

## Domain Generalization with Correlated Style Uncertainty

Zheyuan Zhang<sup>1</sup>, Bin Wang<sup>1</sup>, Debesh Jha<sup>1</sup>, Ugur Demir<sup>1</sup>, Ulas Bagci<sup>1</sup>  
<sup>1</sup> Machine & Hybrid Intelligence Lab, Northwestern University, USA

ulas.bagci@northwestern.edu

### Abstract

Domain generalization (DG) approaches intend to extract domain invariant features that can lead to a more robust deep learning model. In this regard, style augmentation is a strong DG method taking advantage of instance-specific feature statistics containing informative style characteristics to synthetic novel domains. While it is one of the state-of-the-art methods, prior works on style augmentation have either disregarded the interdependence amongst distinct feature channels or have solely constrained style augmentation to linear interpolation. To address these research gaps, in this work, we introduce a novel augmentation approach, named *Correlated Style Uncertainty (CSU)*, surpassing the limitations of linear interpolation in style statistic space and simultaneously preserving vital correlation information. Our method’s efficacy is established through extensive experimentation on diverse cross-domain computer vision and medical imaging classification tasks: PACS, Office-Home, and Camelyon17 datasets, and the Duke-Market1501 instance retrieval task. The results showcase a remarkable improvement margin over existing state-of-the-art techniques. The source code is available <https://github.com/freshman97/CSU>.

### 1. Introduction

Recent years have witnessed the remarkable success of Deep learning (DL) in computer vision domain operating under the premise that the training (source) and testing (target) datasets adhere to a principle of independent and identically distributed (iid) data [48]. However, this oversimplified assumption often fails in practice when there is a distribution drift between training and testing datasets. The violation of this assumption induces the phenomenon that well-trained model in the source domain degrades dramatically in the target domain. If domain generalization is successfully implemented within a DL model, it would inherently solve issues related to domain shift. This would not only signify a significant advancement in the field, but also streamline the practical deployment of DL models across various domains. For example, a car detector should perform accurately both on sunny and cloudy days. DL based

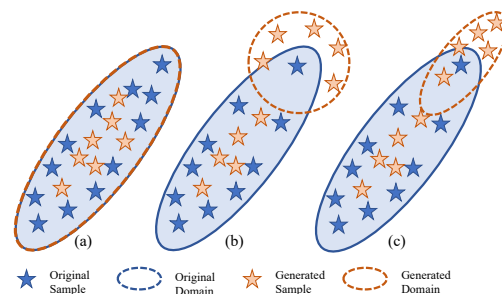


Figure 1. Visualization of synthetic feature statistics samples using (a) MixStyle [50], (b) DSU [18], and (c) Our proposed Correlated Style Uncertainty (CSU) method. CSU preserves the correlation among feature channels.

medical image segmentation algorithm, for another example, should generate stable segmentation regardless of the acquisition and scanner differences, and so on.

**Domain adaptation** and **domain generalization** are two distinct approaches for addressing the challenge of domain shift in machine learning. A widely adopted approach to counteract domain shift problems entails acquiring unlabeled data from the target domain. By leveraging this data, we can adapt a model, initially trained on the source domain, to align better with the characteristics of the target domain. This strategy is called Domain Adaptation (DA), which has been the subject of much systematic investigation in the last few years and achieved promising results in many fields [48]. To ensure robust model adaptation, regularization algorithms such as entropy regularization [24, 26] are also applied to target domain during training. Domain adaptation assumes that the target domain is known and that some data from the target domain is available for training. However, accessing the target domain data can be quite challenging. Specifically in high-risk applications (e.g., medical data), target domain data might not be available at all. This is a primary research concern of Domain Generalization (DG) as a strong alternative path to DA.

In DG, the objective is to develop models that can perform well across a wide range of domains without explicit training on each individual domain. Unlike DA, DG assumes that the target domain is unknown and that

no data from the target domain is available for training [34, 48]. Since target data is unavailable for investigating domain shift, DG relies solely on extracting robust, domain-invariant feature representations from diverse training distributions. As a result, when domain information is feasible, feature alignments among different domains could significantly leverage the model’s out-of-distribution generalization ability, as shown in previous research [16, 19].

**What do we propose?** In this paper, we propose a novel DG method, called *Correlated Style Uncertainty (CSU)*. The method is specifically designed to retain the relationship among distinct feature space channels while addressing the distribution drift between target and source domains. The effectiveness of our approach, as depicted in Figure 1(c), is clearly evident when compared with widely-used techniques like MixStyle and DSU. Like MixStyle and DSU, our newly introduced algorithm, CSU, belongs to the family of *style augmentation methods* employed for DG, but it presents distinctive advantages.

**How is CSU different from previous style augmentation?** In the context of style augmentation, MixStyle [50] operates by randomly choosing two samples and performing linear interpolation between their corresponding style attributes, to put it succinctly. While MixStyle is useful, it’s confined to generating in-distribution samples, potentially limiting network’s generalization capabilities. pAdaIN [22] rearranges the instance-specific feature statistics within a batch, thus sharing the same problem with MixStyle. DSU [18], contrarily, addresses this limitation by creating out-of-distribution samples via uncertainty modeling of feature statistics. However, DSU’s effectiveness is hinged upon the assumption that each channel operates independently, implying that inter-channel correlations bear no impact on the task at hand. This assumption might not always hold in practical scenarios. Style Neophile [11] select style prototypes that represent the distribution of source styles stored in the source style queue using MMD distance from style storage during training. [42, 46] apply the adversarial attack on the feature statistics for domain generalization while introducing expensive computational burden.

