

# Single Domain Generalization via Normalised Cross-correlation Based Convolutions

WeiQin Chuah<sup>1\*</sup> Ruwan Tennakoon<sup>1\*</sup> Reza Hoseinnezhad<sup>1</sup> David Suter<sup>2</sup> Alireza Bab-Hadiashar<sup>1</sup>  
RMIT University, Australia<sup>1</sup> Edith Cowan University (ECU), Australia<sup>2</sup>  
{wei.qin.chuah, ruwan.tennakoon, rezah, abh}@rmit.edu.au, d.suter@ecu.edu.au

## Abstract

Deep learning techniques often perform poorly in the presence of domain shift, where the test data follows a different distribution than the training data. The most practically desirable approach to address this issue is Single Domain Generalization (S-DG), which aims to train robust models using data from a single source. Prior work on S-DG has primarily focused on using data augmentation techniques to generate diverse training data. In this paper, we explore an alternative approach by investigating the robustness of linear operators, such as convolution and dense layers commonly used in deep learning. We propose a novel operator called “XCNorm” that computes the normalized cross-correlation between weights and an input feature patch. This approach is invariant to both affine shifts and changes in energy within a local feature patch and eliminates the need for commonly used non-linear activation functions. We show that deep neural networks composed of this operator are robust to common semantic distribution shifts. Furthermore, our empirical results on single-domain generalization benchmarks demonstrate that our proposed technique performs comparably to the state-of-the-art methods.<sup>1</sup>

## 1. Introduction

Deep learning techniques have achieved practical success in a variety of fields, including computer vision, natural language processing, and speech processing. However, this success is often limited to settings where the test data follows the same distribution as the training data. In many real-world situations, this assumption breaks down due to shifts in data distribution, known as *domain-shift* [2], which can significantly degrade performance [24].

Dealing with domain-shift is a challenging problem with important practical implications. There are two main ap-




Figure 1. Comparison of accuracy for common normalization methods in image classification on CIFAR-10-C, considering five levels of domain discrepancy caused by corruptions. The evaluated methods include BN, IN, and our proposed approaches (XCNorm and R-XCNorm).

proaches to address domain shift and enable the transfer of knowledge from previously seen environments (*source* domains) to a new environment (*target* domain) without using any labeled data of the target domain: (1) Domain Adaptation [31] (DA) where a model trained with source data is recalibrated using unlabeled data from the target domain, and (2) Domain generalisation [36] (DG) where a model is trained on multiple source domains but no target domain data is available for recalibration. The most *data-efficient* domain generalisation technique is the single domain generalisation (S-DG), which requires data from only a single source domain to train a model that is robust against unforeseen data shifts. Although practically desirable, S-DG has received little attention in the past.

S-DG presents a significant challenge due to two main factors. Firstly, the input data, derived from only one source domain, does not provide sufficient opportunity to observe the possible diversity in out-of-domain data. Secondly, the presence of spurious correlations or shortcuts can further complicate the issue by introducing biases and hindering generalization. Prior work on S-DG has primarily focused on increasing the diversity of input data using adaptive data

\*WeiQin Chuah and Ruwan Tennakoon assert equal contributions.

<sup>1</sup>Our implementation is available at [GitHub](https://github.com).

augmentation techniques. These include creating fictitious examples that mimic anticipated shifts in data distribution using random [34], adversarial [8, 21, 26, 35] or causality [5, 20] based data augmentation, as well as image style diversification [29].

The generalization of a model is largely influenced by its support, which refers to the diversity of its input data and its inductive biases [30]. While not explicitly stated as such, the success of the above-mentioned S-DG methods hinges on increasing the input diversity via data augmentation. An alternative and complementary approach that has received less attention is to incorporate inductive biases into the network components to make them more robust to domain shifts. In this paper, we explore the above approach and propose a robust alternative to linear operators, such as convolution and dense layers, which are fundamental components of most neural networks.

We draw on the classical idea of template matching and consider linear layers in a neural network as computing a matching between the template (represented by the weights) and the signal (represented by input feature maps) using *cross-correlation*, as detailed in Equation (4). Early works in template matching have shown that cross-correlation is not ideal for pattern matching as it fails when the local energy of the input (i.e.,  $\sum_{u,v} z[u,v]^2$ ) varies, and is not robust to affine shifts in the input signal [14]. More recently, Jin *et al.* [12] empirically showed that domain shift primarily causes variation in the local energy of feature representations. This suggests that the linear operator, which is sensitive to local energy, might degrade out-of-domain (OOD) generalization.

The above perspective enables us to use more robust template-matching techniques such as normalized cross-correlation [14] to replace convolutions or dense layers in neural networks and recover the underlying invariant features in the input. We call our method “*XCNorm*”, which performs cross-correlation between standardized (i.e., Z-score normalized) weights and patch-wise standardized input feature maps. This reduces the influence of input feature magnitude on the output and makes the operator invariant to affine transformations of the input. Moreover, we leverage robust statistics to improve the resilience of *XCNorm* to outliers and introduce a refined version of our method named *R-XCNorm*. As Figure 1 demonstrates, our methods achieve significantly better robustness to semantic distribution shifts on CIFAR-10-C, in contrast to other normalization techniques. Moreover, the advantage of our methods becomes more pronounced as the domain discrepancy increases. The contributions of this paper include:

- We propose a novel nonlinear operator called “*XCNorm*”, based on normalized cross-correlation, that reduces the influence of input feature magnitude on the output and invariant to affine transformations.

