

Decompose, Adjust, Compose: Effective Normalization by Playing with Frequency for Domain Generalization

Sangrok Lee^{*1}, Jongseong Bae^{*2}, and Ha Young Kim^{†1}

¹Graduate School of Information, Yonsei University

²Department of Artificial Intelligence, Yonsei University
 {lsrock1, js.bae, hayoung.kim}@yonsei.ac.kr

Abstract

Domain generalization (DG) is a principal task to evaluate the robustness of computer vision models. Many previous studies have used normalization for DG. In normalization, statistics and normalized features are regarded as style and content, respectively. However, it has a content variation problem when removing style because the boundary between content and style is unclear. This study addresses this problem from the frequency domain perspective, where amplitude and phase are considered as style and content, respectively. First, we verify the quantitative phase variation of normalization through the mathematical derivation of the Fourier transform formula. Then, based on this, we propose a novel normalization method, *PCNorm*, which eliminates style only as the preserving content through spectral decomposition. Furthermore, we propose advanced *PCNorm* variants, *CCNorm* and *SCNorm*, which adjust the degrees of variations in content and style, respectively. Thus, they can learn domain-agnostic representations for DG. With the normalization methods, we propose ResNet-variant models, *DAC-P* and *DAC-SC*, which are robust to the domain gap. The proposed models outperform other recent DG methods. The *DAC-SC* achieves an average state-of-the-art performance of 65.6% on five datasets: *PACS*, *VLCS*, *Office-Home*, *DomainNet*, and *TerraIncognita*.

1. Introduction

Deep learning has performed remarkably well in various computer vision tasks. However, the performance decreases when distribution-shifted test data are given [39]. As training and testing datasets are assumed to be identically and independently distributed, common vision models are not

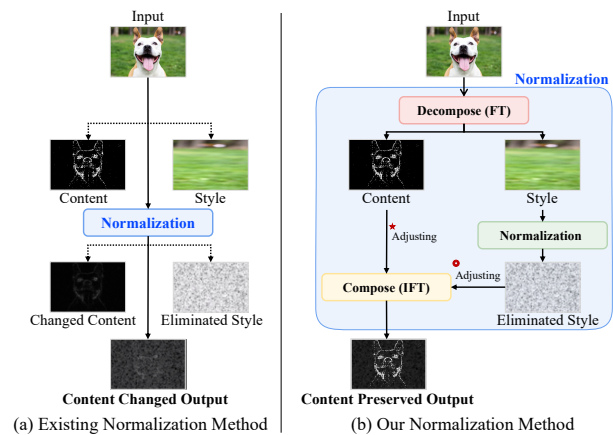


Figure 1. Concepts of (a) the existing normalization and (b) the proposed methods. Our methods prevent or adjust the content change caused by existing normalization using spectral decomposition. The solid line marks the feedforward process and the dashed line conceptually represents the content and style of the feature. Red-colored star and doughnut in (b) indicate the content and style adjusting terms, respectively.

as robust as the human vision system, which is not confused by affordable changes in the image style [14]. Domain generalization (DG) aims to learn models that are robust to the gap between source domain and unseen target domain to address this problem [49]. Moreover, DG is challenging because models should learn domain-irrelevant representation in an unsupervised manner.

The style-based approach is widely studied for DG, which defines the domain gap as the difference in style [19, 31, 51, 56, 57]. Typically, normalization methods, such as batch normalization (BN) [18], layer normalization (LN) [18], and instance normalization (IN) [45], which are well-known in style transfer, are used in this approach. Normalization statistics contain style information, and normalization can successfully extract the style from a feature. However, the content is also changed when the style is eliminated [15, 19, 34].

^{*}Equal contribution

[†]Corresponding author

Moreover, another method for style-based DG [7, 48, 50, 51] is the frequency domain-based method. Input images are decomposed into amplitude and phase using the Fourier transform (FT) [4]. The amplitude and phase are each regarded as the style and content of the input image, respectively [7, 33, 37, 51, 53]. Each component is manipulated independently to generate the style-transformed image. In this context, this method has an advantage of the separation between style and content [56]. Nevertheless, most previous studies have applied it to just the input-level data augmentation [7, 51, 53] for DG.

Thus, the normalization is expected to be complemented by the frequency domain-based method if the method is also applicable at the feature level. To identify the feasibility of this, we conduct a style transfer experiment. We replace IN in AdaIN [15], a milestone work that uses normalization in style transfer, with spectral decomposition. The qualitative results in Fig. 2 indicate that the frequency domain-based method can work as a feature-level style-content separator instead of normalization.

Motivated by this, we aim to overcome the content change problem in normalization by combining the normalization with spectral decomposition. The overall concept of our proposed method is visualized in Fig. 1. For this, we investigate the effect of the existing normalization in DG from the standpoint of the frequency domain. We verify how normalization transforms the content of a feature by mathematically deriving the FT formula. This is the first to present such an analysis.

Then, based upon the analysis, we introduce a novel normalization method, phase-consistent normalization (*PCNorm*), which preserves the content of a pre-normalized feature. The *PCNorm* synthesizes a content-invariant normalized feature by composing the phase of pre-normalized feature and the amplitude of post-normalized feature. The experimental results reveal the effectiveness of *PCNorm* in DG compared to existing normalization.

Along with the success of *PCNorm*, we take a step further and propose two advanced *PCNorm* variants: content-controlling normalization (*CCNorm*) and style-controlling normalization (*SCNorm*). The main idea of both methods is not to preserve the content or style but to adjust the change in it. *CCNorm* and *SCNorm* regulate the changes in content and style, respectively, so they can synthesize more robust representations of the domain gap.