Our proposed method, CSU, is fundamentally different from these methods because i) it produces out-of-distribution samples, ii) correlation among different channels are retained and iii) the sampling process requires a negligible computational burden. Our proposed method, CSU, allows us to generate more meaningful style perturbations in discriminatory tasks, thus enhancing the model’s generalization capacity. Similar to hypotheses in prior research in the literature, we posit that the feature statistics follow a multivariate Gaussian distribution; however, we acknowledge and account for correlations among the variates, unlike many traditional methods. This perspective allows us to consider a more realistic distribution that could poten-

tially improve model performance in real-world scenarios.

In our specific approach, we first compute the covariance matrix at the mini-batch level and then estimate the distribution from the covariance matrix. This allows us to sample correlated feature statistics from the established distribution. This sampling allows us to generate the style statistics outside the linear interpolation while maintaining an identical correlation. In this way, more diverse but meaningful style augmentation can be applied during the training and increase the model’s generalization ability. We highlight our main contributions in this study as follows:

- Our proposed strategy (CSU) is a well-calibrated method that goes beyond the interpolation strategies by preserving correlation between different feature spaces. This allows us to generate more diverse and meaningful style augmentation during training which helps in building a more generalizable model. To the best of our knowledge, such a simple yet effective style augmentation strategy has never been explored.
- To evaluate the effectiveness of the proposed CSU model, we conducted extensive experiments on multi-domain classification benchmarking datasets, including PACS [14], Office-Home [31], Camelyon17 [1] and the Duke-Market1501 dataset for instance retrieval tasks [25, 45]. The quantitative results show that the CSU can significantly improve the model’s generalizability over other state-of-the-art (SOTA) methods.
- We have performed several ablation studies to investigate the optimal parameters for the CSU model. These investigations have covered factors such as the ideal position for model integration, the most efficient sampling hyperparameters, and the batch size that yields the greatest level of generalization.

## 2. Background and Related Works

Several DG strategies have been proposed in the literature; we briefly cover the mostly used ones under sub-categories as follows.

**Data Augmentation:** is highly valued for its role in exposing the model to an expanded set of instances, an essential element for successful Deep Learning. Many methods have been proposed to achieve strong data augmentation, including traditional image augmentation like BigAug [40], deep neural network-based image generation like in Rand-Conv [36], and adversarial data augmentations [33]. These methods are suitable specifically when the domain tags of samples are agnostic. While data augmentation is a powerful technique that can improve the generalization performance of machine learning models, it also has some potential drawbacks in the context of domain generalization. For instance, one drawback is that data augmentation may not

be effective when the variations between the source and target domains are too significant. **Feature Alignment:** is a popular method in representation learning category of DG approaches. Given domain tags, the model will add regularization terms into loss functions to force the extracted features from all source domains to align to the same distribution. For instance, Li [15] introduced the Maximum Mean Discrepancy as a regularization term to achieve feature alignment across multiple domains. Zhao [44] proposed an entropy regularization term that measures the dependency between the learned features and corresponding labels. This regularization method can ensure the conditional invariance of learned features. One major drawback of feature alignment is that it can be difficult to determine the best way to align the features across domains. **Meta-learning:** has also attracted attention from DG communities [4, 29]. Meta-learning aims to learn the learning algorithm itself by learning from previous experience or tasks. By splitting the source domain samples into pseudo-train and pseudo-test, meta-learning mimics the potential domain shift of the actual target domain. Despite its promise, meta-learning can be computationally expensive and time-consuming. Since meta-learning involves training a model on a large number of tasks or domains, there is a risk that the model may overfit to the training data and not generalize well to new domains.

**Style Augmentation:** The final category of DG is very recent: *style augmentation*. This method comes from the simple observation that instance-specific feature statistics such as mean and standard deviation, contain informative style characteristics and can be applied to the style-transferring model [9]. This phenomenon allows us to generate different style images while maintaining the same semantic concept. For example, Seo et al. [27] proposed one domain-specific normalization method by calculating the feature statistics of each domain. Pan et al. [17] add perturbation on the feature embedding with simple Gaussian noise during training. Zhou et al. [50] presented mixing styles (MixStyle) of training instances, and increased the source domain diversity. As a result, authors leveraged the trained model’s generalizability. Nuriel et al. [22] alternatively proposed a Permuted Adaptive Instance Normalization (pAdaIN) method to rearrange the instance-specific feature statistics within a batch, thus improving the model’s generalizability. In a slightly different angle, Li et al. [18] quantified feature statistics’ uncertainty (DSU) and sampled new style feature statistics from the uncertainty distribution, resulting in novel out-of-distribution domains being synthesized implicitly. Kang et al. [11] select style prototypes from style storage during training using the MMD distance and these prototypes represent the distribution of source styles stored in the source style queue. [42, 46] apply the adversarial attack on the feature statistics for domain

generalization while introducing expensive computational burden. Our work is directly related to MixStyle and DSU as we are using the same hypothesis but addressing the major limitations of them, and sharing some similarities with pAdaIN, from the same efforts for synthesizing novel domains. Unlike them, our proposed CSU generates out-of-distribution feature statistics while maintaining the correlation between features.

