

Multi-Source Collaborative Style Augmentation and Domain-Invariant Learning for Federated Domain Generalization

Yikang Wei

College of Intelligence and Computing, Tianjin University, Tianjin, China
yikang@tju.edu.cn

Abstract

Federated domain generalization aims to learn a generalizable model from multiple decentralized source domains for deploying on the unseen target domain. The style augmentation methods have achieved great progress on domain generalization. However, the existing style augmentation methods either explore the data styles within isolated source domain or interpolate the style information across existing source domains under the data decentralization scenario, which leads to limited style space. To address this issue, we propose a Multi-source Collaborative Style Augmentation and Domain-invariant learning method (MCSAD) for federated domain generalization. Specifically, we propose a multi-source collaborative style augmentation module to generate data in the broader style space. Furthermore, we conduct domain-invariant learning between the original data and augmented data by cross-domain feature alignment within the same class and classes relation ensemble distillation between different classes to learn a domain-invariant model. By alternatively conducting collaborative style augmentation and domain-invariant learning, the model can generalize well on unseen target domain. Extensive experiments on multiple domain generalization datasets indicate that our method significantly outperforms the state-of-the-art federated domain generalization methods.

source domains to against the domain-specific bias, e.g. some Single-Domain Generalization (Single-DG) methods[Kang *et al.*, 2022; Zhou *et al.*, 2020b; Zhou *et al.*, 2020a] expand the style space of single source domain and the Multi-Domain Generalization (Multi-DG) methods[Xu *et al.*, 2021; Zhou *et al.*, 2021b] conduct style interpolation between multiple source domains to generate data with novel styles.

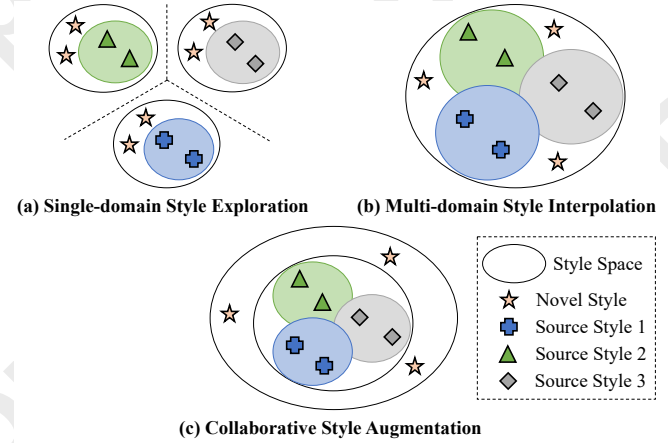


Figure 1: (a) Single-domain style exploration on the isolated source domains ignores the styles of other domains, which leads to limited style diversity. (b) Multi-domain style interpolation mixes the shared style information across decentralized source domains, which leads to the generated styles within the existing source domains. (c) Collaborative style augmentation proposed in this work generates data with novel styles out of the existing source domains, which can explore the broader style space.

1 Introduction

Domain generalization is a challenging task in machine learning, which involves training a model on source domains to generalize well on the unseen target domain. As the data from different domains exist domain shift[Huang *et al.*, 2024; Wu *et al.*, 2024; Li *et al.*, 2024; Tang *et al.*, 2024b; Tang *et al.*, 2025], the model trained on source domains tends to be domain-specific, leading to the degraded performance on unseen target domain. The domain-specific bias between source domains and the unseen target domain partially stems from the differences of data styles. Based on this assumption, the existing domain generalization methods diversify the style of

Conventional Multi-DG[Dou *et al.*, 2019; Du *et al.*, 2022; Zhou *et al.*, 2021b; Tang *et al.*, 2024a] assumes that the data from different source domains are centralized and available for learning a domain-invariant model. However, considering the data privacy and communication cost, the data from different source domains are decentralized and forbidden to share across domains[Yuan *et al.*, 2023; Wu and Gong, 2021; Wei *et al.*, 2023; Wei and Han, 2023; Wei and Han, 2022]. In this work, we consider the Federated Domain Generalization (FedDG) scenario[Yuan *et al.*, 2023; Wei and Han, 2024; Park *et al.*, 2023; Xu *et al.*, 2023b], where the data from different source domains are decentralized.

Under the data decentralization scenario, either single-domain style exploration within each isolated source domain or multi-domain style interpolation across decentralized source domains, could be used to diversify the source domains. However, the diversity of styles generated by existing methods are still limited. Firstly, the single-domain style exploration methods [Zhong *et al.*, 2022; Wang *et al.*, 2021] employed by Single-DG ignore the styles of other source domains, which degrades the generalization performance, as shown in Figure 1 (a). Secondly, the multi-domain style interpolation methods [Liu *et al.*, 2021; Chen *et al.*, 2023] used by Multi-DG result in limited styles within the existing source domains, as shown in Figure 1 (b). Furthermore, the style information across decentralized source domains should be shared for conducting the multi-domain style interpolation under the data decentralization scenario, which have the risk of privacy leakage and lead to additional communication cost. How to explore the out-of-distribution styles between multiple decentralized source domains becoming the key challenge.

For exploring the broader style space across multiple decentralized source domains, we propose a FedDG method to collaboratively explore the out-of-distribution styles across decentralized source domains. Inspired by the existing methods [Huang and Belongie, 2017; Zhong *et al.*, 2022], we conduct channel-wise statistic features transformation to explore the out-of-distribution styles with the collaboration of other source domains. Specifically, the classifier heads from other source domains are used as the discriminators. The augmented data which cannot be correctly classified by the existing classifiers from other domains, tend to be out-of-distribution. Different from the existing single-domain exploration methods and multi-domain style interpolation methods, the proposed Collaborative Style Augmentation (CSA) method can collaboratively explore the broader style space between the decentralized source domains, as shown in Figure 1(c). And the risk of privacy leakage and the cost of communication are smaller than sharing style information across decentralized domains [Liu *et al.*, 2021; Chen *et al.*, 2023].

Furthermore, we conduct domain-invariant learning between the original data and augmented data to learn the domain-invariant information for generalizing well on the unseen target domain. Specifically, we align the features of original data and augmented data by contrastive loss to increase the compactness of representations within the same class. And we distill the classes relationships from multiple classifier heads to improve the generalizable ability of model. On the decentralized source domains, the collaborative style augmentation and domain-invariant learning are conducted alternatively to improve the generalizable ability of models on the unseen target domain.

