2017) is a medium-scale benchmark dataset containing four domains: *Art* (A), *Clipart* (C), *Product* (P), and *Real-World* (R). Following (Li et al. 2021; Yang, Wang, and Wang 2023), we select the top 10 classes for our experiment on DomainNet and OfficeHome. We

adopt the leave-one-domain-out protocol, treating one domain as the target while using the others as source domains.

**Baselines** We compare our method with three types of state-of-the-art methods: 1) Traditional federated learning framework including **FedAvg** (McMahan et al. 2017) and **FedProx** (Li et al. 2020) that enables clients to collaboratively train a global model from scratch. 2) Conventional FedDG methods that integrate DG with FL to collaboratively train a generalized model that perform well on unseen domains, including **ELCFS** (Liu et al. 2021), **FedSR** (Nguyen, Torr, and Lim 2022), **FedADG** (Zhang et al. 2021b) and **FedDG-GA** (Zhang et al. 2021a). 3) Prompt tuning-based methods that introduce prompt tuning into FL to enhance the generalization ability of CLIP, including **PromptFL** (Guo et al. 2023), **FedAPT** (Su et al. 2024), and **DiPrompt** (Bai et al. 2024).

**Implement Details** We conduct our experiment with Python 3.11 on NVIDIA RTX A40 GPU. We use the OpenAI CLIP model with a ViT-L/14 backbone, which encodes both images and text into 768-dimensional feature vectors. For the MST stage, we utilize a 2-layer MLP as the transform network  $Q^{i \rightarrow j}$  with a hidden dimension of 384. we utilize the Adam optimizer with a learning rate of  $1e^{-3}$  and a weight decay of 0.05. The loss weight  $\lambda$  is set as 0.5. For the prompt training stage, we utilize a fully-connection layer as the domain classifier  $f(\cdot, \varphi)$ . We use the SGD optimizer with a learning rate of 0.005 and 0.01 for training of the prompt and domain classifier, respectively. We set the communication round  $R$  as 20 and local training epoch  $E_G = E_L = 5$ . We set the batch size as 16 and prompt length as 4. We use the same hyperparameters and backbones across all experiments to ensure a fair comparison.

### Evaluation

We demonstrate the experimental results of FaST-PT and other methods in Table 1 and Table 2, which consists of totally 18 domain generalization tasks on PACS, VLCS, OfficeHome and DomainNet. All experimental results are reported as the average over three runs with different random seeds. Generally, we observe that although conventional FDG methods have demonstrated significantly better generalization performance than traditional FL approaches,

Methods	OfficeHome (%)					DomainNet (%)						
	A	C	P	R	Avg	C	I	P	Q	R	S	Avg
FedAvg	82.56	79.20	88.28	91.15	85.30	78.04	62.69	80.13	45.78	82.49	84.17	72.22
FedProx	83.19	77.91	90.28	91.60	85.74	79.11	64.77	82.82	51.65	84.65	83.94	74.49
ELCFS	88.72	82.64	89.37	91.75	88.12	83.28	67.18	85.78	58.65	88.31	86.37	78.26
FedSR	88.17	84.47	88.89	90.65	88.04	85.17	66.57	86.01	60.95	87.26	90.10	79.34
FedADG	86.51	87.25	91.52	92.30	89.40	83.95	69.68	84.05	61.51	90.78	86.79	79.46
FedDG-GA	87.86	86.90	92.42	93.25	90.11	86.69	70.93	85.69	60.72	89.82	87.18	80.17
PromptFL	92.35	88.73	95.20	98.80	93.77	89.92	72.75	87.46	60.85	93.54	90.47	82.50
FedAPT	94.28	90.76	97.08	98.95	95.27	91.44	72.04	90.62	63.88	95.27	94.23	84.58
DiPrompt	95.48	91.26	97.90	99.10	95.94	93.89	75.34	93.72	64.16	96.31	97.62	86.84
<i>FaST-PT</i>	<b>98.57</b>	<b>92.72</b>	<b>99.24</b>	<b>99.40</b>	<b>97.48</b>	<b>97.48</b>	<b>78.14</b>	<b>97.17</b>	<b>66.88</b>	<b>97.90</b>	<b>98.46</b>	<b>89.17</b>