### 3. Methods

#### 3.1. Correlation within the style statistics

Given batch level feature maps  $x \in \mathbb{R}^{B \times C \times H \times W}$  of the network  $f(in, \phi)$  where  $in$  denotes the batch-wise inputs and  $\phi$  denotes the network parameters. We can formulate the instance-specific feature statistics mean  $\mu \in \mathbb{R}^{B \times C}$  and standard deviation  $\sigma \in \mathbb{R}^{B \times C}$  as follows

$$\mu(x) = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W x_{b,c,h,w}, \quad (1)$$

$$\sigma^2(x) = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W (x_{b,c,h,w} - \mu(x))^2. \quad (2)$$

Thus, we can formulate the channel-wise covariance matrix  $\Sigma_\mu \in \mathbb{R}^{C \times C}$ ,  $\Sigma_\sigma \in \mathbb{R}^{C \times C}$  of  $\mu, \sigma$ :

$$\Sigma_\mu = \frac{1}{B} (\mu - E(\mu))^T (\mu - E(\mu)), \quad (3)$$

$$\Sigma_\sigma = \frac{1}{B} (\sigma - E(\sigma))^T (\sigma - E(\sigma)), \quad (4)$$

where  $E(\mu), E(\sigma)$  represents the mean value of  $\mu, \sigma$  over batch dimension. It is worth noting that the rank of  $\Sigma_\mu, \Sigma_\sigma$  is strictly limited by  $\min(B, C) \leq C$ . Previous research indicated that the feature maps are unlikely to be linearly independent over the channel dimension [8, 41]. Without any form of regularization, it becomes difficult to assume a diagonal covariance matrix and a zero correlation across each channel, as suggested in the top row of Figure 2.

We observed that the correlation matrix of style statistics (regardless of the mean or the standard deviation values) is not diagonal. Indeed, there exists a strong correlation among the channels. We applied eigenvalue decomposition over the calculated correlation matrix and find that very few eigenvectors dominate most variance, as shown in the bottom row of the Figure 2. This inspired us to rethink the augmentation of style statistics. The correlation matrix indicates that the combinations of feature statistics are not arbitrary but limited by task objectives and training procedures. Most variances happen within the specific principal directions. Arbitrary augmentation over the style statistics might damage the training itself. Previous research about why InstanceNorm can not outperform the BatchNorm in the discriminative tasks also proves this finding [20].

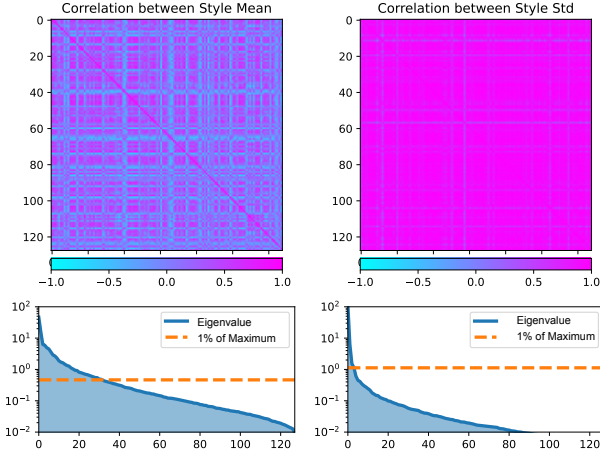


Figure 2. Visualization of feature statistics correlation. We calculate the style statistics (mean and standard deviation, respectively) on the PACS dataset. We extract the feature using the second residual block output from the ImageNet pretrained on ResNet18 [6] with a channel size of 128. For 4 domains of the PACS dataset, including Art, Cartoon, Photo, and Sketch, we select 64 cases from every category (7 categories in total) under each domain. Therefore, the data samples to calculate the correlation matrix is  $7 \times 4 \times 64 = 1792 \gg 128$ .

In reevaluating feature statistics augmentation, we observe that both MixStyle [50] and pAdaIN [22] limit their augmentations within the original feature space, preserving either channel information or channel combinations respectively. DSU [18] attempts to exceed this limitation via uncertainty quantification, but hinges on strict orthogonality between channel feature statistics. Our proposed model, CSU, breaks these boundaries by generating feature statistics beyond the training domain while preserving channel correlations, offering a more robust feature augmentation. CSU builds upon the strengths of these strategies and addresses their limitations, significantly advancing the field of feature statistics augmentation. There are some other style augmentation methods like AdverStyle [42] or Style Neophile [11]. However, direct mathematical correlation relationship analysis for these methods is not feasible due to complex adversarial training or distribution selection.

### 3.2. Modeling correlated style uncertainty

Given that the correlation matrix is real, symmetric, and positive semi-definite, then we can apply eigenvalue decomposition on  $\Sigma_\mu, \Sigma_\sigma$  to analyze its subspaces as:

$$\Sigma_\mu = Q_\mu \text{diag}(\Lambda_\mu) Q_\mu^T, \quad (5)$$

$$\Sigma_\sigma = Q_\sigma \text{diag}(\Lambda_\sigma) Q_\sigma^T, \quad (6)$$

$$Q_\mu Q_\mu^T = Q_\sigma Q_\sigma^T = I, \quad (7)$$

$$\Lambda_\mu, \Lambda_\sigma \in \mathbb{R}^C, Q_\mu, Q_\sigma \in \mathbb{R}^{C \times C}, \quad (8)$$

where  $\Lambda_{\mu,i} \geq \Lambda_{\mu,j} \geq 0$ ,  $\Lambda_{\sigma,i} \geq \Lambda_{\sigma,j} \geq 0$  ( $i > j$ ) represent the sorted eigenvalues, and  $Q_\mu, Q_\sigma$  are the corresponding eigenvectors. The eigenvector corresponding to the largest eigenvalue represents the direction that we could apply dense augmentation. Eigenvectors corresponding to the eigenvalues of 0 or close to 0 are not considered in the data augmentation process due to low variance across such directions within the dataset.