With the proposed normalization methods, we propose ResNet [13] variant models, DAC-P and DAC-SC. DAC-P is the initial model with *PCNorm*, and DAC-SC is the primary model using *CCNorm* and *SCNorm*. In DAC-P, the existing BN in the downsample layer is replaced with *PCNorm*. In contrast, DAC-SC applies *CCNorm* instead of *PCNorm*, and *SCNorm* is inserted at the end of each stage. We evaluate DAC-P and DAC-SC on five DG bench-

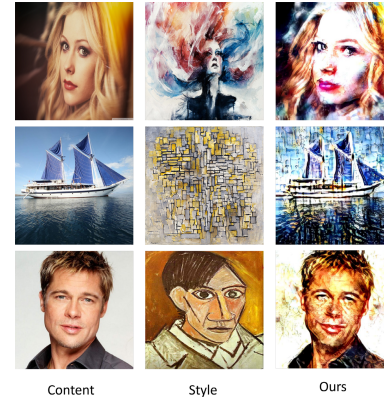


Figure 2. Examples of style transfer with spectral decomposition. Only the amplitude of the target images are transferred instead of their normalization statistics in AdaIN.

marks: PACS, VLCS, Office-Home, DomainNet and TerraIncognita, and DAC-P outperforms other recent DG methods with average performance of 65.1%. Furthermore, the primary model, DAC-SC, achieves state-of-the-art (SOTA) performance of 65.6% on average, and displays the highest performance at 87.5%, 70.3% and 44.9% on the PACS, Office-Home, and DomainNet benchmarks.

The contributions of this paper are as follows:

- For the first time, we analyze the quantitative shift in the phase caused by normalization using mathematical derivation.
- We introduce a new normalization, *PCNorm*, which can remove style only through spectral decomposition.
- We propose the advanced *PCNorm* variants, *CCNorm* and *SCNorm*, which can learn domain-agnostic features for DG by adjusting the degrees of the changes in content and style, respectively.
- We propose ResNet-variant models, DAC-P and DAC-SC, which applies our proposed normalization methods. We experimentally show that our methods are effective for DG and achieve SOTA average performances on five benchmark datasets.

2. Related Work

Domain-invariant learning for DG. The main purpose of DG methods is to learn domain-invariant features. There are numerous approaches for DG. Adversarial learning methods [9, 10, 25, 26, 26, 27, 55] prevent models from fitting domain-specific representations, allowing the extraction of domain-invariant representations only. Regularization methods introduce various regularization strategies, such as the reformulation of loss functions [22, 43], gradient-based methods [16, 40], attention-based method [28], and contrastive learning [20] for DG. Optimization methods aim to

reduce the distributional variance and domain-specific differences using kernel methods [3, 29] and penalty on variance [21]. Meta-learning methods [24, 54] simulate diverse domain changes for various splits of training and testing data from source datasets.

Style-based learning for DG. Style-based learning methods define the domain gap as the style difference between each domain and try to extract style-invariant features. There are two widely used methods that enable handling style and content of images, normalization and frequency domain-based methods. Normalization [6, 31, 42, 57] has been widely used to filter the style of features [15, 17, 30]. There have been various works for DG by applying normalization methods such as BN [42], IN [31, 57], and both of them [9, 35]. Jin *et al.* [19] points out the downside of normalization that it causes the loss of content.

In frequency-based methods, the style and content of the image are represented by the amplitude and phase in frequency domain [12, 32, 33, 38]. Through FT, many works manipulate them to make models robust to the distribution shift. FDA [53] proposes the spectral transfer, which is similar to the style transfer, that the low-frequency component of source amplitude is transferred to that of the target amplitude. In [51], the amplitude swap (AS) and amplitude mix (AM) strategies are introduced for data augmentation, the former is exactly the same as the spectral transfer [53] and the latter is to mix amplitude of source and target domain. Similarly, [50] implements data augmentation by applying multiplicative and additive Gaussian noises to both of the amplitude and phase of the source domain. Most frequency-based works have focused on input-level data augmentation. Different from those, we propose novel frequency domain-based normalization methods that are applicable at the feature level.

3. Analysis

Prior to analysis, we clarify the terms of normalization. We consider that normalization includes just BN, LN, and IN because these are all the existing normalization methods for DG, as far as we know. As mentioned in Sec. 1, normalization has a content change problem. However, there is still no explicit verification of the problem. Motivated by this, we examine it through a mathematical analysis from the frequency domain perspective. Specifically, we derive the quantitative change in the phase caused by normalization by expanding the FT formula.

3.1. Spectral Decomposition

We first explain the spectral decomposition [5, 37]. Generally, the term spectral decomposition has various meanings. In this work, it refers to decomposing a feature into amplitude and phase using the discrete FT (DFT) [44]. The DFT transforms the input spatial feature $f \in \mathbb{R}^{h \times w}$ into the

Algorithm 1 Pseudocode of the Proposed Normalization Methods in PyTorch-like Style.

```
# decompose: frequency feature to amplitude and phase
# compose: amplitude and phase to frequency feature
# weight: learnable parameters, T_c, T_s: temperatures
def pcnorm(f): # f: a spatial feature
    f_norm = batch_norm(f) # batch_norm: BN
    F = FT(f) # FT: fourier transform
    F_norm = FT(f_norm)
    a, p = decompose(F)
    a_norm, p_norm = decompose(F_norm)
    f = IFT(compose(a_norm, p)) # IFT: inverse FT
    return f

def ccnorm(f):
    weight_c = softmax(weight/T_c, dim=0)
    # mean: batch mean(train), cumulative mean(test)
    fc = f - mean * weight_c[0]
    f_norm = batch_norm(fc)
    FC = FT(fc)
    F_norm = FFT(f_norm)
    ac, pc = decompose(FC)
    a_norm, p_norm = decompose(F_norm)
    f = IFT(compose(a_norm, pc))
    return f

def scnorm(f):
    weight_s = softmax(weight/T_s, dim=0)
    f_norm = instance_norm(f) #instance_norm: IN
    F = FT(f)
    F_norm = FFT(f_norm)
    a, p = decompose(F)
    a_norm, p_norm = decompose(F_norm)
    f = IFT(compose(a_norm * weight_s[0] + a * weight_s
                    [1], p))
    return f
```

corresponding frequency feature $\mathcal{F} \in \mathbb{C}^{h \times w}$ where h and w are the height and the width of f . $f(x, y)$ is an element of f at image pixel (x, y) , and $\mathcal{F}(u, v)$ represents the Fourier coefficient at frequency component (u, v) . We omit the round bracket term (\cdot, \cdot) when we denote a feature composed of that element. The DFT of f is defined as follows:

$$\begin{aligned} \mathcal{F}(u, v) &= \frac{1}{wh} \sum_{x=0}^{w-1} \sum_{y=0}^{h-1} f(x, y) \exp^{i2\pi(\frac{ux}{w} + \frac{vy}{h})} \\ &= \mathcal{F}_{real}(u, v) + i \mathcal{F}_{img}(u, v), \end{aligned} \quad (1)$$

where i is imaginary unit. In addition, $\mathcal{F}_{real}(u, v)$ and $\mathcal{F}_{img}(u, v)$ are the real and imaginary parts of $\mathcal{F}(u, v)$, respectively, as follows:

$$\begin{aligned} \mathcal{F}_{real}(u, v) &= \frac{1}{wh} \sum_{x=0}^{w-1} \sum_{y=0}^{h-1} f(x, y) \cos 2\pi(\frac{ux}{w} + \frac{vy}{h}), \\ \mathcal{F}_{img}(u, v) &= \frac{1}{wh} \sum_{x=0}^{w-1} \sum_{y=0}^{h-1} f(x, y) \sin 2\pi(\frac{ux}{w} + \frac{vy}{h}). \end{aligned} \quad (2)$$

The process in which the spatial feature f is decomposed into the amplitude α and phase ρ is called spectral decomposition, where α and ρ are calculated as follows:

$$\begin{aligned} \alpha &= \sqrt{\mathcal{F}_{real}^2 + \mathcal{F}_{img}^2}, \\ \rho &= \arctan \frac{\mathcal{F}_{img}}{\mathcal{F}_{real}}. \end{aligned} \quad (3)$$

In addition, \mathcal{F} can be reassembled from α and ρ :

$$\mathcal{F} = \alpha \cos(\rho) + i \alpha \sin(\rho). \quad (4)$$

We define a function of the disassembling frequency feature in the amplitude and phase as $decompose(\cdot)$ and the opposite process as $compose(\cdot)$:

$$\begin{aligned} \alpha, \rho &= decompose(\mathcal{F}), \\ \mathcal{F} &= compose(\alpha, \rho). \end{aligned} \quad (5)$$

3.2. Content Variation by Normalization

In this section, we mathematically verify how normalization changes the content of the feature in the frequency domain. We develop the FT formula of normalized feature $f^{norm} \in \mathbb{R}^{h \times w}$ and identify the relationship with that of the original feature f . Moreover, f^{norm} is as follows:

$$f^{norm} = \frac{f - \mu}{\sigma}, \quad (6)$$

where μ and σ denote the statistical mean and standard deviation of f , respectively. Statistics are calculated differently as with normalization methods. In this study, however, μ and σ are set as constants because they are computed in the same way in a channel.

Like Eq. 1, the DFT of f^{norm} is as follows:

$$\mathcal{F}^{norm}(u, v) = \mathcal{F}_{real}^{norm}(u, v) + i \mathcal{F}_{img}^{norm}(u, v). \quad (7)$$

In Eq. 2, by the linearity property of FT, $\mathcal{F}_{real}^{norm}(u, v)$ and $\mathcal{F}_{img}^{norm}(u, v)$ are represented as follows:

$$\begin{aligned} \mathcal{F}_{real}^{norm}(u, v) &= \frac{1}{wh} \sum_{x=0}^{w-1} \sum_{y=0}^{h-1} \left\{ \frac{f(x, y) - \mu}{\sigma} \right\} \cos 2\pi \left(\frac{ux}{w} + \frac{vy}{h} \right), \\ \mathcal{F}_{img}^{norm}(u, v) &= \frac{1}{wh} \sum_{x=0}^{w-1} \sum_{y=0}^{h-1} \left\{ \frac{f(x, y) - \mu}{\sigma} \right\} \sin 2\pi \left(\frac{ux}{w} + \frac{vy}{h} \right). \end{aligned} \quad (8)$$

Then we derive a relationship between \mathcal{F} and \mathcal{F}^{norm} by presenting Eq. 8 in terms of Eq. 2:

$$\begin{aligned} \mathcal{F}_{real}^{norm} &= \frac{\mathcal{F}_{real} - \mathcal{F}_{real}^\mu}{\sigma}, \\ \mathcal{F}_{img}^{norm} &= \frac{\mathcal{F}_{img} - \mathcal{F}_{img}^\mu}{\sigma}, \end{aligned} \quad (9)$$

where \mathcal{F}_{real}^μ and \mathcal{F}_{img}^μ are real and imaginary parts of $\mathcal{F}^\mu \in \mathbb{C}^{h \times w}$. \mathcal{F}^μ is the frequency feature of $f^\mu \in \mathbb{R}^{h \times w}$, which is a feature whose elements are all μ .

In the same way as Eq. 3, the amplitude and phase of f^{norm} , α^{norm} and ρ^{norm} , respectively, can be computed as follows:

$$\begin{aligned} \alpha^{norm} &= \frac{\sqrt{(\mathcal{F}_{real} - \mathcal{F}_{real}^\mu)^2 + (\mathcal{F}_{img} - \mathcal{F}_{img}^\mu)^2}}{\sigma}, \\ \rho^{norm} &= \arctan \frac{\mathcal{F}_{img} - \mathcal{F}_{img}^\mu}{\mathcal{F}_{real} - \mathcal{F}_{real}^\mu}. \end{aligned} \quad (10)$$

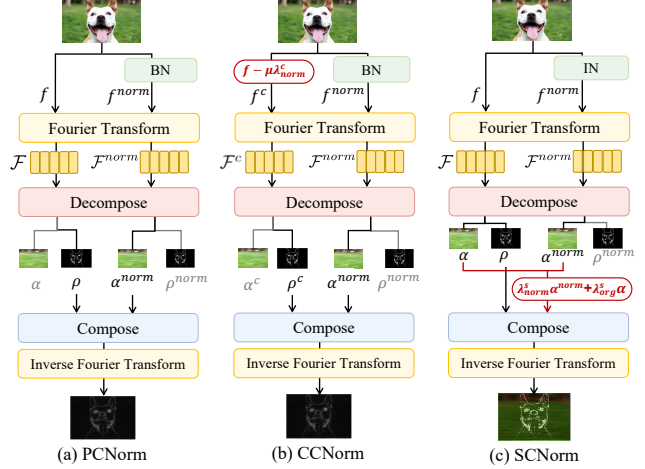


Figure 3. Illustrations of the proposed normalization methods. Included notations are the same as in Sec. 3 and 4. Red denotes adjusting operation in Sec. 4.2 and 4.3.