Table 2: Performance comparison of different methods on the OfficeHome dataset and the DomainNet dataset.

there remains a notable performance gap compared to recent prompt-based FL methods. This substantial improvement in generalization can be largely attributed to the strong cross-modal and generalization capabilities of CLIP. Even when all methods are built upon the same pretrained visual backbone (ViT-L/14), prompt tuning enables more efficient adaptation to the distributional characteristics of each domain. Specifically, FaST-PT achieves state-of-the-art performance with an average accuracy of 98.16% on PACS, 85.52% on VLCS, 97.48% on OfficeHome, and 89.17% on DomainNet, surpassing the SOTA method DiPrompt by up to +3.30% on PACS and +2.33% on DomainNet. Among the total of 18 generalization tasks across all benchmarks, FaST-PT achieves the best performance on 17 tasks, demonstrating remarkable consistency. Compared to existing prompt-based FL methods, FaST-PT still demonstrates a substantial performance advantage, highlighting the effectiveness of the proposed MST and DAP strategy, both of which play a critical role in enhancing cross-domain adaptability and mitigating the adverse effects of distributional shifts in FL.

## Ablation Study

**Effectiveness of Components** To investigate the effectiveness of different components in FaST-PT, we evaluated various variants of FaST-PT on the PACS dataset, as shown in Table 3. In variants ① and ②, we only employ global prompts or domain prompts combined with the DPG mechanism. Compared to variant ④, the performance decreased noticeably, which demonstrates the effectiveness of integrating global and domain knowledge for generalization tasks. In addition, through two comparative analyzes (① vs. ③ and ④ vs. FaST-PT), we observed that the introduction of the MST strategy can mitigate the domain shift problem in FDG and effectively improve model performance, thus improving the generalization capability of the global model. Variant ⑤ exhibits a performance decline compared to FaST-PT, which also reflects that the contrastive loss we employed can strengthen the separation between global and domain knowledge, thereby indirectly enhancing adaptation to the target domain. In Variant ⑥, we conducted experiments by supplementing the textual descriptions of the target domain, allowing clients to generate image features for the target domain to assist the training of the global model.

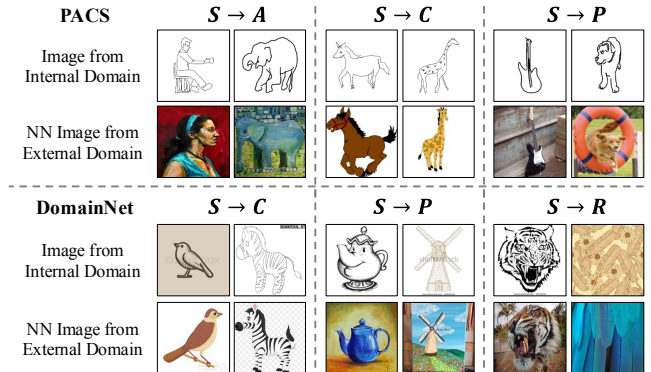


Figure 2: Nearest Neighbors on PACS and DomainNet.

Its performance only surpasses FaST-PT by 0.24%.

**Visualization of MST** We evaluate the style transfer effect of the MST strategy, with the results presented in Figure 2. Since existing methods do not support converting image features back into images, we adopt a Nearest Neighbors (NN) approach to assess the generated augmented features. Specifically, we use all data from the extended domain as the reference set. For each enhanced embedding, we compute its cosine similarity with embeddings in the reference set and select the one with the highest similarity as its NN. In Figure 2, the top row displays the original images used for style transfer, while the bottom row shows the NNs corresponding to their enhanced features. We observe that although there is a clear difference in style between the top and bottom images, they consistently belong to the same semantic category. This observation aligns with the objectives of the two-part loss function employed during training and demonstrates the effectiveness of the MST strategy.