Assuming that the  $\mu, \sigma$  still follows the multi-variable Gaussian distribution, and  $k_\mu, k_\sigma$  represent the independent variable number (or the rank of the corresponding covariance matrix), then we can represent the probability distribution function as:

$$f_\mu = \frac{1}{(2\pi)^{k_\mu} \det^*(\Sigma_\mu)} \exp^{-(\mu - E(\mu))^T \Sigma_\mu^+ (\mu - E(\mu))}, \quad (9)$$

$$f_\sigma = \frac{1}{(2\pi)^{k_\sigma} \det^*(\Sigma_\sigma)} \exp^{-(\sigma - E(\sigma))^T \Sigma_\sigma^+ (\sigma - E(\sigma))}, \quad (10)$$

where the  $\det^*$  is the pseudo-determinant and  $\Sigma^+$  is the generalized inverse. Based on this distribution function, we further derive the correlated uncertainty augmentation after we sample  $\epsilon_\mu, \epsilon_\sigma \in \mathbb{R}^{N \times C}$  from the standard Gaussian distribution  $Y \sim \mathcal{N}(0, I)$  as follow:

$$P_\mu = Q_\mu \text{diag}(\Lambda_\mu)^{\frac{1}{2}} Q_\mu^T, \quad (11)$$

$$P_\sigma = Q_\sigma \text{diag}(\Lambda_\sigma)^{\frac{1}{2}} Q_\sigma^T, \quad (12)$$

$$\hat{\epsilon}_\mu = \epsilon_\mu P_\mu, \quad \hat{\epsilon}_\sigma = \epsilon_\sigma P_\sigma. \quad (13)$$

Essentially, we determine the transform matrix  $P$  such that the covariance matrix  $\Sigma = PP^T$  as shown in Eq.11-12, and these transformation matrices allow us to maintain the correlation. We sample the correlated perturbations  $\hat{\epsilon} \in \mathbb{R}^{N \times C}$  from independent Gaussian noise  $\epsilon \in \mathbb{R}^{N \times C}$  as shown in Eq.13. Note the covariance matrix:

$$\begin{aligned} \Sigma = PP^T &= Q \text{diag}(\Lambda_\mu)^{\frac{1}{2}} Q^T (Q \text{diag}(\Lambda_\mu)^{\frac{1}{2}} Q^T)^T \\ &= Q \text{diag}(\Lambda_\mu) Q^T. \end{aligned}$$

These two transformations  $Q \text{diag}(\Lambda_\mu)^{\frac{1}{2}} Q^T$ ,  $Q \text{diag}(\Lambda_\sigma)^{\frac{1}{2}} Q^T$  both work in principle. However, in practice, if we use the standard format of eigen-decomposition, stochastic axis swapping in the eigenvector will not influence this decomposition, but it can lead to unstable training. Thus, we apply the decomposition trick in the format of  $Q \text{diag}(\Lambda_\mu)^{\frac{1}{2}} Q^T$  rather than using only one set of eigenvectors  $Q \text{diag}(\Lambda_\mu)^{\frac{1}{2}}$  to avoid the random flip issue in traditional eigenvalue decomposition [8]. While we derive  $\epsilon_\mu, \epsilon_\sigma$  from  $C$  independent normal variables, the resulting  $\hat{\epsilon}_\mu, \hat{\epsilon}_\sigma$  contain only  $k_\mu, k_\sigma$  independent components, corresponding to the non-zero components of eigenvalues. It is also worth noticing that the computation of eigenvalue decomposition is conducted by the torch operator with a complexity of  $\sim O(N^3)$



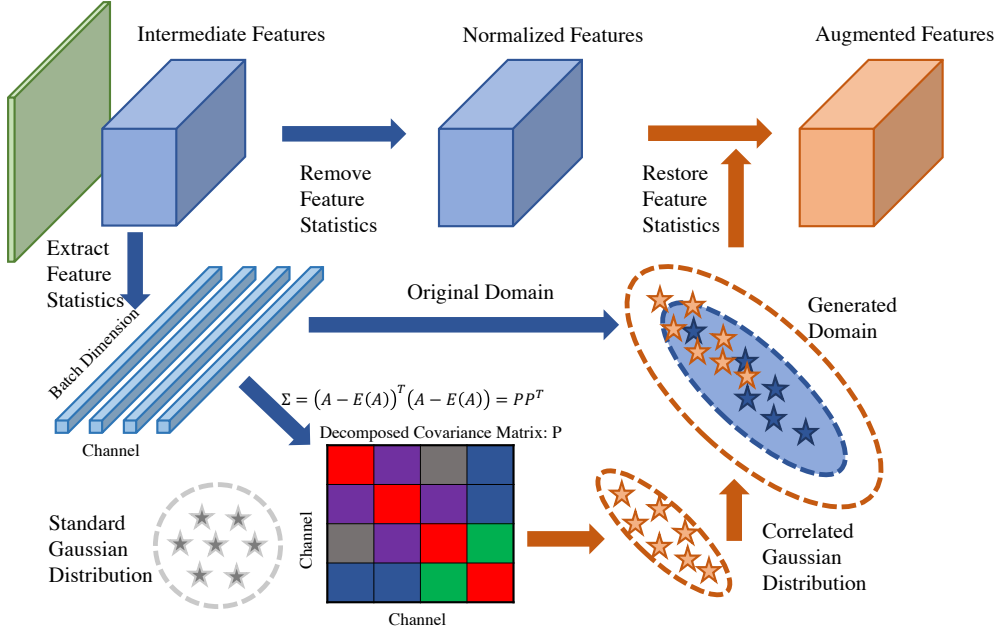


Figure 3. Visualization of feature statistics augmentation using correlated style uncertainty (CSU). Given the intermediate features extracted from the network, we first estimate the covariance matrix of feature statistics and decompose the covariance matrix as described in Sec 3.2. Based on this decomposition, we could generate correlated augmentation from the standard Gaussian distribution that shares identical distribution as the original domain. Then, we update raw data feature statistics by adding this correlated augmentation. Finally, we restore the feature statistics back to the normalized features and achieve the augmented features.

in practice [30]. Given the computation is conducted in polynomial time complexity and  $N < 512$  for most cases, the extra time required is negligible.