The experiments on multiple domain generalization datasets indicate the effectiveness of our method. The major contributions of this work can be summarized as follows:

- We propose a Collaborative Style Augmentation module to explore the out-of-distribution styles under the data decentralization scenario with the collaboration of other source domains.

- We conduct domain-invariant learning between the original data and augmented data by contrastive alignment and ensemble distillation for learning domain-invariant model which can generalize well on the unseen target domain.
- The results and analysis on multiple domain generalization datasets indicate that our method outperforms the state-of-the-art FedDG methods significantly.

2 Related Works

Multi-source Domain Generalization. Multi-DG aims to learn a generalizable model by utilizing multiple labeled source domains. Conventional Multi-DG methods can be categorized into (1) data augmentation methods, (2) domain-invariant representation learning methods, and (3) other learning strategies. Data augmentation methods aim to diversify the source domains for improving the generalization ability of model on unseen target domain. For example, L2A-OT [Zhou *et al.*, 2020b] and DDAIG [Zhou *et al.*, 2020a] learn the image generator to generate images with novel styles, FACT [Xu *et al.*, 2021], MixStyle [Zhou *et al.*, 2021b], and EFDMix [Zhang *et al.*, 2022] conduct style interpolation between different source domains to generate data with novel styles. Domain-invariant representation learning methods aim to learn the intrinsic semantic representation from multiple source domains for applying to unseen target domain. Such as the self-supervised learning methods, JiGen [Carlucci *et al.*, 2019] and EISNet [Wang *et al.*, 2020] utilize the Jigsaw auxiliary task to learn the domain-invariant representations. And RSC [Huang *et al.*, 2020], CDG [Du *et al.*, 2022], I²-ADR [Meng *et al.*, 2022], and DomainDrop [Guo *et al.*, 2023] learn the intrinsic semantic representations by removing the spurious correlation features. Learning strategies, e.g. MASF [Dou *et al.*, 2019] utilizes meta-learning to learn intrinsic semantic representations across different domains. And DAEL [Zhou *et al.*, 2021a] utilizes ensemble learning to learn the complementary knowledge from different domains. However, these conventional Multi-DG methods assume that the data from multiple source domains can be accessed simultaneously, which cannot be satisfied under the data decentralization scenario.

Single-source Domain Generalization. Single-DG directly learns a generalizable model from single source domain. The existing Single-DG methods [Wang *et al.*, 2021; Li *et al.*, 2021a; Xu *et al.*, 2023a; Zhou *et al.*, 2021b; Wang *et al.*, 2021] usually synthesize images or features with novel styles and keep the semantic information invariant, which can expand the diversity of source domain to achieve better generalization ability. For example, StyleNeophile [Kang *et al.*, 2022] and DSU [Liu *et al.*, 2021] explore novel styles in feature space to expand the source domain. And the AdvStyle [Zhong *et al.*, 2022] generates images with novel styles by adversarial augmentation. However, these Single-DG methods cannot collaboratively utilize multiple decentralized source domains and lead to limited generalization performance. We will give a detailed comparison and analysis in experiments by applying these single-domain style exploration methods on FedDG scenario.

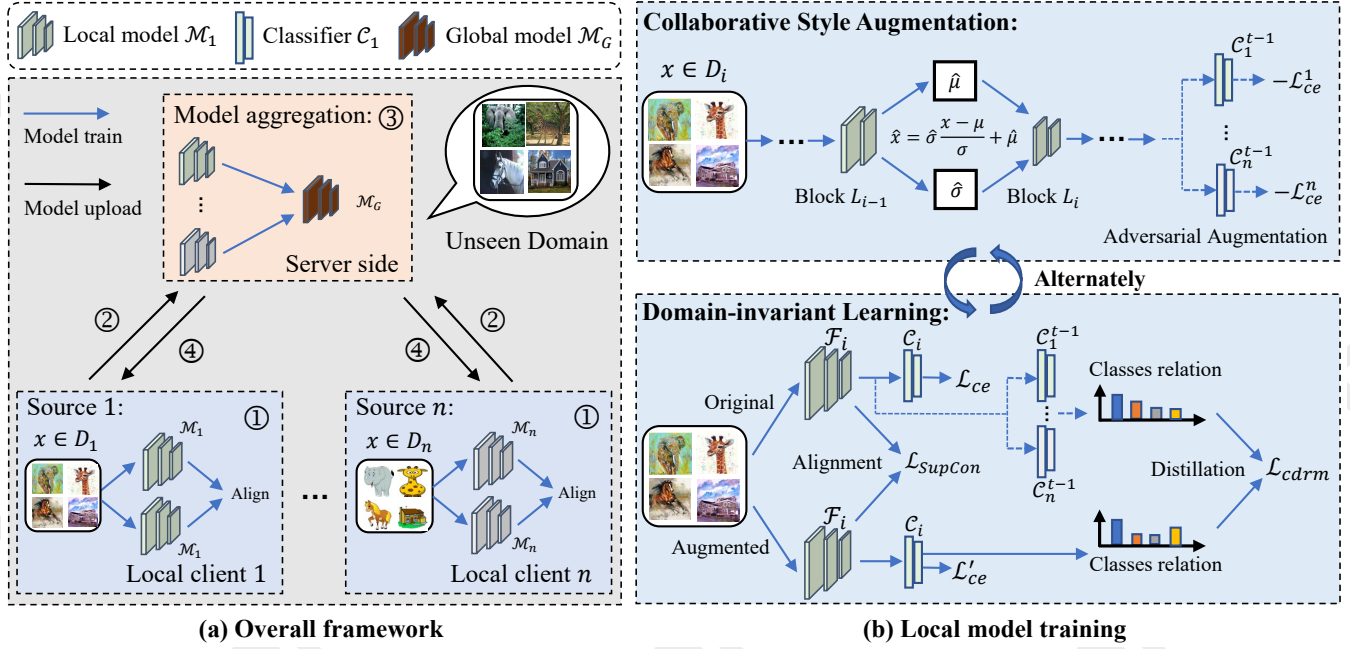


Figure 2: (a) The overall framework of our method. The local source domain models $\{\mathcal{M}_i\}_{i=1}^n$ are trained on the isolated local clients and aggregated on the server side to obtain the global model \mathcal{M}_G , which will be used on the unseen domain for better generalizable ability. (b) The local model training on decentralized source domains, where collaborative style augmentation and domain-invariant learning are conducted alternately to improve the generalizable ability of local source domain models.