- Leveraging robust statistics, we further enhance the robustness of *XCNorm* to outliers. Our experiments on several commonly used benchmarks in S-DG show that the proposed robust operator (“*R-XCNorm*”) is also complementary to augmentation-based methods and achieves state-of-the-art performance.
- We empirically show that a neural network based on “*XCNorm*” or “*R-XCNorm*”, is significantly more robust to semantic shifts compared to a network based on a typical linear operator.

## 2. Related Work

**Domain Generalization:** Domain generalization (DG) methods aim to learn robust models from several source domains that can generalize to unseen target domains, addressing the issue of domain shifts.

A particularly challenging yet practical variant of domain generalization (DG) is a single-domain generalization (S-DG), where only one source domain is available during training. S-DG is especially challenging because, unlike in DG, there is no access to multiple source domains that would allow for the observation of possible shifts in data and invariances between domains. To address this challenge, researchers have primarily focused on using data augmentation techniques to generate diverse training data and increase input diversity. A common technique is posing S-DG as a “distributionally robust optimization” problem and solving it using adversarial data augmentation (ADA) [26]. ADA lacks the ability to produce large semantic shifts that are common in real data. As a result, subsequent works have added additional constraints to adversarial augmentation [8, 15, 21, 35] or incorporated background knowledge about anticipated semantic shifts via random augmentations [34], causality based data augmentations [5, 20], or image style diversification [29].

An alternative that has received little attention is to incorporate inductive biases into the network components to make them more robust to domain shifts. The most closely related work in this direction is the Meta Convolution Neural Networks (Meta-CNN) proposed by Wan *et al.* [27], where the output feature-maps of each layer are reconstructed using templates learned from training data, resulting in universally coded images without biased information from unseen domains. Our proposed operator, *XCNorm*, offers a simpler implementation compared to [27]. While their method involves more complex operations, our approach simply replaces the convolution function with our *XCNorm* method. This simplicity makes our method more straightforward to implement and integrate into existing frameworks.

**Normalization in Neural Networks:** During the training process of a deep neural network, the input distribu-

tion of an intermediate layer continuously changes, a phenomenon known as covariate shift. This makes it challenging to set the hyper-parameters of a layer, such as the learning rate. To tackle this issue, various normalization techniques have been proposed, including Batch Norm [11], Instance Norm [25], GroupNorm [33], and Layer Norm [1], which aim to normalize the output of each layer using batch statistics, feature channels, groups of features, or the entire layer’s output, respectively. Instead of operating on features, Weight Norm [22] proposes normalising the filter weights.

While most of the aforementioned work has focused on in-domain generalization, there are a few studies that have examined generalization ability under domain shift. For instance, BN-Test [19] computed batch normalization statistics on the test batch while DSON [23] used multi-source training data to compute the statistics. Fan *et al.* [8] investigated normalization for single-domain generalization, where adaptive normalization statistics for each individual input are learned. These adaptive statistics are learned by optimizing a robust objective with adversarial data augmentation. The above works view normalization as being independent of the base operator (e.g., convolution, fully-connected). In contrast, our approach considers normalization to be an integral part of the base operator. We normalize both weights and input for each local spatial region of the input.

**Non-linear Transforms:** Several works have explored the use of non-linear transforms to replace the linear transform in the typical convolution operator [9, 16–18, 28, 37]. The works most closely related to ours are those by [3, 16–18] assess the cosine similarity between weights and inputs to improve both model performance and interpretability. In those methods, the convolution is viewed as an inner product between an input patch  $\mathbf{z}$  and the weights  $\mathbf{w}$ :

$$\begin{aligned} \text{Conv}(\mathbf{z}; \mathbf{w}) &= \langle \mathbf{z}, \mathbf{w} \rangle = \|\mathbf{z}\| \|\mathbf{w}\| \cos(\varphi) & (1) \\ &= h(\|\mathbf{z}\|, \|\mathbf{w}\|) g(\cos(\varphi)) & (2) \end{aligned}$$

where  $\cos(\varphi)$  is the angle between  $\mathbf{z}$  and  $\mathbf{w}$ . This view allows for the separation of norm terms from the angle term (decoupling), and to change the form of  $h()$  and  $g()$  independently. Liu *et al.* [16] derived several decoupled variants of the functions  $h()$  and  $g()$ . They demonstrated that the decoupled reparameterizations lead to significant performance gains, easier convergence, and stronger adversarial robustness. Later [3] introduced the “B-cos” operator and showed that it lead to better neural network interpretations.

Our proposed *XCNorm* operator can also be seen within this framework, where the dot product is taken between the centered and normalized versions of the input patch and the weights. While the general formula presented in Equation (2) encompasses a range of potential frameworks as




Figure 2. Probabilistic graphical model representing the data generation process. Circles denote random variables, such as  $X$  as input (covariates) and  $Y$  as target variable, and solid arrows represent direct dependencies between variables. The shaded circles  $R$  and  $X$  indicate that the distribution of those random variables shifts between source and target environments.

demonstrated in [16], it’s important to note that our *XCNorm* and its robust variant have not been previously proposed or thoroughly examined. Moreover, unlike the methods mentioned above, we investigate the use of *XCNorm* for out-of-domain generalization in a single source domain setting.

### 3. Method

Single Domain Generalization (SDG) aims to learn a robust predictive function  $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$  from a single source domain  $\mathcal{S} \sim P_{XY}^S$  that can achieve a minimum prediction error on an unseen target domain  $\mathcal{T} \sim P_{XY}^T$ , where  $P_{XY}^S \neq P_{XY}^T$ . However,  $X$  is constructed from a mix of domain-invariant factors  $X_0$  and domain rendering factors  $R$ , and only the former corresponds to the label  $Y$  as illustrated in Figure 2. Moreover, as  $R$  differs between source and target domains, causing domain shifts, a predictive function  $f_\theta$  that is invariant to  $R$  would greatly improve SDG.