### 3.3. Style augmentation with CSU

Based on the previous two sections, we now present the style augmentation with correlated style uncertainty as follows:

$$\beta(x) = \mu(x) + \lambda * \hat{\epsilon}_\mu, \quad (14)$$

$$\gamma(x) = \sigma(x) + \lambda * \hat{\epsilon}_\sigma, \quad (15)$$

where  $\lambda \sim Beta(\alpha, \alpha)$  represents the augmentation intensity generated from the Beta distribution. Hyperparameter  $\alpha$  controls the shape of the distribution. In the ablation experiments, we further show the influence of hyperparameter selections on the final performance. We can understand the equation in one more intuitive way, the first part is to provide in-domain samples to cover the whole training domain, and the second is to provide the extrapolation while maintaining the same data distribution.

$$\beta(x) = \underbrace{\mu(x)}_{\text{In Domain Sample}} + \underbrace{\lambda * \hat{\epsilon}_\mu}_{\text{Out Domain extrapolation}}.$$

The final augmented instance feature can be defined as:

$$CSU(x) = \gamma(x) \left( \frac{x - \mu(x)}{\sigma(x)} \right) + \beta(x). \quad (16)$$

This plug-and-play module can be easily inserted into any current framework. We provide pseudo-code (PyTorch) in the supplementary materials.

## 4. Experiments

### 4.1. Multi-domain Classification Tasks

We validate our model’s performance on various multidomain classification tasks, including PACS, Office-Home, and Camelyon17. Figure 4 shows some examples with observable domain shifts within the same class. In all experiments, the domain tags are agnostic. Following the MixStyle, we adopt the ResNet-18 [7] with ImageNet [5] pre-training as the backbone for classification. We follow the Leave-One-Domain-Out strategy, which leaves one domain out for evaluation and the rest of the domains participating in the training. Adhering to a fair evaluation framework, we implement the multi-domain classification DASSL setup widely embraced in the works of Zhou et al. [48] for comparison. The batch size is set as 64. We conduct all the experiments on 2 NVIDIA A6000 GPU based on PyTorch [23] framework.

#### 4.1.1 PACS classification

PACS [14] is a widely used benchmark dataset for DG, which contains four domains: Photo (1,670 images), Art Painting (2,048 images), Cartoons (2,344 images), and



Figure 4. Some examples from multi-domain classification including (a) PACS, (b) Office-Home, and (c) Camelyon17 dataset.

Method	Reference	Art	Cartoon	Photo	Sketch	Average(%)
Baseline	-	74.3	76.7	<b>96.4</b>	68.7	79.0
Mixup [39]	ICLR 2018	76.8	74.9	95.8	66.6	78.5
Manifold Mixup [32]	ICML 2019	75.6	70.1	93.5	65.4	76.2
CutMix [38]	ICCV 2019	74.6	71.8	95.6	65.3	76.8
JiGen [2]	CVPR 2019	79.4	75.3	96.0	71.6	80.5
RSC [10]	ECCV 2020	78.9	76.9	94.1	76.8	81.7
L2A-OT [49]	ECCV 2020	83.3	78.2	96.2	76.3	82.8
SagNet [21]	CVPR 2021	83.6	77.7	95.5	76.3	83.3
pAdaIN [22]	CVPR 2021	81.7	76.6	96.3	75.1	82.5
SFA-A [17]	ICCV2021	81.2	77.8	93.9	73.7	81.7
MixStyle [50]	ICLR 2021	82.3	79.0	96.3	73.8	82.8
DSU [18]	ICLR 2022	83.6	79.6	95.8	77.6	84.1
CSU (Ours)	-	<b>85.0</b>	<b>81.0</b>	96.3	<b>78.4</b>	<b>85.2</b>

Table 1. Experimental results on the PACS multi-domain classification task. CSU achieves around highly significant improvements over the baseline in Art, Cartoon, and Sketch domains, respectively, as: 14.3%, 5.6%, and 12.2%. Besides, CSU also shows superiority over other methods, which demonstrates its effectiveness.  $\alpha$  is set as 0.1.

Sketches (3,929 images). Each domain consists of seven categories for classification tasks. These domain shifts are highly suitable for validating the effectiveness of the DG algorithms. Here, we compare our model’s performance with other SOTA methods, and all the evaluation metrics indicate the reported value by default.

As detailed in Table 1, our experimental outcomes reveal a significant improvement over other methodologies, underscoring the robustness of the CSU. We observe an increase of 14.3%, 5.6%, and 12.2% over the baseline in Art, Cartoon, and Sketch domains, respectively. Overall, CSU has a nearly 7.8% improvement in average accuracy across four domains. It should be also noted that we adopted the pre-trained model on ImageNet; therefore, it would be hard to generate significant improvement over the baseline in the Photo domain (As discussed in [36]), as this is expected. Nevertheless, we can still preserve the most dominant features by taking advantage of correlation modeling. Consequently, we achieve a performance drop of only around 0.1% compared with other methods. To guarantee the reliability of the reported value, we conduct training stability analysis in the supplementary materials.

We also notice that there are some domain generalization works are developed based on the DomainBed framework. While direct results comparison is not feasible, we still provide some results discussion here. Among these works,

like [35], the author introduced amplitude mix which linearly interpolates between the amplitude spectrum of two images to achieve data augmentation. The author achieved one classification accuracy of 84.51% with ResNet18 which is lower than our performance, while the baseline method achieved one accuracy of 79.9% which is higher than our performance. [11] achieve one classification accuracy of 84.5% with ResNet18 by adaptively synthesizing diverse style information with a greedy algorithm while noting that the baseline model performance is also 79.9% [43] proposes to perform exact feature distribution matching in the image feature space and achieve one classification performance of 83.9 with ResNet18 while the baseline method achieves one performance of 79.5% [3] shows that the effect of the pre-trained model with mutual-information regularization can improve the out-of-distribution improvement, which is orthogonal to all previous methods. [37] propose a novel proxy-based contrastive learning method by replacing the sample-to-sample relations with proxy-to-sample relations. This method achieves one classification accuracy of 88.7% with ResNet50, with ERM showing one performance of 85.5%.