Federated Domain Generalization. FedDG collaboratively train multiple decentralized source domains for obtaining a model generalizing well on the unseen target domain. The existing FedDG methods utilize the federated learning framework e.g. FedAvg [McMahan *et al.*, 2017] to collaboratively train the local source models on the decentralized source domains and aggregate the local source models on the server side. The existing FedDG methods can be categorized into (1) domain-invariant learning methods and (2) model aggregation strategies. The domain-invariant learning methods focus on exploring the data with out-of-distribution styles for generalizing on the unseen target domain. For example, ELCFS [Liu *et al.*, 2021] and CCST [Chen *et al.*, 2023] interpolate the shared style information across decentralized source domains to generate data with novel styles, the StableFDG [Park *et al.*, 2023] explores novel styles based on the shared style information across decentralized source domains. And FADH [Xu *et al.*, 2023b] trains additional image generators on the local clients to generate images with novel styles. COPA [Wu and Gong, 2021] augments the images by pre-defined augmentation pool, e.g. RandAug [Cubuk *et al.*, 2020] for learning the domain-invariant model. Different from these multi-domain style interpolation methods, our method can explore the broader style space without sharing the style information across decentralized source domains. The model aggregation strategies e.g. CASC [Yuan *et al.*, 2023] and GA [Zhang *et al.*, 2023] calibrate the aggregation weights of local models on server side to obtain a fairness global model for the better generalization ability. Different from these methods, our method aims to learn domain-

invariant model on local clients.

3 Method

Problem Definition. We focus on the generic image classification task for FedDG. Given multiple decentralized source domains $\{X_i, Y_i\}_{i=1}^n$, each domain $\{X_i, Y_i\}$ contains N_i samples. Data from different source domains are located on the isolated clients and forbidden to share across clients. The goal of FedDG is to learn a generalizable global model \mathcal{M}_G from multiple decentralized source domains $\{X_i, Y_i\}_{i=1}^n$ for deploying on the unseen target domain X_t . The different source domains $\{X_i, Y_i\}_{i=1}^n$ and unseen target domain X_t exist the covariate shift, e.g. the marginal distribution of images $P(X)$ differs but the conditional label distribution $P(Y|X)$ keeps same across domains.

3.1 Overall Framework

We propose a multi-source collaborative style augmentation and domain-invariant learning method to explore out-of-distribution styles and learn the domain-invariant model across decentralized source domains. As shown in Figure 2(a), each client contains a source domain, there are four steps to collaboratively train the decentralized source domains. On Step 1, the local models $\{\mathcal{M}_i\}_{i=1}^n$ are trained on the isolated source domains respectively. Each local model \mathcal{M}_i contains a feature extractor \mathcal{F}_i and a classifier head \mathcal{C}_i . On Step 2, the local models $\{\mathcal{M}_i\}_{i=1}^n$ are uploaded to the server side. On Step 3, the local models $\{\mathcal{M}_i\}_{i=1}^n$ are aggregated by parameter averaging to obtain a global model \mathcal{M}_G . Then, the global model \mathcal{M}_G and the classifier heads $\{\mathcal{C}_i^{t-1}\}_{i=1}^n$ of dif-

ferent source domains are broadcasted to local clients on Step 4, which are used to conduct local training with the collaboration of other domains on the next t -th round training. The four steps are conducted iteratively until the global model convergence. Finally, the global model \mathcal{M}_G is deployed on the unseen domain $\{X_t\}$.

For learning a domain-invariant model from multiple decentralized source domains to generalize well on unseen target domain, we propose (1) Collaborative Style Augmentation (CSA) to explore the out-of-distribution styles with the collaboration of other domain classifier heads, and (2) domain-invariant learning between the original data and augmented data to learn the intrinsic semantic information within class by contrastive alignment and the relationship between classes by cross-domain relation matching.

3.2 Collaborative Style Augmentation

Inspired by the existing style augmentation methods, e.g. MixStyle[Zhou *et al.*, 2021b], the style information of the feature $f \in \mathbb{R}^{H \times W \times C}$ with the spatial size $H \times W$ can be revealed by the channel-wise mean μ and standard deviation σ :

$$\mu = \frac{1}{HW} \sum_{h \in H, w \in W} f_{h,w}, \quad (1)$$

$$\sigma = \sqrt{\frac{1}{HW} \sum_{h \in H, w \in W} (f_{h,w} - \mu)^2}. \quad (2)$$

For generating features with novel style, the original feature f is normalized by the channel-wise mean μ and standard deviation σ , then the normalized feature is scaled by the novel standard deviation $\hat{\sigma}$ and added by the novel mean $\hat{\mu}$:

$$\hat{f} = \hat{\sigma} \frac{f - \mu}{\sigma} + \hat{\mu}. \quad (3)$$

For generating novel statistic mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$, the existing style augmentation methods e.g. single-domain style exploration method DSU [Li *et al.*, 2021b] expands the mean μ and standard deviation σ by sampling perturbations from the estimated Gaussian distribution, the AdvStyle [Zhong *et al.*, 2022] expands μ and σ by maximizing the cross-entropy loss of the current model $\mathcal{F}_i \circ \mathcal{C}_i$ to learn the adversarial perturbations. However, these single-domain style exploration methods only explore the styles based on the current source domain and ignores the other decentralized source domains, which leads to limited diversity of styles. Although some FedDG methods e.g. CCST [Chen *et al.*, 2023] and StableFDG [Park *et al.*, 2023] share the mean μ and standard deviation σ across multiple decentralized source domains to conduct multi-domain style interpolation, the generated novel styles still come from the style space of existing source domains.