**Effect of prompt length** We analyze the effect of prompt length and demonstrate the result in Figure 3. Here, we set the lengths of the global and domain prompts to be equal. On PACS and DomainNet datasets, we observed that in most tasks, model performance exhibited a trend of initially increasing and then decreasing as the prompt length grows. Short prompt constrains the performance due to the limited information conveyed in the prompt. Conversely, long

Variant	$P^G$	$P^D$	$\mathcal{L}_{\text{Con}}$	DPG	MST	$t_{\text{target}}$	A	C	P	S	Average
①	✓						91.21	90.94	94.54	93.61	92.58
②		✓		✓			85.48	83.16	90.79	91.25	87.67
③	✓				✓		92.43	94.14	95.81	95.13	94.38
④	✓	✓	✓	✓			96.03	97.44	98.98	95.77	97.06
⑤	✓	✓		✓	✓		97.74	99.09	99.48	95.48	97.95
⑥	✓	✓	✓	✓	✓	✓	97.93	99.68	99.85	96.27	98.42
FaST-PT	✓	✓	✓	✓	✓		97.32	99.31	99.93	96.09	98.16

Table 3: Effect of different components of Fast-PT on PACS.

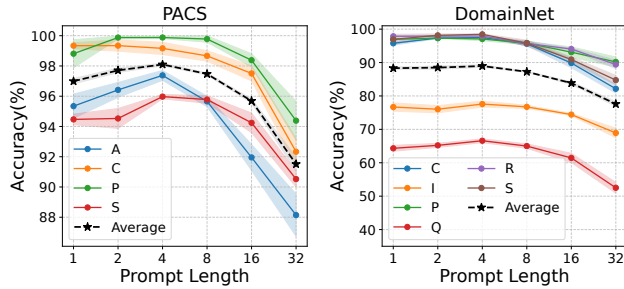


Figure 3: Effect of prompt length on PACS and DomainNet datasets over three runs with different random seeds.

prompts may introduce noise or task-irrelevant information, which can negatively impact the alignment between image and text. Moreover, when using long prompts, the prompt weighting strategy employed in FaST-PT may lead to issues such as semantic dilution and reduced interpretability. We observe that when the prompt length is set to 4, most tasks, as well as the average accuracy, achieve their optimal performance. Therefore, we adopt a prompt length of 4 as the default setting in our experiments.

**Result under Few-shot scenario** To investigate the effectiveness of our framework in a few-shot setting, we compare the generalization ability of FaST-PT with existing prompt-based methods under varying conditions of extremely limited data (shots = 1, 2, 4, 8, 16, 32). As illustrated in Figure 4, we observe that the performance of all methods improves with the increase in training samples. FaST-PT consistently demonstrates significant advantages across all settings on both the PACS and DomainNet datasets. These experimental results indicate that FaST-PT is capable of capturing both global and domain-specific information under few-shot settings, achieving robust generalization to unseen domains.

**Cost Analysis** We analysis the efficiency of FaST-PT in terms of communication and computation cost. We selected one representative method from each of the three categories of baseline approaches for analysis. For fairness, all methods employ ViT-L/14 as the visual encoder, with a consistent batch size of 16 and prompt length set to 4. Regarding communication efficiency, we measure params each client uploads per training round. Compared to traditional FL that requires uploading the entire model, prompt-based FL methods incur significantly lower communication costs.




Figure 4: Few-shot performance on PACS and DomainNet datasets over three runs with different random seeds.

Method	FedAvg	FedADG	DiPrompt	Ours
Parameters	305M	318M	0.012M	0.005M
FLOPs	1.296T	2.593T	1.739T	1.431T

Table 4: Communication and computation cost compared with baseline FDG methods. ( $Batchsize = 16$ )

Moreover, relative to DiPrompt, our method utilizes more lightweight and efficient modules, resulting in fewer than half the number of transmitted parameters. In terms of computational cost, we evaluated the FLOPs for processing a batch during inference. The additional modules introduced by FedADG substantially increase inference cost, whereas our approach only adds a domain classifier. Consequently, its FLOPs are only slightly higher than FedAvg and remain lower than those of DiPrompt.