To this extent, we draw on the classical idea of template matching and propose *XCNorm* that computes the normalized cross-correlation between weights and an input feature patch. Our method is invariant to both affine shifts and changes in energy within a local feature patch, recovering the underlying invariant features  $X_0$  in the input. Additionally, *XCNorm* inherently integrates normalization and non-linear operations, negating the necessity for normalization layers and activation functions commonly employed in convolutional neural networks.

#### 3.1. Normalized Cross-Correlation Layer

Typically, the linear units of a DNN layer compute the cross-correlation between the input feature maps  $\mathbf{z} \in \mathbb{R}^{H \times W \times C_{in}}$ <sup>1</sup> and the weights  $\mathbf{w} \in \mathbb{R}^{K \times K \times C_{in} \times C_{out}}$ <sup>2</sup>. Here  $C_{in(out)}$  is the number of channels in the input (or output),  $[H, W]$  are the feature map spatial dimensions and  $K$

<sup>1</sup>Batch dimension is omitted for simplicity.

<sup>2</sup>We use notations that are consistent with convolutional layers for convenience. For a fully connected (dense) layer, we assume that  $H = W = K = 1$ .

is the kernel size. The pixel  $(u, v)$  of the  $c^{\text{th}}$  output feature channel is computed as:

$$\Phi_{u,v,c}(\mathbf{z}; \mathbf{w}) = \langle \mathbf{z}_{\tilde{u},\tilde{v}}, \mathbf{w}_c \rangle \quad (3)$$

$$= \sum_j z_{\tilde{u},\tilde{v}}^{(j)} \cdot w_c^{(j)} \quad (4)$$

Here  $\mathbf{z}_{\tilde{u},\tilde{v}}$  is a patch of the input feature map centered at  $\tilde{u}, \tilde{v}$  with the same shape as the weight tensor  $\mathbf{w}_c$ , and  $j$  index the pixels within the patch. The map  $\nu : (u, v) \rightarrow (\tilde{u}, \tilde{v})$  is determined by parameters of the convolution (or dense) layer (i.e., stride, kernel width).

The above operator is ineffective for pattern matching when the patch energy (i.e.,  $\|\mathbf{z}_{\tilde{u},\tilde{v}}\| = \sum_j (z_{\tilde{u},\tilde{v}}^{(j)})^2$ ) is not uniform across a feature map. It also lacks robustness to affine transformations of the input, i.e.,  $\langle \mathbf{A}\mathbf{z}_{\tilde{u},\tilde{v}} + \mathbf{b}, \mathbf{w}_c \rangle \neq \langle \mathbf{A}\mathbf{z}_{\tilde{u},\tilde{v}} + \mathbf{b}, \mathbf{w}_c \rangle$  [14]. To overcome the above limitations, we propose using the normalized cross-correlation operator as a replacement for the linear operator. With this new operator, which we coin *XCNorm*, the output feature at position  $(u, v)$  of the  $c^{\text{th}}$  feature channel is computed as:

$$\Psi_{u,v,c}(\mathbf{z}; \mathbf{w}) = \frac{\sum_j [z_{\tilde{u},\tilde{v}}^{(j)} - \mu_{\mathbf{z}(\tilde{u},\tilde{v})}] [w_c^{(j)} - \mu_{\mathbf{w}_c}]}{\|\mathbf{z}_{\tilde{u},\tilde{v}} - \mu_{\mathbf{z}(\tilde{u},\tilde{v})}\| \|\mathbf{w}_c - \mu_{\mathbf{w}_c}\|} \quad (5)$$

$$= \frac{\sum_j z_{\tilde{u},\tilde{v}}^{(j)} \cdot w_c^{(j)} - \alpha \mu_{\mathbf{z}(\tilde{u},\tilde{v})} \mu_{\mathbf{w}_c}}{\alpha \sigma_{\mathbf{w}_c} \sqrt{\frac{1}{\alpha} \sum_j z_{\tilde{u},\tilde{v}}^{(j)} \cdot z_{\tilde{u},\tilde{v}}^{(j)} - \mu_{\mathbf{z}(\tilde{u},\tilde{v})}^2} + \epsilon} \quad (6)$$

Here,  $\mu_{\bullet}$  is the mean of  $\bullet$ ,  $\sigma_{\mathbf{w}_c} = \sqrt{\frac{1}{\alpha} \sum_j [w_c^{(j)} - \mu_{\mathbf{w}_c}]^2}$ ,  $\epsilon$  is a small constant to ensure numerical stability, and  $\alpha = K \times K \times C_{in}$  is the number of pixels in the patch.