Different from these single-domain exploration methods and multi-domain interpolation methods, we propose collaborative style augmentation to generate the out-of-distribution styles with the collaboration of other source domains. Specifically, we use the classifier heads from other source domains as the discriminators to guide the generalization of novel

styles, as shown in the top of Figure 2(b). On the isolated domain D_i , the style statistics $\hat{\mu}$ and $\hat{\sigma}$ can be learned as follows:

$$\hat{\mu} = \mu + \frac{1}{n} \sum_{j=1}^n \eta \nabla_{\mu} \mathcal{L}'_{ce}(\mathcal{F}_i^l \circ \mathcal{C}_j^{t-1}; \hat{f}), \quad (4)$$

$$\hat{\sigma} = \sigma + \frac{1}{n} \sum_{j=1}^n \eta \nabla_{\sigma} \mathcal{L}'_{ce}(\mathcal{F}_i^l \circ \mathcal{C}_j^{t-1}; \hat{f}). \quad (5)$$

The $\mathcal{L}'_{ce} = -y \log(\delta(\mathcal{F}_i^l \circ \mathcal{C}_j^{t-1}(\hat{f})))$, δ is the softmax function. $\{\mathcal{C}_j^{t-1}\}_{j=1}^n$ are the classifier heads of decentralized source domains from the last round $t-1$. For the feature f extracted by hidden layers, we conduct adversarial style augmentation by using the rest neural network layers \mathcal{F}_i^l . In experiments, we analysis where are the best location to conduct feature augmentation.

In Equation 4 and Equation 5, the generated novel styles tend to be away from the decision boundary of existing classifiers from different source domains, so that the features with novel styles are out of the existing source domains.

3.3 Domain-invariant Learning

The original data and augmented data are used to train the local model with cross-entropy:

$$\mathcal{L}_{task} = \frac{1}{2} (\mathcal{L}_{ce}(\mathcal{F}_i \circ \mathcal{C}_i; x) + \mathcal{L}'_{ce}(\mathcal{F}_i^l \circ \mathcal{C}_i; \hat{f})). \quad (6)$$

For further improving the generalization ability of model, we conduct domain-invariant learning between original data and augmented data. Firstly, the cross-domain feature alignment is utilized to learn the compact representations within the class by contrastive alignment. Secondly, Cross-Domain Relationship Matching (CDRM) is proposed to learn the relationship between classes from the ensemble of multiple classifier heads.

Cross-domain feature alignment

We conduct supervised contrastive alignment [Khosla *et al.*, 2020] at feature level to improve the compactness of representations within the same class. For the i -th isolated domain:

$$\mathcal{L}_{SupCon} = - \sum_{j=0}^{2N_i} \frac{1}{|P(j)|} \sum_{p \in P(j)} \log \frac{e^{(z_j \cdot z_p / \tau)}}{\sum_{a \in A(j)} e^{(z_j \cdot z_a / \tau)}}. \quad (7)$$

The z_j are the feature of j -th image extracted by feature extractor \mathcal{F}_i . The $P(j)$ is the set of features with the same label as j -th image, $|P(j)|$ indicates the number of features in set $P(j)$. $A(j)$ indicates the features of original data and augmented data. τ is the temperature parameter, here we set 0.07 following previous work.

Cross domain relation matching

Furthermore, we conduct cross-domain relation matching to keep the intrinsic similarity between classes, e.g. the category of giraffe has a larger similarity with the elephant than the house. For obtaining the intrinsic relationship between

classes, we calculate the ensemble logits l_{ens} of the same class from multiple classifiers $\{C_i^{t-1}\}_{i=1}^n$. For example, on the m -th source domain, we calculate the l_{ens}^k of class k by averaging the logits of original images from different classifier heads within the same class $C(k)$:

$$l_{ens}^k = \sum_{i=1}^n \sum_{j \in C(k)} \mathcal{F}_m \circ C_i^{t-1}(x_j), \quad (8)$$

then l_{ens}^k is scaled with a temperature $\tau = 2.0$ by softmax function for smoothing the probability between classes, which can capture intrinsic semantic relationship between current class k and other classes. The smoothed l_{ens}^k of different classes are $\{p_{ens}^k\}_{k=1}^K$.

We also calculate the class relation $\{p_{cur}^k\}_{k=1}^K$ of the original images and $\{\hat{p}_{cur}^k\}_{k=1}^K$ of the stylized images from current model $\mathcal{F}_m \circ C_m$, which are used to match the ensemble class relation $\{p_{ens}^k\}_{k=1}^K$:

$$\mathcal{L}_{cdrm} = \sum_{k=1}^K \sum_{j=1}^K (p_{ens}^{k(j)} \log p_{cur}^{k(j)} + p_{ens}^{k(j)} \log \hat{p}_{cur}^{k(j)}), \quad (9)$$

The overall loss of semantic learning on local clients is as follows:

$$\mathcal{L}_{local} = \mathcal{L}_{task} + \lambda_{con} \mathcal{L}_{con} + \lambda_{cdrm} \mathcal{L}_{cdrm}, \quad (10)$$

where λ_{con} and λ_{cdrm} are the hyper-parameters to balance different losses.

3.4 Model Aggregation

On the server side, we aggregate different source domain models $\{\mathcal{M}_i\}_{i=1}^n$ by parameter averaging:

$$\mathcal{M}_G = \sum_{i=1}^n \frac{N_i}{N_{total}} \mathcal{M}_i, \quad (11)$$

the N_{total} is the sum of all samples from different source domains. The aggregated global model \mathcal{M}_G is used as the initial local model on local clients for the next round of training.

4 Experiments

4.1 Datasets

In this work, we follow the previous DG works [Zhou *et al.*, 2021b; Wu and Gong, 2021] to evaluate our method on three image classification datasets: **PACS**, **Office-Home**, and **VLCS**. **PACS** contains 9,991 images of 7 categories from four domains, Art-Painting (Art), Cartoon, Photo, and Sketch. Following previous works [Zhou *et al.*, 2021b], we split the data of each domain into 80% for training and 20% for testing. **Office-Home** contains 15,500 images of 65 categories from four domains, Artistic (Art), Clipart, Product, and Real-World (Real). Following the previous works [Zhou *et al.*, 2021b], we split each domain into 90% as the training set and 10% as the test set. **VLCS** contains 10,729 images of 5 categories from four domains, Pascal, LabelMe, Caltech, and

Sun. Following previous works [Chen *et al.*, 2023], we split the data of each domain into 80% for training and 20% for testing.

We utilize the leave-one-domain-out protocol [Zhou *et al.*, 2021b] to select a domain as the unseen target domain for evaluation and the rest as the source domains for training.