By comparing ρ^{norm} with ρ , we verify the numerical variation of the content information caused by normalization. We determine that ρ^{norm} , the phase of spatial feature $(f - \mu)/\sigma$, is same as the phase of $(f - \mu)$ because it is not affected by σ (Eq. 10). That is, the difference between ρ^{norm} and ρ in the frequency domain is simply caused by the mean shift in normalization in the spatial domain. Hence, we consider μ to be the content variation factor. As observed, the degree of content variation becomes greater when μ is larger.

4. Proposed Method

In this section, based upon the analysis in Sec. 3.2, we explain the novel normalization methods, *PCNorm*, *CCNorm*, and *SCNorm*, whose Pytorch-like pseudocode is described in Algorithm. 1. Next, we describe the ResNet [13]-variant models: DAC-P and DAC-SC.

4.1. Phase Consistent Normalization (PCNorm)

In Sec. 3.2, we verify that the difference between ρ and ρ^{norm} is caused by the mean shift by μ in normalization. Thus, we infer that a simple method to prevent content from changing is to avoid a mean shift in normalization. We can circumvent content variation using the phase of pre-normalized feature, ρ , instead of ρ^{norm} . In this context, we propose our *PCNorm*, depicted in Fig. 3 (a). *PCNorm* is a normalization method that maintains an original content. *PCNorm* decomposes pre- and post-normalized features into their phase and amplitude, respectively. Then, it combines ρ with α^{norm} . Finally, the composed frequency feature is transformed to spatial feature through inverse FT

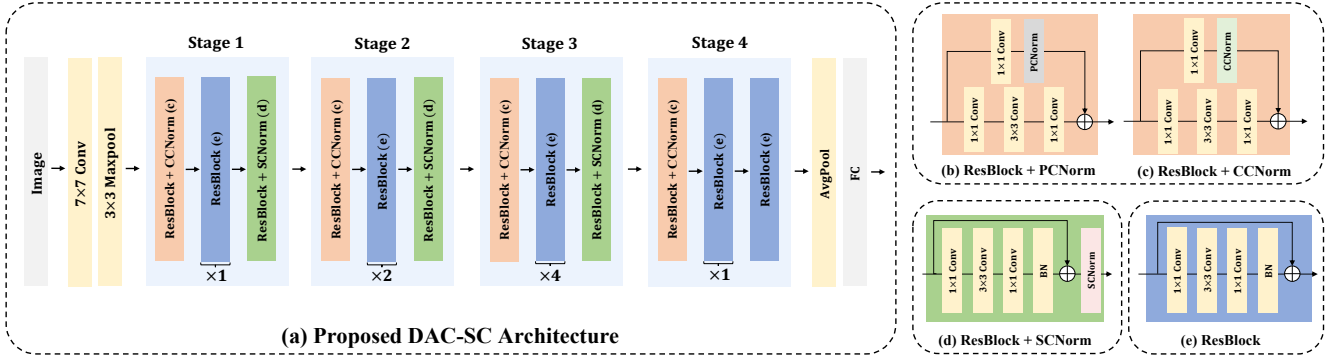


Figure 4. Overall architecture of DAC-SC. (a): DAC-SC is composed of ResNet50 with $CCNorm$ and $SCNorm$. (b), (c): The proposed $PCNorm$ and $CCNorm$ replace BN in the downsample layer, respectively. We use $CCNorm$ for DAC-SC and $PCNorm$ for DAC-P. (d): $SCNorm$ is attached at the end of stage1 to stage3. (e): A residual block is represented.

(IFT). The $PCNorm$ is defined as follows:

$$PCNorm(f) = IFT(\text{compose}(\alpha^{norm}, \rho)), \quad (11)$$

where $IFT(\cdot)$ denotes IFT.

4.2. Content Controlling Normalization (CCNorm)

$PCNorm$ prevents content from changing. Further, we have a fundamental question regarding whether the variation of content is absolutely negative in DG. We believe that the answer can be clarified by making the model learn itself to reduce the content change. If the change is harmful to DG, the model would gradually decrease it during training.

As we explain in Sec. 3.2, content change occurs due to μ , which is the mean shift in normalization in the spatial domain. Thus, we adjust the degree of content change by introducing a learnable parameter $\lambda^c \in \mathbb{R}^2$. The content adjusting terms, $(\lambda_{norm}^c, \lambda_{org}^c) = \text{softmax}(\lambda^c/T_c)$, where T_c denotes the temperature value, determine the proportion of the normalized and original content, respectively. Then, the content-adjusted feature f^c is defined as follows:

$$f^c = f - \mu \lambda_{norm}^c. \quad (12)$$

In the equation, λ_{org}^c is omitted because it is a dummy variable for the performance, and we only consider the normalized content. If λ_{norm}^c is 0, the phase of f^c (ρ^c) is the same as ρ , and if λ_{norm}^c is 1, ρ^c is the same as ρ^{norm} . Then, we propose the first advanced $PCNorm$, $CCNorm$, which employs ρ^c instead of ρ in $PCNorm$. The $CCNorm$ is defined as follows:

$$CCNorm(f) = IFT(\text{compose}(\alpha^{norm}, \rho^c)), \quad (13)$$

which is illustrated in Fig. 3 (b). The main idea of $CCNorm$ is to mitigate the content variation of normalization, not to entirely offset it. Interestingly, $CCNorm$ performs better than $PCNorm$ in many experiments. Sec. 5 discusses the effects of $CCNorm$ and insights from it.

4.3. Style Controlling Normalization (SCNorm)

It is common to remove style information through IN in DG [19, 34]. Similar to the motivation of $CCNorm$, an essential question occurs regarding whether completely eliminating style is ideal for DG. Thus, we identify the effect of adjusting the style elimination degree.