#### 4.1.2 Office-Home classification

Office-Home [31] is another benchmark dataset for DG, containing four domains: Art, Clipart, Product, and Real-World, and each domain consists of 65 categories. The dataset contains 15,500 images with an average of around 70 photos per class. Similarly, we compare our model’s performance with other SOTA methods.

Method	Art	Clipart	Product	Real	Average(%)
Baseline	58.8	48.3	74.2	76.2	64.4
Mixup [39]	58.2	49.3	74.7	76.1	64.6
CrossGrad [28]	58.4	49.4	73.9	75.8	64.4
Manifold Mixup [32]	56.2	46.3	73.6	75.2	62.8
CutMix [38]	57.9	48.3	74.5	75.6	64.1
RSC [10]	58.4	47.9	71.6	74.5	63.1
L2A-OT [49]	60.6	50.1	74.8	<b>77.0</b>	65.6
MixStyle [50]	58.7	53.4	74.2	75.9	65.5
DSU [18]	60.2	54.8	74.1	75.1	66.1
CSU (Ours)	<b>61.3</b>	<b>54.9</b>	<b>74.9</b>	76.1	<b>66.8</b>

Table 2. Experimental results on Office-Home multi-domain classification task. We achieve around 4.3%, 13.6%, 0.9% improvement over the baseline in the Art, Clipart, and Product domain, respectively. CSU consistently outperforms the other strong baseline models with considerable margins ( $\alpha = 0.4$ )

As shown in Table 2, CSU achieves around 4.3%, 13.6%, 0.9% improvement over the baseline in Art, Clipart, and Product domain, respectively. On average, CSU shows 3.7% improvement over the baseline across four domains. On the other hand, improving the Real-world images in the PACS dataset is hard due to the same reason given before. Despite this difficulty, CSU remains with a strong performance with only 0.1% drop.

Method	H1	H2	H3	H4	H5	Average(%)
Baseline	95.3	91.4	89.5	96.2	94.6	93.4
MixStyle [50]	96.1	91.2	93.0	95.0	92.7	93.6
pAdaIN [22]	96.6	93.0	<b>94.7</b>	95.2	94.0	94.7
DSU [18]	<b>96.8</b>	93.3	91.7	<b>96.4</b>	94.4	94.5
CSU (Ours)	96.7	<b>93.8</b>	94.2	95.5	<b>95.5</b>	<b>95.1</b>

Table 3. Experimental results on Camelyon17 multi-domain classification task. H1-H5 represents five different hospitals. We can find that CSU outperforms other methods ( $\alpha = 0.3$ ).

Model	Market To Duke				Duke To Market			
	mAP	R1	R5	R10	mAP	R1	R5	R10
ResNet-50	19.3	35.4	50.4	56.4	20.4	45.2	63.6	70.9
Baseline	19.3	35.4	50.4	56.4	20.4	45.2	63.6	70.9
RandomEraser [47]	14.3	27.8	42.6	49.1	16.1	38.5	56.8	64.5
DropBlock [50]	18.2	33.2	49.1	56.3	19.7	45.3	62.1	69.1
MixStyle [50]	23.8	42.2	58.8	64.8	24.1	51.5	69.4	76.2
pAdaIN [22]	22.0	41.4	56.4	62	24.1	52.1	68.8	75.5
DSU [18]	21.2	40.5	56	62.5	24.0	51.7	70.6	77.3
CSU (Ours)	<b>24.5</b>	<b>44.1</b>	<b>60.3</b>	<b>65.9</b>	<b>24.4</b>	<b>52.4</b>	<b>71.4</b>	<b>78.2</b>
OSNet	25.9	44.7	59.6	65.4	24.0	52.2	67.5	74.7
Baseline	25.9	44.7	59.6	65.4	24.0	52.2	67.5	74.7
RandomEraser [47]	20.5	36.2	52.3	59.3	22.4	49.1	66.1	73.0
DropBlock [50]	23.1	41.5	56.5	62.5	21.7	48.2	65.4	71.3
MixStyle [50]	27.2	48.2	62.7	68.4	27.8	58.1	74.0	81.0
pAdaIN [22]	28.3	48.8	62.7	68.1	27.6	57.5	74.2	80.3
DSU [18]	29.0	51.0	65.0	70.4	26.1	57.2	74.6	80.7
CSU (Ours)	<b>31.1</b>	<b>53.1</b>	<b>67.9</b>	<b>76.3</b>	<b>29.8</b>	<b>60.1</b>	<b>77.3</b>	<b>83.4</b>

Table 4. Experimental results on the Duke-Market1501 Instance Retrieval Datasets. CSU achieves around 26.9%, 19.6% advancement over the baseline in mAP value using the ResNet-50 model in the Market1501 to Duke and the Duke to Market1501 experiment, correspondingly. Likewise, CSU achieves around 20.1%, 24.2% improvement for the OSNet model experiment. We could also observe similar advancements in ranking accuracy, and CSU achieves impressive improvement over other methods.

### 4.1.3 Camelyon17 classification

Medical image analysis often suffers the most from domain shifting, given that multiple parameters, like the image acquisition device, and protocol can induce significant domain shift. We validate the model’s performance on the challenging Camelyon17 dataset [1], containing images from five medical centers. This dataset consists of the histopathological images as input and the labels indicating whether the central region includes any tumor tissue. Due to lacking reported performance from the current literature, we conduct this experiment from scratch based on the WILDS framework proposed by Koh [13]. Besides the baseline, we compare our model with three state-of-art strategies, including the MixStyle [50], pAdaIN [22], DSU [18]. For a fair comparison, we directly use the official implementation of each method without any modifications.