4.2 Implementation Details

Following previous works [Yuan *et al.*, 2023; Wu and Gong, 2021], we use the pre-trained ResNet-18 on ImageNet as the backbone for PACS, Office-Home, and VLCS dataset. The SGD optimizer is used to optimize the network with momentum 0.9 and weight decay $5e-4$. The initial learning rate is 0.001 decayed by the cosine schedule to 0.0001 for the PACS and VLCS datasets. For the Office-Home dataset, the initial learning rate is 0.002 and decayed to 0.0001. The batch size is 16 for PACS and VLCS, and 30 for Office-Home. The adversarial learning rate η is 1.0 for PACS and VLCS dataset, 0.3 for Office-Home datasets. The λ_{SupCon} is 1.0 for PACS and Office-Home dataset, 0.3 for VLCS dataset. The λ_{cdrm} is 4.0 for PACS, 0.7 for Office-Home, and 0.3 for VLCS dataset. The τ of \mathcal{L}_{cdrm} is 1.5 for all dataset. The values of hyper-parameters are set according to the performance on validation set of source domains.

We train the local model on each client 1 epoch in Step 1 and then upload the local models to the server side for model aggregation. The total communication rounds between clients and server is 40 for PACS, Office-Home, VLCS. All experiments are conducted three times with different random seeds, and the mean accuracy (%) is reported.

For evaluating the effectiveness of our method, we make a comparison with the state-of-the-art DG methods. The DeepAll indicates training the model by cross-entropy on the centralized source domains. The FedAvg [McMahan *et al.*, 2017] indicates collaborative training the decentralized source domains. As shown in Table 1, Table 2, and Table 3, we report the accuracy (%) of state-of-the-art methods under the data centralized and decentralized scenario, where the **Dec.** indicates the data decentralization scenario.

4.3 Comparison With the State-of-the-Art Methods

As shown in Table 1, Table 2, and Table 3, our method MCSAD achieves the best average accuracy under the data decentralization scenario on three datasets, PACS, Office-Home, and VLCS. Compared with the style interpolation methods, e.g. ELCFS and CCST, our method can explore the broader style space, which leads to large performance improvement. FADH trains the additional image generators on each source domain by maximizing the entropy of the global model and minimizing the cross-entropy of the local model. Compared with FADH, our method MCSAD is simple but effective. As shown in Table 1 and Table 2, our method outperforms FADH significantly. COPA learns the domain-invariant model with the collaboration of multiple classifier heads, and the ensemble of classifier heads from different domains is used to deploy on the unseen domain. However, our method only utilizes a global model to deploy on unseen domain and achieves better performance, as shown in Table 1, Table 2, and Table

Methods	Dec.	Art	Cartoon	Photo	Sketch	Avg
DeepAll [Zhou <i>et al.</i> , 2021a]	×	77.0	75.9	96.0	69.2	79.5
JiGen [Carlucci <i>et al.</i> , 2019]	×	79.4	75.3	96.0	71.4	80.5
EISNet [Wang <i>et al.</i> , 2020]	×	81.9	76.4	95.9	74.3	82.2
MASF [Dou <i>et al.</i> , 2019]	×	80.3	77.2	95.0	71.7	81.0
DAEL [Zhou <i>et al.</i> , 2021a]	×	<u>84.6</u>	74.4	95.6	78.9	83.4
L2A-OT [Zhou <i>et al.</i> , 2020b]	×	83.3	78.2	96.2	73.6	82.8
DDAIG [Zhou <i>et al.</i> , 2020a]	×	84.2	78.1	95.3	74.7	83.1
FACT [Xu <i>et al.</i> , 2021]	×	85.4	78.4	95.2	79.2	84.5
MixStyle [Zhou <i>et al.</i> , 2021b]	×	84.1	78.8	96.1	75.9	83.7
EFDMix [Zhang <i>et al.</i> , 2022]	×	83.9	79.4	96.8	75.0	83.9
DSU [Li <i>et al.</i> , 2021b]	×	83.6	79.6	95.8	77.6	84.1
StyleNeo [Kang <i>et al.</i> , 2022]	×	84.4	79.2	94.9	83.2	85.4
RSC [Huang <i>et al.</i> , 2020]	×	83.4	80.3	96.0	80.9	85.2
I ² -ADR [Meng <i>et al.</i> , 2022]	×	82.9	<u>80.8</u>	95.0	83.5	85.6
CDG [Du <i>et al.</i> , 2022]	×	83.5	80.1	95.6	83.8	<u>85.8</u>
FedAvg [McMahan <i>et al.</i> , 2017]	✓	79.7	75.6	94.7	81.1	82.8
CASC [Yuan <i>et al.</i> , 2023]	✓	82.0	76.4	95.2	81.6	83.8
GA [Zhang <i>et al.</i> , 2023]	✓	83.2	76.9	94.0	82.9	84.3
ELCFS [Liu <i>et al.</i> , 2021]	✓	82.3	74.7	93.3	82.7	83.2
CCST [Chen <i>et al.</i> , 2023]	✓	81.3	73.3	95.2	80.3	82.5
FADH [Xu <i>et al.</i> , 2023b]	✓	83.8	77.2	94.4	<u>84.4</u>	85.0
StableFDG [Park <i>et al.</i> , 2023]	✓	83.0	79.3	94.9	79.8	84.2
COPA [Wu and Gong, 2021]	✓	83.3	79.8	94.6	82.5	85.1
MCSAD (ours)	✓	84.2	81.2	95.1	84.6	86.3

Table 1: Accuracy(%) on PACS dataset. We have bolded the best results and underlined the second results.