Since the patch-wise mean,  $\mu_{\bullet}$ , can be computed using linear operations (i.e., convolution of a constant weight tensor with all elements being equal to  $1/\alpha$ ), the *XCNorm* can be realized using linear operators:

$$\Psi(\mathbf{z}; \mathbf{w}) = \frac{\Phi(\mathbf{z}; \mathbf{w}) - \alpha \Phi(\mathbf{z}; \mathbf{w}_\alpha) \mu_{\mathbf{w}}}{\alpha \sqrt{[\Phi(\mathbf{z}^2; \mathbf{w}_\alpha) - \Phi(\mathbf{z}; \mathbf{w}_\alpha)^2] \sigma_{\mathbf{w}} + \epsilon}} \quad (7)$$

$$= \frac{\Phi(\mathbf{z}; \mathbf{w}) - \alpha \mu_{\mathbf{z}} \mu_{\mathbf{w}}}{\alpha \sqrt{[\mu_{\mathbf{z}^2} - \mu_{\mathbf{z}}^2] \sigma_{\mathbf{w}} + \epsilon}} \quad (8)$$

Here,  $\mathbf{w}_\alpha \in \mathbb{R}^{K \times K \times C_{in} \times 1}$  is a constant matrix with all elements equal to  $1/\alpha$ , and  $\mu_{\mathbf{w}} \in \mathbb{R}^{1 \times C_{out}}$  and  $\sigma_{\mathbf{w}} \in \mathbb{R}^{1 \times C_{out}}$  are the mean and variance of the weights within each output channel.

### 3.2. Robust XCNorm

The sensitivity of *XCNorm* to outliers in the input patch [7] can lead to issues when the input distribution changes unpredictably or includes outliers. For instance, salt and pepper noise can cause large variations in the input

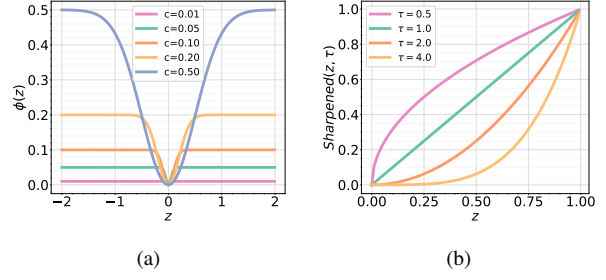


Figure 3. (a) Welsch function with varying  $c$ . (b) Sharpening the output of  $\Upsilon(\mathbf{z}; \mathbf{w})$  (Eq. (11)). Large values of  $\tau$  exaggerate small deviation from perfect alignment between  $\mathbf{z}$  and  $\mathbf{w}$ .

energy (first term in the denominator of Equation (5)) and affect the output of *XCNorm*. To overcome this, we propose a robust version of the operator called *R-XCNorm*, which modifies Equation (5) as follows:

$$\Gamma_{u,v,c}(\mathbf{z}; \mathbf{w}) = \frac{\sum_j [\phi(z_{\tilde{u},\tilde{v}}^{(j)} - \mu_{\mathbf{z}(\tilde{u},\tilde{v})})] [w_c^{(j)} - \mu_{\mathbf{w}_c}]}{\|\phi(\mathbf{z}_{\tilde{u},\tilde{v}} - \mu_{\mathbf{z}(\tilde{u},\tilde{v})})\| \|\mathbf{w}_c - \mu_{\mathbf{w}_c}\| + \epsilon} \quad (9)$$

Here,  $\phi(\cdot)$  can be any robust function such as the Huber, Cauchy (aka Lorentzian), Tukey, or Welsch function. In this work, we use the Welsch function due to its simplicity, which is defined as follows [6]:

$$\phi(z) = c \left[ 1 - \exp\left(\frac{-z^2}{2c^2}\right) \right] \quad (10)$$

Here,  $c$  is a learnable parameter that controls the amount of penalty for the outliers. The influence of  $c$  is depicted in Figure 3a. During the training process, we adopt an initialization strategy where we set the value of the parameter  $c$  to a large value. Subsequently, we update the value of  $c$  for each layer by computing the mean of the patchwise standard deviation of the input over the training dataset. Similar to batch normalization, we incorporate a moving average component to enhance the stability and effectiveness of the normalization process.

### 3.3. Improving Convergence

To enhance the convergence of *XCNorm* and *R-XCNorm*, we incorporate some modifications to their base formulations. We use the notation  $\Upsilon(\mathbf{z}; \mathbf{w})$  to denote either  $\Psi(\mathbf{z}; \mathbf{w})$  or  $\Gamma(\mathbf{z}; \mathbf{w})$  in the following text.

**Sharpening:** The peaks and valleys of the output  $\Upsilon(\mathbf{z}; \mathbf{w})$  can be emphasized (or de-emphasized) by raising it to power  $\tau$ :

$$\tilde{\Upsilon}_{[\tau]}(\mathbf{z}; \mathbf{w}) = \{\max[0, \Upsilon(\mathbf{z}; \mathbf{w})]\}^\tau \quad (11)$$

Here, we solely consider the positive outputs, as our empirical observations indicate that this choice stabilizes the training process and prevents convergence to undesirable optima. Figure 3b shows the relationship between the input and output of the above operation. Similar approaches have also been used in cosine similarity-based techniques [32]. However, our approach differs from cosine similarity-based methods in that we do not pre-determine the value of  $\tau$ . Instead, we treat it as a learnable parameter and optimize it alongside the weights.

**Gradient Scaling:** The weight normalization in Equation (9) tends to reduce the gradient magnitude, which in turn leads to slower convergence [16]. To mitigate this issue, we propose *gradient scaling* using a learnable scaling factor  $\mathbf{A} \in \mathbb{R}^{C_l}$ , where  $C_l$  is the feature channel dimension of layer  $l$ . More specifically, we apply different  $\mathbf{A}$  to the output of  $\tilde{\Upsilon}_{[1]}(\mathbf{z}; \mathbf{w})$  at every layer, as shown below:

$$\tilde{\Upsilon}_{[2]}(\mathbf{z}; \mathbf{w}) = \mathbf{A} \odot \tilde{\Upsilon}_{[1]}(\mathbf{z}; \mathbf{w}). \quad (12)$$

Given that scaling the output of a function by a constant is equivalent to augmenting the gradient of the function by that constant, the proposed method effectively addresses the reduced gradient issue and accelerates the convergence process.