Therefore, we propose the other advanced version of $PCNorm$, $SCNorm$, that regulates the degree of style elimination. It is illustrated in Fig. 3 (c). The $SCNorm$ method mixes α with α^{norm} by applying learnable parameters. In $SCNorm$, the content is preserved in the same way as in $PCNorm$. Analogous to $CCNorm$, the style-adjusting terms are λ_{norm}^s and λ_{org}^s , which are the outputs of the learnable parameter, $\lambda^s \in \mathbb{R}^2$, obtained using the softmax. That is, $(\lambda_{norm}^s, \lambda_{org}^s) = \text{softmax}(\lambda^s/T_s)$, where T_s represents temperature value. The $SCNorm$ is formulated as follows:

$$SCNorm(f) = IFT(\text{compose}(\lambda_{norm}^s \alpha^{norm} + \lambda_{org}^s \alpha, \rho)). \quad (14)$$

Both adjusting terms, λ_{norm}^s and λ_{org}^s , are learned to make the model independently determine the ratios of pre- and post-normalized style, respectively. The IN is used by default to obtain α^{norm} . Similar to λ^c in $CCNorm$, $SCNorm$ becomes an identity function when λ_{norm}^s is 0, and if λ_{norm}^s is 1, $SCNorm$ completely removes the style. Sec. 5 discusses its effectiveness.

4.4. DAC-P and DAC-SC

In this section, we introduce the ResNet-variant models, DAC-P, and DAC-SC. DAC-P is the early model for feasibility, where $PCNorm$ is applied. DAC-SC is the primary model that adopts both $CCNorm$ and $SCNorm$. The overall architecture of DAC-SC is described in Fig. 4.

In DAC-P, we replace BN in the downsample layer with $PCNorm$ (Fig. 4 (b)). The downsample layer is included in the residual block, which is represented as $\mathcal{H}(x) + x$,

where x is the input feature, and $\mathcal{H}(\cdot)$ is the residual function. Unlike other layers, the shape of x in downsample layer changes to match that of $\mathcal{H}(x)$. That is, the phase of x inevitably changes, although for identity mapping it should be invariant. Thus, the residual $\mathcal{H}(x)$ is approximated to the biased input whose content information is changed. As it contains BN with the content variation problem, the content change in this layer becomes larger. Consequently, the biased approximation of $\mathcal{H}(x)$ degrades DG performance.

To relieve this, we substitute the BN in the downsample layer with $PCNorm$. It is to take advantage of the fact that $PCNorm$ preserves content. First, the existing downsample layer of ResNet [13], $downsample(\cdot)$, is represented as follows:

$$downsample(x) = BatchNorm(Conv(x)), \quad (15)$$

where $BatchNorm(\cdot)$ indicates the BN, and $Conv(\cdot)$ is a 1×1 convolution layer with a stride 2, and x is the input feature. On the other hand, the downsample layer of our DAC-P, $downsample_p(\cdot)$, is formulated as follows:

$$downsample_p(x) = PCNorm(Conv(x)). \quad (16)$$

Next, we introduce the primary model, DAC-SC, where both $CCNorm$ and $SCNorm$ are applied. As explained above, the content change in the downsample layer especially needs to be relieved. In DAC-SC, traditional BN is replaced with $downsample_c$, which uses $CCNorm$ instead of $PCNorm$ (Fig. 4 (c)). Then, $downsample_c$ is defined as follows:

$$downsample_c(x) = CCNorm(Conv(x)). \quad (17)$$

In addition, DAC-SC also exploits $SCNorm$, which adjusts the degree of style elimination. Previous studies [19,57] found that putting a style regularizer between residual blocks enhances performance. Inspired by this, we inserted $SCNorm$ between residual blocks (Fig. 4 (d)). Hence, the primary model, DAC-SC, determines the proper intensities of the content and style changes for DG.

5. Experiment

5.1. Dataset

Dataset	Class	Image	Domain
VLCS [23]	7	9,991	4
PACS [23]	5	10,729	4
Office-Home [47]	65	15,588	4
DomainNet [36]	345	586,575	6
TerraIncognita [2]	10	24,788	4

Table 2. Description of five datasets, VLCS, PACS, Office-Home, DomainNet, and TerraIncognita.

We evaluated the proposed methods on five DG benchmarks: VLCS, PACS, Office-Home, DomainNet, and TerraIncognita. Table 2 summarizes the number of classes, images, and domains in each dataset.

5.2. Experimental Details

We chose ResNet50 as a backbone network, the same as the baseline model (ERM) [46]. In DAC-P, the BN in all four downsample layers was replaced with $PCNorm$ layers. For DAC-SC, four $CCNorm$ were inserted at the same locations as $PCNorm$ in DAC-P, and three $SCNorm$ were added at the ends of the first to third stages, respectively. The affine transform layer of base normalization is transferred to the end of the proposed normalization layer. The model was initialized with ImageNet [8] pre-trained weight. The elements of λ^s and λ^c were initialized with 0, and the temperatures T_s and T_c were set to 1e-1 and 1e-6, respectively. For data augmentation, we randomly cropped images on a scale from 0.7 to 1.0 and resized them to 224×224 pixels. Then, we applied random horizontal flip, color jittering, and gray scaling. In training, we used a mini-batch size of 32, and the Nesterov SGD optimizer with a weight decay of 5e-4, a learning rate of 1e-4, and momentum of 0.9. For DomainNet dataset only, the learning rate of 1e-2 was applied. We trained the proposed model for 20 epochs with 500 iterations each, except for DomainNet, which we trained with 7500 iterations and adopted a cosine annealing scheduler with early stopping (tolerance of 4). All experiments were conducted four times and each performance was evaluated on the training-domain validation set, which reserved 20% of source domain data. The performance values were reported using the average performance for the entire domains, which was evaluated using a single out-of-training domain. We conducted an exhaustive hyperparameter search for model selection and evaluated the models based on accuracy. The hardware and software environments were Ubuntu 18.04, Python 3.8.13, PyTorch 1.12.1+cu113, CUDA 11.3, and a single NVIDIA A100 GPU.