Methods	Dec.	Art	Clipart	Product	Real	Avg
DeepAll [Zhou <i>et al.</i> , 2021a]	×	57.9	52.7	73.5	74.8	64.7
JiGen [Carlucci <i>et al.</i> , 2019]	×	53.0	47.5	71.5	72.8	61.2
EISNet [Wang <i>et al.</i> , 2020]	×	56.8	53.3	72.3	73.5	64.0
DAEL [Zhou <i>et al.</i> , 2021a]	×	59.4	55.1	74.0	<u>75.7</u>	<u>66.1</u>
L2A-OT [Zhou <i>et al.</i> , 2020b]	×	60.6	50.1	74.8	77.0	65.6
DDAIG [Zhou <i>et al.</i> , 2020a]	×	59.2	52.3	74.6	<u>76.0</u>	65.5
FACT [Xu <i>et al.</i> , 2021]	×	60.3	54.9	74.5	<u>76.6</u>	66.6
MixStyle [Zhou <i>et al.</i> , 2021b]	×	58.7	53.4	74.2	75.9	65.5
StyleNeo [Kang <i>et al.</i> , 2022]	×	59.6	55.0	73.6	75.5	65.9
RSC [Huang <i>et al.</i> , 2020]	×	58.4	47.9	71.6	74.5	63.1
CDG [Du <i>et al.</i> , 2022]	×	59.2	54.3	74.9	75.7	66.0
DomainDrop [Guo <i>et al.</i> , 2023]	×	59.6	55.6	74.5	76.6	66.6
FedAvg [McMahan <i>et al.</i> , 2017]	✓	58.2	51.6	73.1	73.8	64.2
GA [Zhang <i>et al.</i> , 2023]	✓	58.8	54.3	73.7	74.7	65.4
ELCFS [Liu <i>et al.</i> , 2021]	✓	57.8	54.9	71.1	73.1	64.2
CCST [Chen <i>et al.</i> , 2023]	✓	59.1	50.1	73.0	71.7	63.6
FADH [Xu <i>et al.</i> , 2023b]	✓	<u>59.9</u>	55.8	73.5	74.9	66.0
StableFDG [Park <i>et al.</i> , 2023]	✓	<u>57.2</u>	<u>57.9</u>	72.8	72.2	65.0
COPA [Wu and Gong, 2021]	✓	59.4	55.1	74.8	75.0	<u>66.1</u>
MCSAD (ours)	✓	59.4	58.8	75.0	75.4	67.2

Table 2: Accuracy(%) on Office-Home dataset. We have bolded the best results and underlined the second results.

Methods	Dec.	Pascal	LabelMe	Caltech	Sun	Avg
DeepAll [Zhou <i>et al.</i> , 2021a]	×	71.4	59.8	97.5	69.0	74.4
JiGen [Carlucci <i>et al.</i> , 2019]	×	<u>74.0</u>	61.9	97.4	66.9	75.1
L2A-OT [Zhou <i>et al.</i> , 2020b]	×	72.8	59.8	98.0	70.9	75.4
MixStyle [Zhou <i>et al.</i> , 2021b]	×	72.6	58.5	97.7	73.3	75.5
RSC [Huang <i>et al.</i> , 2020]	×	75.3	59.8	97.0	71.5	75.9
DomainDrop [Guo <i>et al.</i> , 2023]	×	76.4	64.0	98.9	73.7	78.3
FedAvg [McMahan <i>et al.</i> , 2017]	✓	72.0	63.3	96.5	72.4	76.0
CASC [Yuan <i>et al.</i> , 2023]	✓	72.0	63.5	97.2	72.1	76.2
ELCFS [Liu <i>et al.</i> , 2021]	✓	71.1	59.5	96.6	<u>74.0</u>	75.3
COPA [Wu and Gong, 2021]	✓	71.5	61.0	93.8	71.7	74.5
MCSAD (ours)	✓	76.1	65.6	<u>98.6</u>	75.0	78.8

Table 3: Accuracy(%) on VLCS dataset. We have bolded the best results and underlined the second results.

3. Compared with the methods which focus on the model aggregation stage, e.g. CASC and GA, our method learns the domain-invariant representations on local clients and outperforms these methods largely.

Although the data from different domains are kept decentralized, our method even outperforms the state-of-the-art methods e.g. MixStyle, RSC, and StyleNeo, which can access multiple source domain data simultaneously.

4.4 Ablation Study

Contributions of Different Components. As shown in Table 4, we conduct ablation study about the different components of our method on PACS dataset. Compared with the baseline, the Collaborative Style Augmentation (CSA) can improve the average accuracy largely. Combining with the \mathcal{L}_{SupCon} and \mathcal{L}_{cdrm} , the performance can be further improved by learning the compact representations and the class relationship. By alternative conduct style augmentation and semantic learning with \mathcal{L}_{SupCon} and \mathcal{L}_{cdrm} , we achieve the best average accuracy.

The CSA proposed in this work augments the feature in hidden layer. We also conduct experiments to validate the effectiveness of conducting CSA on different layers. As shown in Table 5, we conduct CSA after the different blocks of ResNet-18. Compared with the baseline method, conducting CSA after different blocks can improve the accuracy on unseen target domain. Due to the shallow layer of Deep Neural Network extracting the features with more style information, conducting CSA after the first and second Blocks can achieve better accuracy. When conducting CSA after the first and second Blocks, we can achieve the best average accuracy 85.0%.

Comparison with Different Style Augmentation Methods. We also make a comparison with other data augmentation methods under the FedAvg framework, including (1) single-domain style expanding methods, e.g. DSU and AdvStyle, (2) multi-domain style interpolation method, e.g. AM, MixStyle, and EFDMix. (3) pre-defined augmentation pool, e.g. RandAug, and (4) the methods FADH, which train image generator.

CSA	\mathcal{L}_{SupCon}	\mathcal{L}_{cdrm}	Art	Cartoon	Photo	Sketch	Avg
-	-	-	79.7	75.6	94.7	81.1	82.8
✓	-	-	82.3	80.2	95.4	82.2	85.0
✓	✓	-	83.4	80.9	95.7	83.9	86.0
✓	-	✓	82.5	80.1	95.5	83.6	85.4
✓	✓	✓	84.2	81.2	95.1	84.6	86.3

Table 4: Accuracy(%) of each component on PACS dataset.

Block1	Block2	Block3	Art	Cartoon	Photo	Sketch	Avg
-	-	-	79.7	75.6	94.7	81.1	82.8
✓	-	-	81.4	78.3	95.1	83.0	84.5
-	✓	-	81.7	79.8	95.4	81.3	84.6
-	-	✓	82.7	79.3	95.3	80.3	84.4
✓	✓	-	82.3	80.2	95.4	82.2	85.0
✓	✓	✓	80.9	78.1	94.7	83.2	84.2

Table 5: Ablation study about the location of collaborative style augmentation on PACS dataset.