**Norm-Based Attention Mask (NBAM):** The input norm  $\|\tilde{\mathbf{z}}\|$  signifies the importance of the local patch within the input. Here,  $\tilde{\mathbf{z}} = \mathbf{z}_{\tilde{u}, \tilde{v}} - \mu_{\mathbf{z}(\tilde{u}, \tilde{v})}$  for *XCNorm* and  $\tilde{\mathbf{z}} = \phi(\mathbf{z}_{\tilde{u}, \tilde{v}} - \mu_{\mathbf{z}(\tilde{u}, \tilde{v})})$  for *R-XCNorm*. Normalizing with  $\|\tilde{\mathbf{z}}\|$  assigns equal importance to all patches, but this may be problematic when  $\|\tilde{\mathbf{z}}\|$  is very small. Such small values indicate low-variation (low-information) input areas that should not be equally weighted as high-information areas, as this may cause spurious matches with templates.

To address this issue, we propose the *Norm-based Attention Mask* technique, which leverages a lightweight single convolution layer with sigmoid activation, denoted as  $\psi$ . This function learns to dynamically assign importance weights to image patches. Specifically, by taking  $\|\tilde{\mathbf{z}}\|$  as input,  $\psi$  learns to generate patch-wise importance weights  $m \in [0, 1]$ . Subsequently, the output is redefined using the following computation:

$$\begin{aligned} \tilde{\Upsilon}_{[3]}(\mathbf{z}; \mathbf{w}) &= \psi(\|\tilde{\mathbf{z}}\|) \odot \tilde{\Upsilon}_{[2]}(\mathbf{z}; \mathbf{w}) \\ &+ (1 - \psi(\|\tilde{\mathbf{z}}\|)) \odot \left\{ \tilde{\Upsilon}_{[2]}(\mathbf{z}; \mathbf{w}) \odot \|\tilde{\mathbf{z}}\| \right\} \end{aligned} \quad (13)$$

where  $\Gamma(\mathbf{z}; \mathbf{w}) \odot \|\tilde{\mathbf{z}}\|$  represents the output without normalization.

To further ensure that each feature is weighted equally, we normalize each channel in the output tensor:

$$\tilde{\Upsilon}_{[4]}(\mathbf{z}; \mathbf{w}) = \frac{\tilde{\Upsilon}_{[3]}(\mathbf{z}; \mathbf{w}) - \mu_{\tilde{\Upsilon}}}{\sigma_{\tilde{\Upsilon}}} \quad (14)$$

Method	SVHN	MNIST-M	SYN	USPS	Avg
ERM	27.83	52.72	39.65	76.94	49.29
CCSA	25.89	49.29	37.31	83.72	49.05
d-SNE	26.22	50.98	37.83	<b>93.16</b>	52.05
JiGen	33.80	57.80	43.79	77.15	53.14
ADA	35.51	60.41	45.32	77.26	54.63
M-ADA	42.55	67.94	48.95	78.53	59.49
ME-ADA	42.56	63.27	50.39	81.04	59.32
RandConv	57.52	<u>87.76</u>	62.88	83.36	72.88
L2D	62.86	87.30	63.72	83.97	74.46
MetaCNN	<b>66.50</b>	<b>88.27</b>	70.66	<u>89.64</u>	<b>78.77</b>
XCNorm	59.92	66.67	67.33	85.24	69.79
R-XCNorm	<u>65.42</u>	71.00	<u>71.25</u>	89.04	74.18
R-XCNorm <sup>+RC</sup>	64.25	83.22	<u>72.25</u>	88.89	<u>77.15</u>

Table 1. Results of single domain generalization experiments on the Digits-DG benchmark. Our proposed method demonstrates superior performance compared to adversarial augmentation-based approaches while achieving comparable results with the state-of-the-art. Additionally, our method complements existing data augmentation techniques, including Random Convolution (RC). The **best results** are denoted as bold and underlined, while the **second best results** are indicated by being underlined.

where  $\mu_{\tilde{\Upsilon}}$  and  $\sigma_{\tilde{\Upsilon}}$  are the mean and the variance of  $\tilde{\Upsilon}_{[3]}(\mathbf{z}; \mathbf{w})$  along each channel.

## 4. Experiments

### 4.1. Datasets and Settings

**Digits-DG:** consists of five distinct subsets: MNIST, MNIST-M, SVHN, SYN and USPS. Each subset represents a different domain with variations in writing styles and quality, scales, backgrounds, and strokes. We mainly utilise the Digits-DG benchmark for single-source domain evaluation and ablation studies. Following [21, 26, 27, 29], we chose the first 10,000 images from both the MNIST training and validation sets as the source dataset.

**CIFAR-10-C [10]:** is typically used for corruption robustness benchmarking. It contains 15 different corruption types that mimic real-world scenarios, such as noise, blur, weather, and digital artifacts. Each corruption type has five levels of severity. We follow the setup of [27, 29] and use the CIFAR-10 training set as the source dataset while images in CIFAR-10-C are used for evaluation.

**Camelyon-17-Wilds [13]:** comprises 455k histopathology image patches extracted from 1000 whole-slide images (WSIs) of sentinel lymph nodes [4]. These WSIs were obtained from five distinct medical centres, with each centre representing a unique domain within the dataset. The primary objective of this dataset is to classify input patches and determine whether the central region contains any tumour tissue. In our experimental setup, we selected each of the domains as the source domain in turn and used the rest of the domains as target domains.