5.3. Comparison with SOTA Methods

The test accuracies of DAC-P, DAC-SC, and recent DG methods on five benchmarks are reported in Table 1. Specifically, DAC-SC, which exhibited the highest average performance improvement, achieved new SOTA results at 87.5%, 70.3%, and 44.9% on the PACS, Office-Home, and DomainNet benchmark datasets. These experimental results indicate that there can be proper degrees of content and style changes for DG. In the TerraIncognita dataset, DAC-P outperformed DAC-SC, which was expected due to the dataset characteristics. The TerraIncognita dataset consists of camera trap images in which objects are captured in four locations. The content between the domains is closely distributed, resulting in fewer domain gaps compared to other datasets. In these cases, preserving the content can be more effective than adjusting changes under a slight domain shift. From the results of this experiment, it is critical to preserve content in DG, but higher performance can be

Model	VLCS	PACS	Office-Home	DomainNet	TerraIncognita	Avg
ERM [46]	77.4 ± 0.3	85.7 ± 0.5	67.5 ± 0.5	41.2 ± 0.2	47.2 ± 0.4	63.8
IRM [1]	78.5 ± 0.5	83.5 ± 0.8	64.3 ± 2.2	33.9 ± 2.8	47.6 ± 0.8	61.6
GroupDRO [41]	76.7 ± 0.6	84.4 ± 0.8	66.0 ± 0.7	33.3 ± 0.2	43.2 ± 1.1	60.7
Mixup [52]	77.4 ± 0.6	84.6 ± 0.6	68.1 ± 0.3	39.2 ± 0.1	47.9 ± 0.8	63.4
MLDG [24]	77.2 ± 0.4	84.9 ± 1.0	66.8 ± 0.6	41.2 ± 0.1	47.7 ± 0.9	63.6
CORAL [43]	78.8 ± 0.6	86.2 ± 0.3	68.7 ± 0.3	41.5 ± 0.1	47.6 ± 1.0	64.5
MMD [25]	77.5 ± 0.9	84.6 ± 0.5	66.3 ± 0.1	23.4 ± 9.5	42.2 ± 1.6	58.8
DANN [10]	78.6 ± 0.4	83.6 ± 0.4	65.9 ± 0.6	38.3 ± 0.1	46.7 ± 0.5	62.6
CDANN [26]	77.5 ± 0.1	82.6 ± 0.9	65.8 ± 1.3	38.3 ± 0.3	45.8 ± 1.6	62.0
MTL [3]	77.2 ± 0.4	84.6 ± 0.5	66.4 ± 0.5	40.6 ± 0.1	45.6 ± 1.2	62.9
ARM [54]	77.6 ± 0.3	85.1 ± 0.4	64.8 ± 0.3	35.5 ± 0.2	45.5 ± 0.3	61.7
VREx [21]	78.3 ± 0.2	84.9 ± 0.6	66.4 ± 0.6	33.6 ± 2.9	46.4 ± 0.6	61.9
RSC [16]	77.1 ± 0.5	85.2 ± 0.9	65.5 ± 0.9	38.9 ± 0.5	46.6 ± 1.0	62.7
IIB [22]	77.2 ± 1.6	83.9 ± 0.2	68.6 ± 0.1	41.5 ± 2.3	45.8 ± 1.4	63.4
SelfReg [20]	77.8 ± 0.9	85.6 ± 0.4	67.9 ± 0.7	42.8 ± 0.0	47.0 ± 0.3	64.2
SagNet [31]	77.8 ± 0.5	86.3 ± 0.2	68.1 ± 0.1	40.3 ± 0.1	48.6 ± 1.0	64.2
DAC-P (ours)	77.0 ± 0.6	85.6 ± 0.5	69.5 ± 0.1	43.8 ± 0.3	49.8 ± 0.2	<u>65.1</u>
DAC-SC (ours)	<u>78.7 ± 0.3</u>	87.5 ± 0.1	70.3 ± 0.2	44.9 ± 0.1	46.5 ± 0.3	65.6

Table 1. Comparison with recent DG methods. For the performances of IIB and SelfReg, we refer to each paper. The other performance results are the reported numbers from DomainBed [11]. The best performance values are in bold and the second-best performance values are underlined.

	Method	P	VL	PA	OH	DN	TI	Avg
1	ResNet		77.4	85.7	67.5	41.2	47.2	63.8
2	+ <i>CCNorm</i>	D	77.1	86.0	69.9	44.7	48.1	65.1(+1.3)
3	+ <i>SCNorm</i>	E	78.7	87.4	70.3	44.9	46.5	65.6(+0.5)
4	+ <i>PCNorm</i>	D	77.0	85.6	69.5	43.8	49.8	65.1(+1.3)
5	+ <i>SCNorm</i>	E	77.1	86.0	68.8	44.6	48.0	64.9(-0.2)
6	+ <i>SCNorm</i>	D	77.3	86.4	69.2	44.5	45.3	64.5(+0.7)
7	+ <i>SCNorm</i>	E	77.3	86.8	69.2	44.9	43.0	64.3(-0.2)

Table 3. Ablation study on our methods. VL, PA, OH, DN, and TI denote VLCS, PACS, Office-Home, DomainNet, and TerraIncognita, respectively. P indicates module position, and D and E of column P mean downsample layer and end of the stage in order.

achieved when the preservation degree can be appropriately controlled. A more detailed discussion of this is presented in Sec. 5.4.

5.4. Ablation Study

In this section, we conducted various ablation studies focusing on the DAC-SC because it is the main model and adopts many advanced factors.

Effect of CCNorm and SCNorm. To investigate the performance improvement effect with the two proposed modules, *CCNorm* and *SCNorm*, we added them sequentially and individually to ResNet and compared their performance (Rows 2 and 3 of Table 3). As described in Sec. 4.4, *CCNorm* is applied to the downsample layer and *SCNorm* is inserted at the end of the stage. As the two modules are inserted in different positions, we added Column P, indicating the position of the applied module for

clarity, where D and E denote the downsample layer and end of the stage, respectively. An average performance improvement of 1.3% compared to the baseline ResNet (ERM) occurs when only *CCNorm* is added. Furthermore, we added both *SCNorm* and *CCNorm*, the DAC-SC, providing an extra improvement of 0.5% to the average performance. It achieves the highest increase of 3.7% on the DomainNet dataset.

In contrast, applying *PCNorm* and *SCNorm* together displayed an average performance degradation of 0.2% compared to applying *PCNorm* alone (Row 5). Considering the performance gains when *SCNorm* is combined with *CCNorm*, it implies that *SCNorm* works better with *CCNorm* than *PCNorm*. These experimental results demonstrate that there is an efficient combination of adjusted content and style, which DAC-SC can successfully determine.