Table 3 proves the effectiveness of our model. CSU achieves impressive improvement compared with the baseline or other style augmentation methods. This indicates that by taking advantage of correlation modeling, CSU can help induce a more generalized model even with extremely challenging medical data.

## 4.2. Instance Retrieval Experiments

The undertaking of person re-identification, which involves the cross-camera recognition of individuals, poses a substantial domain generalization challenge. Considering each distinct camera output as its own unique domain amplifies the intricacy inherent to the person re-identification task. Following previous research, we conducted this experiment on the commonly used Duke [25] and Market1501 [45] datasets. To evaluate the model’s generalizability, we take one dataset as training and test the performance on the other domain. The camera data from the test domain does not participate in any training process. We adopt the exact framework implementation of MixStyle and test the CSU influence on the ResNet50 [7] and OSNet [49]. Similarly, ranking accuracy and mean average precision (mAP) are performance measures. For a fair comparison, we repeat the pAdaIN and DSU experiments on the same framework with the MixStyle and use the best configuration reported in the original paper.

Table 4 shows the experiment results using two models in the two domains. We could observe that CSU outperforms other methods by a large margin. CSU achieves around 26.9%, 19.6% advancement in mAP using the ResNet-50 model in the Market1501 to Duke and the Duke to Market1501 experiment, correspondingly. Similarly, CSU achieves around 20.1%, 24.2% improvement for the OSNet model experiment. We could also observe similar advancements in ranking accuracy, and CSU achieves impressive improvement over other methods. Nevertheless, to show the effectiveness of CSU rather than position fine-tuning, we insert the permutation in all positions as described in Sec 4.3. The supplementary materials show that changing the inserting position can achieve even more significant performance advancement.

## 4.3. Ablation Experiments

**Insert Position Selection:** To answer the question of where we should insert the CSU, we conduct comprehensive experiments on the PACS and Office-Home datasets using the ResNet18 structure. We investigate all possible positions of ResNet18, including the first Convolution, first Max Pooling, and 1, 2, 3, and 4 Res-block, which are named 0-5, respectively. We divide the experiment into several groups according to the inserted CSU number. Within each group, we shift the start position one by one from 0 to end. For example, for the group containing 2 CSU blocks, we will have 01, 12, 23, 34, and 45 potential combinations and five comparison experiments in total. To avoid the influence of hyperparameters, we set  $\alpha = 0.3$  for all experiments. Thus, we can reasonably and adequately compare the inserting position’s influence on final performance.

Figure 5 shows the ablation experiment results. Inserting 6 blocks of CSU in all potential positions achieves the best results for the PACS dataset while inserting 4 blocks of CSU




Figure 5. Influence of inserting position. Inserting 6 blocks of CSU in all potential positions achieves the best results for the PACS dataset while for the Office-Home dataset, inserting 4 blocks of CSU in the last 4 positions achieves the best results. The performance after inserting the CSU model always shows superiority over the baseline by a large margin regardless of the inserting number or position.

in the last 4 positions achieves the best results for the Office-Home dataset. Within each group of a fixed number of CSU blocks, the performance tends to increase when we start the inserting position at the medium blocks. This trend is different with the MixStyle which the model prefers the first several blocks [50]. We explain this phenomenon as CSU can provide more reasonable feature (statistics) augmentation due to correlation preservation. This preservation will avoid information loss in the medium or last blocks. We can also notice that compared with inserting 4, 5, and 6 blocks of CSU, inserting 2, or 3 blocks of CSU can not achieve comparable performance. This indicates that a more significant number of CSU blocks can be helpful to increase the model’s generalization ability due to accumulated correlations over the blocks. It is also worth noting that no matter how we choose the inserting position, the performance of the CSU model always shows superiority over the baseline by a large margin. This firmly proves the effectiveness of the proposed model.

**Hyper-parameter Selection:** As described in the previous section, the hyper-parameter alpha determines the intensity of augmentation during training by manipulating the shape of the Beta distribution. Here, we show the influence of alpha on the PACS, Office-Home using the ResNet18 structure. Similarly, to avoid the influence of different inserting positions, we insert CSU block in all positions for every experiment. We select  $\alpha$  from 0.1, 0.2, 0.3, 0.4, 0.5, 0.7, and 0.9 for one comprehensive experiment. As shown in Figure 6, we can find that a smaller number of  $\alpha < 0.5$  always performs better than the relatively larger number ( $> 0.5$ ). Based on these experiment results, we recommend selecting the alpha from 0.1, 0.2, 0.3, and 0.4, and the best configuration may vary according to the tasks.

**Effect of Batch Size:** The batch size could have a notable impact on the accuracy of correlation data, making it critical to scrutinize its effect on the final model’s generalizability. Accordingly, we’ve introduced the CSU at every stage, as per the earlier section, and set a fixed value of  $\alpha = 0.3$  in all the experiments. We compare the model’s performance with a batch size of 16, 32, 64, 128, 256, and 512. Figure 7 shows the experimental result. We found that




Figure 6. Influence of hyper-parameters selection. A smaller number of  $\alpha < 0.5$  performs better. Red indicates the best result.




Figure 7. Effect of batch size on the two classification tasks. Too small or large batch sizes can be toxic for DG. Red indicates the best result.

it might be hard to estimate an accurate correlation when the batch size is too small. Similarly, when the batch size is too large, the network tends to converge to sharp minimizers of the training and testing functions, leading to poorer generalizations, as shown in previous research [12].