## Conclusion

In this work, we propose a new framework named FaST-PT, which integrates CLIP into FL to tackle the FDG problem in terms of local feature augmentation and efficient unseen domain adaptation. Specifically, FaST-PT utilizes the MST strategy that allows clients to generate image embedding of external domains for mitigating the domain shift among clients. Besides, we design a prompt tuning module that partitions the prompt for capturing general and domain-specific knowledge, along with a domain-aware prompt generation mechanism to achieve the similarity-based domain knowledge mixing for efficient adaptation to unseen domains. Extensive experiments demonstrate that FaST-PT outperforms existing FDG methods across all the generalization tasks on four benchmark cross-domain datasets.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62572311, 62202302, U21B2019, 62406347, and 62572260.

## References

- Bai, S.; Zhang, J.; Guo, S.; Li, S.; Guo, J.; Hou, J.; Han, T.; and Lu, X. 2024. Diprompt: Disentangled prompt tuning for multiple latent domain generalization in federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27284–27293.
- Bui, M.-H.; Tran, T.; Tran, A.; and Phung, D. 2021. Exploiting domain-specific features to enhance domain generalization. *Advances in Neural Information Processing Systems*, 34: 21189–21201.
- Chen, J.; Jiang, M.; Dou, Q.; and Chen, Q. 2023. Federated domain generalization for image recognition via cross-client style transfer. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 361–370.
- Fang, C.; Xu, Y.; and Rockmore, D. N. 2013. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE international conference on computer vision*, 1657–1664.
- Guan, H.; Yap, P.-T.; Bozoki, A.; and Liu, M. 2024. Federated learning for medical image analysis: A survey. *Pattern Recognition*, 110424.
- Guo, T.; Guo, S.; Wang, J.; Tang, X.; and Xu, W. 2023. Promptfl: Let federated participants cooperatively learn prompts instead of models-federated learning in age of foundation model. *IEEE Transactions on Mobile Computing*.
- Huang, W.; Ye, M.; Shi, Z.; Li, H.; and Du, B. 2023. Rethinking federated learning with domain shift: A prototype view. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16312–16322. IEEE.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19113–19122.
- Konečný, J.; McMahan, H. B.; Yu, F. X.; Richtárik, P.; Suresh, A. T.; and Bacon, D. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6027–6036.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, 5542–5550.
- 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, X.; Jiang, M.; Zhang, X.; Kamp, M.; and Dou, Q. 2021. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*.
- Li, Y.; Gong, M.; Tian, X.; Liu, T.; and Tao, D. 2018. Domain generalization via conditional invariant representations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Liu, Q.; Chen, C.; Qin, J.; Dou, Q.; and Heng, P.-A. 2021. Feddgn: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1013–1023.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Muandet, K.; Balduzzi, D.; and Schölkopf, B. 2013. Domain generalization via invariant feature representation. In *International conference on machine learning*, 10–18. PMLR.
- Nguyen, A. T.; Torr, P.; and Lim, S. N. 2022. Fedsr: A simple and effective domain generalization method for federated learning. *Advances in Neural Information Processing Systems*, 35: 38831–38843.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1406–1415.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Su, S.; Yang, M.; Li, B.; and Xue, X. 2024. Federated adaptive prompt tuning for multi-domain collaborative learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15117–15125.
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5018–5027.
- Wu, G.; and Gong, S. 2021. Collaborative optimization and aggregation for decentralized domain generalization and adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6484–6493.
- Yang, F.-E.; Wang, C.-Y.; and Wang, Y.-C. F. 2023. Efficient model personalization in federated learning via client-specific prompt generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19159–19168.
- Zhang, H.; Zhang, Y.-F.; Liu, W.; Weller, A.; Schölkopf, B.; and Xing, E. P. 2022. Towards principled disentanglement for domain generalization. In *Proceedings of the IEEE/CVF*

*conference on computer vision and pattern recognition*, 8024–8034.

Zhang, L.; Lei, X.; Shi, Y.; Huang, H.; and Chen, C. 2021a. Federated learning with domain generalization. *arXiv preprint arXiv:2111.10487*.

Zhang, L.; Lei, X.; Shi, Y.; Huang, H.; and Chen, C. 2021b. Federated learning with domain generalization. *arXiv preprint arXiv:2111.10487*.

Zhao, H.; Du, W.; Li, F.; Li, P.; and Liu, G. 2023. Fed-Prompt: Communication-Efficient and Privacy-Preserving Prompt Tuning in Federated Learning. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16816–16825.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.