**Implementation Details:** As mentioned in Section 3, XC-

Method	Weather	Blur	Noise	Digits	Average
ERM	67.28	56.73	30.02	62.30	54.08
ADA	72.67	67.04	39.97	66.62	61.58
M-ADA	75.54	63.76	54.21	65.10	64.65
ME-ADA	74.44	71.37	66.47	70.83	70.78
RandConv*	74.90	74.95	55.71	76.80	70.59
L2D	75.98	69.16	73.29	72.02	72.61
MetaCNN	<b>77.44</b>	76.80	<b>78.23</b>	81.26	<b>78.43</b>
XCNorm	72.98	73.76	48.15	77.52	68.10
R-XCNorm	75.57	75.72	60.81	77.61	72.43
R-XCNorm <sup>+RC</sup>	76.99	<b>79.61</b>	63.95	<b>82.27</b>	75.70

Table 2. Experiments of single domain generalization on CIFAR-10 classification. Models are trained on the CIFAR-10 train set and evaluated on the CIFAR-10-C benchmark. “\*” represents our implementation of existing work. Our method is also complementary to existing data augmentation methods such as Random Convolution (RC). The **best results** are denoted as bold and underlined, while the second best results are indicated by being underlined.

Norm incorporates normalization and non-linear operations by default. Thus, in our implementation, we replace all *Conv-Norm-ReLU* blocks as well as the final classifier layer within our chosen convolutional neural network architecture with our proposed XCNorm or R-XCNorm.

Complete details regarding the experimental setup, including the network architecture and model selection, can be found in the supplementary document, providing a comprehensive understanding of our methodology.

#### 4.2. Comparisons on Digits-DG

**Results:** Table 1 provides a comprehensive comparison of the out-of-domain generalization performance between our proposed method and state-of-the-art approaches. The results highlight the significant improvements achieved by our proposed method. Both our base method (XCNorm) and robust variant (R-XCNorm) showcase substantial enhancements over the ERM baseline (49.3% → 69.8% and 49.3% → 74.2%, respectively), without the need for data augmentation techniques or extensive network modifications. Notably, our methods outperform adversarial augmentation-based domain generalization approaches, including ADA, M-ADA, and ME-ADA, by a considerable margin. Also, Table 1 highlights the complementary nature of our method with data augmentation techniques, such as Random Convolution (RC) [34]. The combined approach achieves even higher performance, positioning it competitively alongside the complex state-of-the-art method, MetaCNN.

#### 4.3. Comparisons on CIFAR-10-C

**Results:** The average accuracy for four categories of level-5 severity corruptions is presented in Table 2. Our proposed method showcases substantial improvements over the baseline model (ERM), achieving an impressive accu-




Figure 4. Comparison of model performance on the Camelyon-17 dataset for single domain generalization. Our proposed method surpasses the baseline ERM method, while integrating with the Random Convolution (RC) approach leads to further improvements in domain generalization performance.

racy boost from 54.08% to 68.10% without relying on data augmentation techniques. Notably, our approach demonstrates a remarkable 17.03% improvement for *blur* corruptions and a noteworthy 15.22% improvement for *digits* corruptions. Furthermore, our R-XCNorm method effectively enhances model robustness against noise corruption, elevating the accuracy from 48.15% to 60.81%. It is worth highlighting that our R-XCNorm outperforms most of the data augmentation-based approaches by a respectable margin. Additionally, when combined with RC [34], our method exhibits exceptional performance across all four categories and demonstrates competitive results compared to the current state-of-the-art method.

#### 4.4. Comparisons on Camelyon-17

**Results:** While the Camelyon17 dataset is commonly employed for conventional domain generalization tasks (generalizing from multiple source domains to a single target domain), it has not been extensively explored for single domain generalization. In this study, we investigate the single domain generalization performance on the Camelyon17 dataset using the AUROC metric, as shown in Figure 4. Notably, the ERM model exhibits poor generalization when trained on a single source domain, particularly in domains 2 and 3. We attribute this observation to significant variations in the staining agent’s color across different hospitals (refer to the supplementary document for image samples collected from different domains). In contrast, the Random Convolution (RC) approach demonstrates impressive domain generalization capabilities on the Camelyon17 dataset.

Remarkably, our proposed XCNorm method and its robust variant, R-XCNorm, consistently outperform the ERM baseline across all domains, without relying on data augmentation techniques. In fact, our methods even surpass the performance of the RC method. Moreover, when combined with RC, our approach achieves a robust model with exceptional domain generalization performance. These re-