## 5. Conclusion

In this paper, we introduce the Correlated Style Uncertainty (CSU), a novel domain generalization approach that transcends linear interpolation while upholding the correlation among feature channels. With detailed and exhaustive ablation studies, we’ve determined the optimal position for integrating the CSU model, assessed the impact of sampling hyperparameters, and evaluated the ideal batch size for yielding a more universally applicable model. Comprehensive experiments across numerous datasets corroborate that CSU model notably enhances the model’s ability to generalize. This research is expected to catalyze further in-depth studies on feature statistics augmentation in future.

This work is supported by NIH R01-CA246704, R01-CA240639, R15-EB030356, R03-EB032943, U01-DK127384-02S1, and U01-CA268808.



## References

- [1] Peter Bandi, Oscar Geessink, Quirine Manson, Marcorcy Van Dijk, Maschenka Balkenhol, Meyke Hermesen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 38(2):550–560, 2018. [2](#), [7](#)
- [2] 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 (CVPR)*, pages 2229–2238, 2019. [6](#)
- [3] Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*. Springer, 2022. [6](#)
- [4] Seokeon Choi, Taekyung Kim, Minki Jeong, Hyoungseob Park, and Changick Kim. Meta batch-instance normalization for generalizable person re-identification. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3425–3435, 2021. [3](#)
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255, 2009. [5](#)
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [4](#)
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 770–778, 2016. [5](#), [7](#)
- [8] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 791–800, 2018. [3](#), [4](#)
- [9] 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 (CVPR)*, pages 1501–1510, 2017. [3](#)
- [10] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 124–140, 2020. [6](#)
- [11] 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*, 2022. [2](#), [3](#), [4](#), [6](#)
- [12] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. [8](#)
- [13] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 5637–5664, 2021. [7](#)
- [14] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 5542–5550, 2017. [2](#), [5](#)
- [15] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5400–5409, 2018. [3](#)
- [16] Haoliang Li, YuFei Wang, Renjie Wan, Shiqi Wang, Tie-Qiang Li, and Alex Kot. Domain generalization for medical imaging classification with linear-dependency regularization. *Advances in Neural Information Processing Systems*, 33:3118–3129, 2020. [2](#)
- [17] Pan Li, Da Li, Wei Li, Shaogang Gong, Yanwei Fu, and Timothy M Hospedales. A simple feature augmentation for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8886–8895, 2021. [3](#), [6](#)
- [18] 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*, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [19] Wang Lu, Jindong Wang, Haoliang Li, Yiqiang Chen, and Xing Xie. Domain-invariant feature exploration for domain generalization. *arXiv preprint arXiv:2207.12020*, 2022. [2](#)
- [20] Hyeonseob Nam and Hyo-Eun Kim. Batch-instance normalization for adaptively style-invariant neural networks. *Advances in Neural Information Processing Systems (NIPS)*, 31, 2018. [3](#)
- [21] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8690–8699, 2021. [6](#)
- [22] Oren Nuriel, Sagie Benaim, and Lior Wolf. Permuted adain: Reducing the bias towards global statistics in image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9482–9491, 2021. [2](#), [3](#), [4](#), [6](#), [7](#)
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in NIPS*, 32, 2019. [5](#)
- [24] Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. Sentry: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 8558–8567, 2021. [1](#)
- [25] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for

- multi-target, multi-camera tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 17–35, 2016. 2, 7
- [26] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 8050–8058, 2019. 1
- [27] Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 68–83, 2020. 3
- [28] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018. 6
- [29] Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9624–9633, 2021. 3
- [30] Gilbert Strang. *Linear algebra and its applications*. Belmont, CA: Thomson, Brooks/Cole, 2006. 5
- [31] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5018–5027, 2017. 2, 6
- [32] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *Proceedings of the ICML*, pages 6438–6447, 2019. 6
- [33] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *Advances in Neural Information Processing Systems (NIPS)*, 31, 2018. 2
- [34] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE Transactions on Knowledge and Data Engineering*, 2022. 2
- [35] 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. 6
- [36] Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer. Robust and generalizable visual representation learning via random convolutions. *arXiv preprint arXiv:2007.13003*, 2020. 2, 6
- [37] Xufeng Yao, Yang Bai, Xinyun Zhang, Yuechen Zhang, Qi Sun, Ran Chen, Ruiyu Li, and Bei Yu. Pcl: Proxy-based contrastive learning for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 6
- [38] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, pages 6023–6032, 2019. 6
- [39] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 6
- [40] Ling Zhang, Xiaosong Wang, Dong Yang, Thomas Sanford, Stephanie Harmon, Baris Turkbey, Bradford J Wood, Holger Roth, Andriy Myronenko, Daguang Xu, et al. Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation. *IEEE Transactions on Medical Imaging*, 39(7):2531–2540, 2020. 2
- [41] Shengdong Zhang, Ehsan Nezhadarya, Homa Fashandi, Jiayi Liu, Darin Graham, and Mohak Shah. Stochastic whitening batch normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10978–10987, 2021. 3
- [42] Yabin Zhang, Bin Deng, Ruihuang Li, Kui Jia, and Lei Zhang. Adversarial style augmentation for domain generalization. *arXiv preprint arXiv:2301.12643*, 2023. 2, 3, 4
- [43] 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*, 2022. 6
- [44] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Advances in Neural Information Processing Systems (NIPS)*, 33:16096–16107, 2020. 3
- [45] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE International Conference on Computer Vision (ECCV)*, pages 1116–1124, 2015. 2, 7
- [46] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial style augmentation for domain generalized urban-scene segmentation. *arXiv preprint arXiv:2207.04892*, 2022. 2, 3
- [47] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13001–13008, 2020. 7
- [48] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1, 2, 5
- [49] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 561–578, 2020. 6, 7
- [50] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021. 1, 2, 3, 4, 6, 7, 8