Methods	Art	Cartoon	Photo	Sketch	Avg
FedAvg [McMahan <i>et al.</i> , 2017]	79.7	75.6	94.7	81.1	82.8
+ DSU [Li <i>et al.</i> , 2021b]	81.5	77.9	95.5	81.6	84.1
+ AdvStyle [Zhong <i>et al.</i> , 2022]	80.0	75.9	94.7	82.0	83.2
+ AM [Zhang <i>et al.</i> , 2023]	83.4	76.2	95.4	80.8	84.0
+ MixStyle [Zhou <i>et al.</i> , 2021b]	82.1	76.8	95.3	82.3	84.1
+ EFDMix [Zhang <i>et al.</i> , 2022]	81.9	78.3	94.7	82.5	84.4
+ RandAug [Wu and Gong, 2021]	83.6	76.1	95.9	81.4	84.3
+ FADH [Xu <i>et al.</i> , 2023b]	82.8	77.2	93.9	84.1	84.5
+ CSA (ours)	82.3	80.2	95.4	82.2	85.0

Table 6: Accuracy (%) of different style augmentation methods on PACS dataset.

As shown in Table 6, the CSA method proposed in this work outperforms the single-domain style expanding methods and multi-domain style interpolation methods on average accuracy largely. Our method can generate novel styles out of the existing source domains with the collaboration of multiple classifier heads, the single-domain style expanding methods or multi-domain style interpolation methods only lead to limited styles in the style space of the existing source domains.

The RandAug generates images with pre-defined augmentation strategies, e.g. image processing, which cannot dynamically generate novel styles out of the existing source domains. As shown in Table 6, RandAug performs poorly on the Sketch domain. Different from RandAug, our CSA method generates novel styles by adversarial style augmentation, which can explore the out-of-distribution styles in an online manner. As shown in Table 6, our CSA also outperforms RandAug in average accuracy.

The FADH trains the image generators on multiple decentralized source domains and achieves the average accuracy 84.5%. Different from FADH, our method CSA diversify the features by adversarial style augmentation, which leads to better average accuracy. Furthermore, compared with FADH, our method leads to smaller computational and storage cost.

Different Distillation Strategies. In the domain-invariant learning stage, we propose Cross-Domain Relation Matching (CDRM) \mathcal{L}_{cdrm} loss to distill the intrinsic relation between classes.

There are different distillation strategies, including (1) Kullback-Leibler (KL) divergence used by [Kang *et al.*, 2022; Xu *et al.*, 2021], here we align the logits of augmented features from current model to the ensembled logits of original features from multiple classifier heads, (2) Collaboration of Frozen Classifiers (CoFC) used by COPA [Wu and Gong, 2021], the CoFC freezes the classifier heads from other domains to learn the domain-invariant feature extractor. (3) Cross-Domain Relation Matching (CDRM) proposed in this work, the class relation from original feature and augmented feature are aligned to the ensembled classes relation, as shown in Equation 9.

Method	Art	Cartoon	Photo	Sketch	Avg
Baseline	82.3	80.2	95.4	82.2	85.0
+ KL [Kang <i>et al.</i> , 2022]	82.1	78.9	94.5	84.5	85.0
+ CoFC [Wu and Gong, 2021]	81.2	78.2	95.1	84.1	84.7
+ CDRM (ours)	82.5	80.1	95.5	83.6	85.4

Table 7: Accuracy (%) of different distillation strategies.

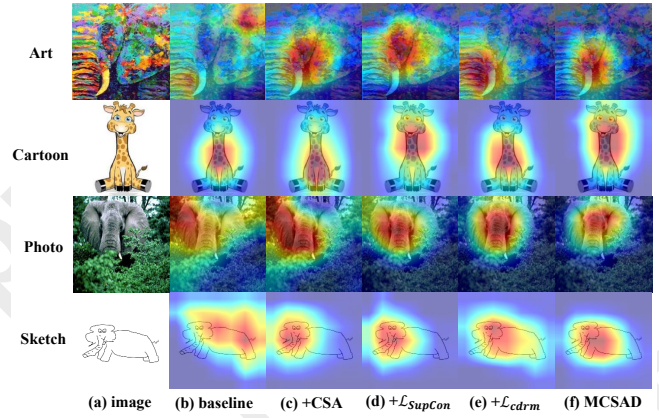


Figure 3: Visualization by Grad-CAM on unseen target domain.

As shown in Table 7, we report the results of different distillation strategies by combining the CSA method proposed in this work. The KL lead to limited improvement on the average accuracy, and the CoFC even leads to degraded performance on average accuracy. On the Art and Photo domains, KL and CoFC cannot achieve improvement. The CDRM strategy proposed in this work outperforms KL and CoFC significantly by distilling the relationship between classes to improve the generalization ability of model.

Visualization. Furthermore, we also visualize the activated region on input images for classification by Grad-CAM [Selvaraju *et al.*, 2017]. As shown in Figure 3, the learned baseline model usually focuses on the texture or background region on the unseen domain. By using CSA, the learned model can capture the region of the object for classification. Combined with \mathcal{L}_{SupCon} and \mathcal{L}_{cdrm} , the learned model tends to capture the discriminative and overall region of the object, which leads to better performance on the unseen domain. By combining all components, the most discriminative region can be focused on for better generalization performance on unseen domain.

5 Conclusion

In this paper, we propose a MCSAD method to solve the multi-source domain generalization problem under the data decentralization scenario. To explore the out-of-distribution styles on the decentralized source domains, we propose a multi-source collaborative style augmentation method to generate features with novel styles for diversifying the source domains. Moreover, we propose the domain-invariant learning between the original data and augmented data to learn the domain-invariant representations and improve the generalization ability of model. The extensive experiments can validate the effectiveness of the proposed MCSAD method.