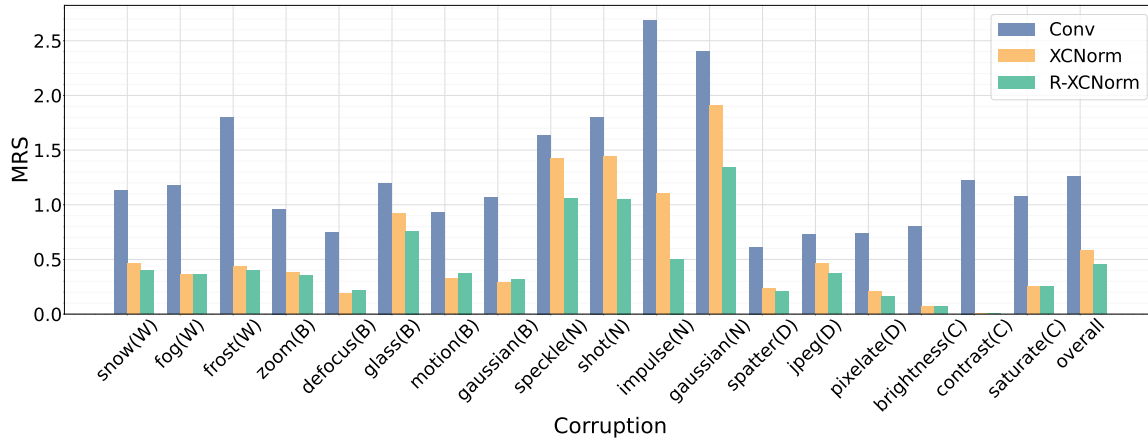


Figure 5. Model Robustness Score (MRS) comparison of different models on the CIFAR-10-C dataset. The MRS values were computed for various categories, including weather (W), blur (B), digital compression (D), noise (N), and chromatic (C) augmentations. Lower MRS values indicate better model robustness against corruption. The baseline ERM model exhibits poor performance, while our proposed method demonstrates significantly better robustness compared to ERM.

XCNorm	NBAM	Robust	SVHN	MNIST-M	SYN	USPS	Avg
✗	–	–	30.0	55.2	39.4	75.9	50.1
✓	✗	✗	59.9	66.7	67.3	85.2	69.8
✓	✓	✗	62.0	67.4	68.8	87.9	71.5
✓	✗	✓	64.0	69.1	70.8	87.3	72.8
✓	✓	✓	65.4	71.0	71.3	89.0	74.2

Table 3. Ablation results on Digits-DG benchmark, and the Top-1 accuracy is reported. Our method can substantially improve the generalization performance of the baseline model.

sults highlight the effectiveness of our method in enhancing domain generalization and its potential to improve model robustness in practical applications, such as medical imaging classification.

## 5. Discussion

### 5.1. Ablation Study

In this section, we present the results of our ablation study conducted on the Digits-DG benchmark to evaluate the efficacy of each component of our proposed method. Specifically, we examine the effectiveness of our XC-Norm method, the robust variant (R-XCNorm), and the norm-based attention mask (NBAM) proposed to improve model convergence, as discussed in Section 3.3.

Table 3 reports the classification results of the four variants of our original framework, including the baseline ERM model for comparison. Our XCNorm method, without any extension, achieves a significant 19.7% performance improvement over the baseline, demonstrating the effectiveness of our approach. Also, the integration of NBAM, which relaxes the normalization of input features, leads to

additional performance gains. This highlights the importance of selectively normalizing important input regions, as a global normalization approach may mistakenly assign equal importance to insignificant regions and hinder training.

Furthermore, we evaluate the effectiveness of the robust variant, R-XCNorm, specifically designed for outlier rejection. As illustrated in Table 3, the integration of the Welsch robust function enhances performance across all domains except for the USPS dataset. Moreover, combining the robust variant with NBAM yields even greater improvements across all domains without sacrificing performance. These results highlight the benefits of incorporating the robust variant and NBAM in our framework, demonstrating their potential for enhancing model performance and domain generalization.

### 5.2. Gradient Scaling

In this section, we investigate the effect of gradient scaling, as proposed in Section 3.3, on the performance of XC-Norm. As mentioned, The weight normalization employed in Equation (9) tends to reduce the gradient magnitude during backpropagation, as shown in Equation (15):

$$\frac{\partial}{\partial \tilde{w}} \Gamma(\tilde{z}; \tilde{w}) = \frac{\hat{z} - \hat{w}^\top \hat{z} \cdot \hat{w}}{\|\hat{w}\|} \quad (15)$$

where  $\hat{z} = \tilde{z}/|\tilde{z}|$  and  $\hat{w} = \tilde{w}/|\tilde{w}|$ . From Equation (15), we observe that the large norm  $|\hat{w}|$  can diminish the gradients, significantly slowing down the convergence of the optimization process.

To address this issue, we propose a gradient scaling method. As shown in Figure 6, we compare the conver-




Figure 6. Comparison of optimization convergence rates with and without the proposed gradient scaling. The results highlight the superior convergence speed achieved by incorporating our proposed technique, leading to reduced training time and convergence to a lower optimal point.

gence of XCNorm with and without the proposed gradient scaling. The results clearly demonstrate that XCNorm with gradient scaling achieves faster convergence than the variant without gradient scaling. This confirms the effectiveness of our proposed gradient scaling method in overcoming the problem of diminishing gradients caused by weight normalization, ultimately speeding up the training convergence.

### 5.3. Coherence of XCNorm and Random Convolution

Integrating Random Convolution (RC) [34] with XC-Norm further boosts single domain generalization performance, especially with blur, noise, and digital corruptions (see Table 2). We postulate that this is due to the diversification of texture information introduced by RandConv which prevents learning highly “texture-sensitive templates”. RC generates new images with comparable global shapes to the source data while introducing random local textures strengthening the model’s resistance against local texture variations inherent in the mentioned corruptions.

### 5.4. Model Sensitivity to Input Perturbations

In this section, we investigate the sensitivity of our model to input perturbations using the CIFAR-10-C dataset. Our analysis focuses on the variance in predicted class probabilities,  $P(y|x)$ , when using a set of corrupted images with different severity levels (denoted as  $s$ ) ranging from 0 to 5. A severity level of 0 represents no corruption, while a severity level of 5 indicates the most severe corruption.