Content Preserving vs. Content Controlling. Rows 2 and 4 of Table 3 reveal that *PCNorm*, which perfectly preserves content, and *CCNorm*, which controls content change, have the same average performance at 65.1%. However, referring to the results of the individual datasets, *CCNorm* outperforms *PCNorm* on the other four datasets except TerraIncognita. It means that controlling content changes (*CCNorm*) is more beneficial in terms of generalization performance.

Content Controlling vs. Style Controlling in Downsample Layer. We replace *CCNorm* with *SCNorm* in

Method	VL	PA	OH	DN	TI	Avg
InNorm	76.0	87.2	65.6	44.0	42.7	63.1
$SCNorm_{in}$	78.7	87.4	70.3	44.9	46.5	65.6
$SCNorm_{ln}$	75.0	87.3	67.9	44.3	44.3	63.8
$SCNorm_{bn}$	77.0	86.6	70.3	44.8	46.9	65.1

Table 4. Ablation study on base normalization in $SCNorm$. InNorm is the simple IN without $SCNorm$. $SCNorm_{in}$, $SCNorm_{ln}$, and $SCNorm_{bn}$ denote IN, LN, and BN applied to $SCNorm$, respectively. The abbreviations of the datasets are the same as in Table 3.

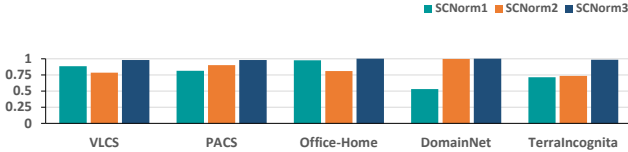


Figure 5. Comparison of the learned values of λ_{org}^s in each $SCNorm$ layer.

DAC-SC to verify that adjusting content changes is more effective than controlling style changes at the downsample layer. Rows 6 and 7 of Table 3 present the results when $SCNorm$ is inserted into the downsample layer. Compared to the baseline ResNet, applying $SCNorm$ to the downsample layer renders a slight average performance improvement of 0.7%. However, this performance enhancement is low compared to the result of applying $CCNorm$ or $PCNorm$ to the downsample layer. From these results, we infer that, although the adjustment to the style change in the downsample layer works lightly, the regulation of the content variation seems more appropriate in this layer.

Which Normalization Fits on $SCNorm$? In $SCNorm$, we apply IN by default. In this experiment, We first compared the performances of our $SCNorm$ with those of the simple IN to clarify that $SCNorm$ is more effective than existing IN. Then, we observed the results when the IN in $SCNorm$ is replaced with the LN or BN to identify that it is optimal to select the IN compared to other methods.

As listed in Table 4, applying IN to $SCNorm$ is more effective for all datasets than using the simple IN (InNorm). Compared with the other two normalization methods (LN and BN) applied to $SCNorm$, the highest average performance is reached when IN is applied to $SCNorm$ except in the TerraIncognita dataset. In contrast, $SCNorm$ with LN consistently performs poorly compared with the other two normalization methods (IN and BN) except for the PACS dataset. These results indicate that it is optimal to apply IN to $SCNorm$.

Adjusting Style Elimination Degree. To identify the degree that style elimination is adjusted through $SCNorm$, we illustrate λ_{org}^s in $SCNorm$ in Fig. 5, where λ_{org}^s is the weight of the original style, which determines the extent of the original style when it is mixed with the normalized style. Moreover, $SCNorm1$, $SCNorm2$, and $SCNorm3$ indi-

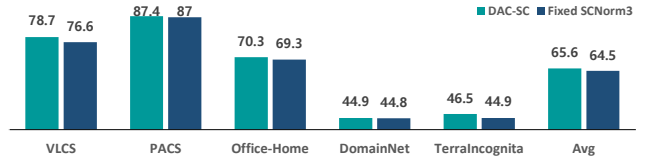


Figure 6. Accuracy comparison of DAC-SC with and without style controlling (λ_{org}^s). Fixed $SCNorm3$ indicates that λ_{org}^s is set to a fixed number of 1.

cate each $SCNorm$ module at the ends of the first, second, and third stages of the DAC-SC, respectively. As demonstrated, all λ_{org}^s were generally distributed above 0.5, but no clear tendency was found for λ_{org}^s of $SCNorm1$ and $SCNorm2$. However, all λ_{org}^s in $SCNorm3$ had an average value of 0.9898, which is very close to 1, in all five datasets. That is, the original style α is nearly preserved at the end of Stage 3.

Then, we deduce that adjusting the style elimination is not required in $SCNorm3$ because the reflection of the normalized style is negligible. To clarify this, we set λ_{org}^s in $SCNorm3$ to a fixed number at 1 and compared the performance. As presented in Fig. 6, the performance consistently drops when λ_{org}^s is 1 in $SCNorm3$. On the VLCS, Office-Home, and TerraIncognita benchmarks, the performance drops by more than 1.0%. This result confirms that even a small change in λ_{org}^s significantly affects performance. Hence, we infer that an adjustment is needed in the style elimination degree for DG, which $SCNorm$ successfully determines.

6. Conclusion

This paper addresses the content change problem in existing normalization methods in DG, and suggests an incipient approach that explores its improvements from the aspect of the frequency domain. For this, we provide a pioneering analysis of the quantitative content change through mathematical derivation. Based on this analysis, we propose novel normalization methods for DG, PCNorm, CCNorm, and SCNorm. With the proposed methods, the ResNet variant models, DAC-P and DAC-SC, both achieve SOTA performance on five DG benchmarks. The experimental results highlight the importance of content preservation in DG, and take a step further in variation adjustments in the existing normalization-based style extraction.

7. Acknowledgement

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2023R1A2C2003379).