References

[Carlucci *et al.*, 2019] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition*, pages 2229–2238, 2019.
- [Chen *et al.*, 2023] Junming Chen, Meirui Jiang, Qi Dou, and Qifeng Chen. Federated domain generalization for image recognition via cross-client style transfer. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 361–370, 2023.
- [Cubuk *et al.*, 2020] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020.
- [Dou *et al.*, 2019] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Du *et al.*, 2022] Dapeng Du, Jiawei Chen, Yuexiang Li, Kai Ma, Gangshan Wu, Yefeng Zheng, and Limin Wang. Cross-domain gated learning for domain generalization. *International Journal of Computer Vision*, 130(11):2842–2857, 2022.
- [Guo *et al.*, 2023] Jintao Guo, Lei Qi, and Yinghuan Shi. Domaindrop: Suppressing domain-sensitive channels for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19114–19124, 2023.
- [Huang and Belongie, 2017] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017.
- [Huang *et al.*, 2020] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 124–140. Springer, 2020.
- [Huang *et al.*, 2024] Wenke Huang, Mang Ye, Zekun Shi, Guancheng Wan, He Li, Bo Du, and Qiang Yang. Federated learning for generalization, robustness, fairness: A survey and benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [Kang *et al.*, 2022] Juwon Kang, Sohyun Lee, Namyup Kim, and Suha Kwak. Style neophile: Constantly seeking novel styles for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7130–7140, 2022.
- [Khosla *et al.*, 2020] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020.
- [Li *et al.*, 2021a] Lei Li, Ke Gao, Juan Cao, Ziyao Huang, Yepeng Weng, Xiaoyue Mi, Zhengze Yu, Xiaoya Li, and Boyang Xia. Progressive domain expansion network for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 224–233, 2021.
- [Li *et al.*, 2021b] Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, and LINGYU DUAN. Uncertainty modeling for out-of-distribution generalization. In *International Conference on Learning Representations*, 2021.
- [Li *et al.*, 2024] Deng Li, Aming Wu, Yaowei Wang, and Yehong Han. Prompt-driven dynamic object-centric learning for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17606–17615, 2024.
- [Liu *et al.*, 2021] Quande Liu, Cheng Chen, Jing Qin, Qi Dou, and Pheng-Ann Heng. Feddg: 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*, pages 1013–1023, 2021.
- [McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.
- [Meng *et al.*, 2022] Rang Meng, Xianfeng Li, Weijie Chen, Shicai Yang, Jie Song, Xinchao Wang, Lei Zhang, Mingli Song, Di Xie, and Shiliang Pu. Attention diversification for domain generalization. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIV*, pages 322–340. Springer, 2022.
- [Park *et al.*, 2023] Jungwuk Park, Dong-Jun Han, Jinho Kim, Shiqiang Wang, Christopher Brinton, and Jaekyun Moon. Stablefdg: Style and attention based learning for federated domain generalization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Selvaraju *et al.*, 2017] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [Tang *et al.*, 2024a] Luyao Tang, Yuxuan Yuan, Chaoqi Chen, Xinghao Ding, and Yue Huang. Mixstyle-entropy: Whole process domain generalization with causal intervention and perturbation. In *35th British Machine Vision Conference 2024, BMVC 2024, Glasgow, UK, November 25–28, 2024. BMVA*, 2024.
- [Tang *et al.*, 2024b] Luyao Tang, Yuxuan Yuan, Chaoqi Chen, Kunze Huang, Xinghao Ding, and Yue Huang. Bootstrap segmentation foundation model under distribution shift via object-centric learning. *arXiv preprint arXiv:2408.16310*, 2024.

- [Tang *et al.*, 2025] Luyao Tang, Yuxuan Yuan, Chaoqi Chen, Zeyu Zhang, Yue Huang, and Kun Zhang. Ocrt: Boosting foundation models in the open world with object-concept-relation triad. *arXiv preprint arXiv:2503.18695*, 2025.
- [Wang *et al.*, 2020] Shujun Wang, Lequan Yu, Caizi Li, Chi-Wing Fu, and Pheng-Ann Heng. Learning from extrinsic and intrinsic supervisions for domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX*, pages 159–176. Springer, 2020.
- [Wang *et al.*, 2021] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh. Learning to diversify for single domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 834–843, 2021.
- [Wei and Han, 2022] Yikang Wei and Yahong Han. Dual collaboration for decentralized multi-source domain adaptation. *Frontiers of Information Technology & Electronic Engineering*, 23(12):1780–1794, 2022.
- [Wei and Han, 2023] Yikang Wei and Yahong Han. Exploring instance relation for decentralized multi-source domain adaptation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [Wei and Han, 2024] Yikang Wei and Yahong Han. Multi-source collaborative gradient discrepancy minimization for federated domain generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 15805–15813, 2024.
- [Wei *et al.*, 2023] Yikang Wei, Liu Yang, Yahong Han, and Qinghua Hu. Multi-source collaborative contrastive learning for decentralized domain adaptation. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(5):2202–2216, 2023.
- [Wu and Gong, 2021] Guile Wu and Shaogang Gong. Collaborative optimization and aggregation for decentralized domain generalization and adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6484–6493, 2021.
- [Wu *et al.*, 2024] Aming Wu, Jiaping Yu, Yuxuan Wang, and Cheng Deng. Prototype-decomposed knowledge distillation for learning generalized federated representation. *IEEE Transactions on Multimedia*, 2024.
- [Xu *et al.*, 2021] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14383–14392, 2021.
- [Xu *et al.*, 2023a] Qinwei Xu, Ruipeng Zhang, Yi-Yan Wu, Ya Zhang, Ning Liu, and Yanfeng Wang. Simde: A simple domain expansion approach for single-source domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4797–4807, 2023.
- [Xu *et al.*, 2023b] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yi-Yan Wu, and Yanfeng Wang. Federated adversarial domain hallucination for privacy-preserving domain generalization. *IEEE Transactions on Multimedia*, 2023.
- [Yuan *et al.*, 2023] Junkun Yuan, Xu Ma, Defang Chen, Fei Wu, Lanfen Lin, and Kun Kuang. Collaborative semantic aggregation and calibration for federated domain generalization. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [Zhang *et al.*, 2022] Yabin Zhang, Minghan Li, Ruihuang Li, Kui Jia, and Lei Zhang. Exact feature distribution matching for arbitrary style transfer and domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8035–8045, 2022.
- [Zhang *et al.*, 2023] Ruipeng Zhang, Qinwei Xu, Jiangchao Yao, Ya Zhang, Qi Tian, and Yanfeng Wang. Federated domain generalization with generalization adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3954–3963, 2023.
- [Zhong *et al.*, 2022] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial style augmentation for domain generalized urban-scene segmentation. *Advances in Neural Information Processing Systems*, 35:338–350, 2022.
- [Zhou *et al.*, 2020a] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13025–13032, 2020.
- [Zhou *et al.*, 2020b] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 561–578. Springer, 2020.
- [Zhou *et al.*, 2021a] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain adaptive ensemble learning. *IEEE Transactions on Image Processing*, 30:8008–8018, 2021.
- [Zhou *et al.*, 2021b] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021.