References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. **7**
- [2] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. *ArXiv*, abs/1807.04975, 2018. **6**
- [3] Gilles Blanchard, Aniket Anand Deshmukh, Ürün Dogan, Gyemin Lee, and Clayton D. Scott. Domain generalization by marginal transfer learning. *J. Mach. Learn. Res.*, 22:2:1–2:55, 2021. **3, 7**
- [4] Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986. **2**
- [5] John P Castagna and Shengjie Sun. Comparison of spectral decomposition methods. *First break*, 24(3), 2006. **3**
- [6] Chaoqi Chen, Jiongcheng Li, Xiaoguang Han, Xiaoqing Liu, and Yizhou Yu. Compound domain generalization via meta-knowledge encoding. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7109–7119, 2022. **3**
- [7] Guangyao Chen, Peixi Peng, Li Ma, Jia Li, Lin Du, and Yonghong Tian. Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 448–457, 2021. **2**
- [8] 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. Ieee, 2009. **6**
- [9] Xinjie Fan, Qifei Wang, Junjie Ke, Feng Yang, Boqing Gong, and Mingyuan Zhou. Adversarially adaptive normalization for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8208–8217, 2021. **2, 3**
- [10] Yaroslav Ganin, E. Ustinova, Hana Ajakan, Pascal Germain, H. Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. Domain-adversarial training of neural networks. In *J. Mach. Learn. Res.*, 2016. **2, 7**
- [11] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *ArXiv*, abs/2007.01434, 2021. **7**
- [12] Bruce C Hansen and Robert F Hess. Structural sparseness and spatial phase alignment in natural scenes. *JOSA A*, 24(7):1873–1885, 2007. **3**
- [13] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. **2, 4, 6**
- [14] Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *ArXiv*, abs/1903.12261, 2019. **1**
- [15] Xun Huang and Serge J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1510–1519, 2017. **1, 2, 3**
- [16] Zeyi Huang, Haohan Wang, Eric P. Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *ECCV*, 2020. **2, 7**
- [17] Eleftherios Ioannou and Steve Maddock. Depth-aware neural style transfer using instance normalization. *arXiv preprint arXiv:2203.09242*, 2022. **3**
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ArXiv*, abs/1502.03167, 2015. **1**
- [19] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3140–3149, 2020. **1, 3, 5, 6**
- [20] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9619–9628, 2021. **2, 7**
- [21] David Krueger, Ethan Caballero, Jörn-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Rémi Le Priol, and Aaron C. Courville. Out-of-distribution generalization via risk extrapolation (rex). In *ICML*, 2021. **3, 7**
- [22] Bo Li, Yifei Shen, Yezhen Wang, Wenzhen Zhu, Colorado Reed, Jun Zhang, Dongsheng Li, Kurt Keutzer, and Han Zhao. Invariant information bottleneck for domain generalization. *ArXiv*, abs/2106.06333, 2022. **2, 7**
- [23] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5543–5551, 2017. **6**
- [24] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. **3, 7**
- [25] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex Chichung Kot. Domain generalization with adversarial feature learning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018. **2, 7**
- [26] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *ECCV*, 2018. **2, 7**
- [27] Toshihiko Matsuura and Tatsuya Harada. Domain generalization using a mixture of multiple latent domains. In *AAAI*, 2020. **2**
- [28] Rang Meng, Xianfeng Li, Weijie Chen, Shicai Yang, Jie Song, Xinchao Wang, Mingli Song Lei Zhang, Di Xie, and Shiliang Pu. Attention diversification for domain generalization. In *European Conference on Computer Vision (ECCV)*, 2022. **2**
- [29] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, 2013. **3**
- [30] Hyeonseob Nam and Hyo-Eun Kim. Batch-instance normalization for adaptively style-invariant neural networks. In *NeurIPS*, 2018. **3**
- [31] Hyeonseob Nam, Hyunjae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. *2021 IEEE/CVF Conference on Computer*

- Vision and Pattern Recognition (CVPR)*, pages 8686–8695, 2021. 1, 3, 7
- [32] A Oppenheim, Jae Lim, Gary Kopec, and SC Pohlig. Phase in speech and pictures. In *ICASSP'79. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, pages 632–637. IEEE, 1979. 3
- [33] Alan V Oppenheim and Jae S Lim. The importance of phase in signals. *Proceedings of the IEEE*, 69(5):529–541, 1981. 2, 3
- [34] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, 2018. 1, 5
- [35] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018. 3
- [36] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. 6
- [37] Leon N Piotrowski and Fergus William Campbell. A demonstration of the visual importance and flexibility of spatial-frequency amplitude and phase. *Perception*, 11:337 – 346, 1982. 2, 3
- [38] Leon N Piotrowski and Fergus W Campbell. A demonstration of the visual importance and flexibility of spatial-frequency amplitude and phase. *Perception*, 11(3):337–346, 1982. 3
- [39] Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. Dataset shift in machine learning. 2009. 1
- [40] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In *ICML*, 2022. 2
- [41] Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *ICLR*, 2020. 7
- [42] Mattia Segu, Alessio Tonioni, and Federico Tombari. Batch normalization embeddings for deep domain generalization. *arXiv preprint arXiv:2011.12672*, 2020. 3
- [43] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pages 443–450. Springer, 2016. 2, 7
- [44] Duraisamy Sundararajan. *The discrete Fourier transform: theory, algorithms and applications*. World Scientific, 2001. 3
- [45] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Instance normalization: The missing ingredient for fast stylization. *ArXiv*, abs/1607.08022, 2016. 1
- [46] Vladimir Naumovich Vapnik. *Statistical learning theory*. 1998. 6, 7
- [47] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5385–5394, 2017. 6
- [48] Haohan Wang, Xindi Wu, Pengcheng Yin, and Eric P. Xing. High-frequency component helps explain the generalization of convolutional neural networks. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8681–8691, 2020. 2
- [49] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, and Tao Qin. Generalizing to unseen domains: A survey on domain generalization. In *IJCAI*, 2021. 1
- [50] Jing-Yi Wang, Ruoyi Du, Dongliang Chang, Kongming Liang, and Zhanyu Ma. Domain generalization via frequency-domain-based feature disentanglement and interaction. *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. 2, 3
- [51] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14378–14387, 2021. 1, 2, 3
- [52] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. *ArXiv*, abs/2001.00677, 2020. 7
- [53] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 2, 3
- [54] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. In *NeurIPS*, 2021. 3, 7
- [55] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. In *NeurIPS*, 2020. 2
- [56] Xingchen Zhao, Chang Liu, Anthony Sicilia, Seong Jae Hwang, and Yun Fu. Test-time fourier style calibration for domain generalization. *arXiv preprint arXiv:2205.06427*, 2022. 1, 2
- [57] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *ArXiv*, abs/2104.02008, 2021. 1, 3, 6