As depicted in Figure 2, the input  $X$  is influenced by the causal factor  $X_0$  (e.g., semantics) and the rendering factor  $R$ , which contributes to domain shift. Since the objective of S-DG is to develop a model that is invariant to  $R$  as much as possible, it is crucial to examine the effect of  $R$  on the predictions of the developed model. We measure this

by computing the Model Robustness Score (MRS), which quantifies the discrepancy between the predictions obtained using clean and perturbed inputs:

$$\text{MRS}(x, \xi) = \frac{1}{5} \sum_{s=1}^5 KL(f_{\theta}(\xi(x; 0)); f_{\theta}(\xi(x; s))) \quad (16)$$

where,  $\xi(x; s)$  represents the perturbation (e.g., blur, noise, compression, weather) included in the CIFAR-10-Corruption dataset. A model that relies solely on  $X_0$  for prediction would exhibit a lower MRS compared to a model that incorporates  $R$  in its predictions.

Figure 5 shows that the ERM model is highly susceptible to  $R$ , leading to large MRS values across all categories. In contrast, the proposed XCNorm effectively improves the model’s robustness, resulting in models that are less reliant on  $R$  for prediction. Specifically, our method significantly reduces MRS for the weather, blur, and compression categories. Moreover, our robust variant R-XCNorm further diminishes the impact of  $R$  on the model’s predictions, particularly in the noise category, thereby enhancing both the model’s robustness and domain generalization performance. These findings highlight the effectiveness of our approach in mitigating the influence of domain-specific factors and improving the model’s sensitivity to the underlying semantics  $X_0$  rather than the rendering factor  $R$ .

## 6. Conclusion

In this paper, we introduce XCNorm, a novel normalization technique based on normalized cross-correlation. XC-Norm exhibits invariance to affine shifts and changes in energy within local feature patches, while also eliminating the need for non-linear activation functions. The robust variant, R-XCNorm, focuses on outlier rejection, resulting in improved performance in challenging domains while maintaining competitiveness in others. Our proposed masking method selectively normalizes important input regions, enhancing model stability and out-of-domain performance. The integration of both methods showcases their complementary nature, leading to further improvements across all domains. We demonstrate the practical applicability of XC-Norm in medical imaging classification, where it enhances model robustness and sensitivity to underlying semantics. Overall, our work provides effective methods for enhancing model performance and robustness, making notable contributions to the field of single-source domain generalization. These contributions have potential implications for various real-world applications.

## Acknowledgements

This research was partially funded by the Australian Government and Ford Motor Company through an Australian Research Council Linkage Project (LP190100165).



## References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 3
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010. 1
- [3] Moritz Böhle, Mario Fritz, and Bernt Schiele. B-cos networks: alignment is all we need for interpretability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10329–10338, 2022. 3
- [4] Péter Bándi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermesen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, Quanzheng Li, Farhad Ghazvinian Zanjani, Svitlana Zinger, Keisuke Fukuta, Daisuke Komura, Vlado Ovtcharov, Shenghua Cheng, Shaoqun Zeng, Jeppe Thagaard, Anders B. Dahl, Huangjing Lin, Hao Chen, Ludwig Jacobsson, Martin Hedlund, Melih Çetin, Eren Halıcı, Hunter Jackson, Richard Chen, Fabian Both, Jörg Franke, Heidi Küsters-Vandeveldel, Willem Vreuls, Peter Bult, Bram van Ginneken, Jeroen van der Laak, and Geert Litjens. 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, 2019. 5
- [5] Jin Chen, Zhi Gao, Xinxiao Wu, and Jiebo Luo. Meta-causal learning for single domain generalization. *arXiv preprint arXiv:2304.03709*, 2023. 2
- [6] John E Dennis Jr and Roy E Welsch. Techniques for nonlinear least squares and robust regression. *Communications in Statistics-simulation and Computation*, 7(4):345–359, 1978. 4
- [7] Susan J Devlin, Ramanathan Gnanadesikan, and Jon R Kettenring. Robust estimation and outlier detection with correlation coefficients. *Biometrika*, 62(3):531–545, 1975. 4
- [8] 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
- [9] Kamaleddin Ghiasi-Shirazi. Generalizing the convolution operator in convolutional neural networks. *Neural Processing Letters*, 50(3):2627–2646, 2019. 3
- [10] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019. 5
- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. 3
- [12] Yujie Jin, Xu Chu, Yasha Wang, and Wenwu Zhu. Domain generalization through the lens of angular invariance. 2022. 2
- [13] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton A. Earnshaw, Imran S. Haque, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning (ICML)*, 2021. 5
- [14] JP LEWIS. Fast template matching. *Proceeding of Vision Interface 95*, pages 120–123, 1995. 2, 4
- [15] Lei Li, Ke Gao, Juan Cao, Ziyao Huang, Yepeng Weng, Xiaoyue Mi, Zhengze Yu, Xiaoya Li, and Boyang Xia. Progressive domain expansion network for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 224–233, 2021. 2
- [16] Weiyang Liu, Zhen Liu, Zhiding Yu, Bo Dai, Rongmei Lin, Yisen Wang, James M Rehg, and Le Song. Decoupled networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2771–2779, 2018. 3, 5
- [17] Weiyang Liu, Yan-Ming Zhang, Xingguo Li, Zhiding Yu, Bo Dai, Tuo Zhao, and Le Song. Deep hyperspherical learning. *Advances in neural information processing systems*, 30, 2017. 3
- [18] Chunjie Luo, Jianfeng Zhan, Xiaohe Xue, Lei Wang, Rui Ren, and Qiang Yang. Cosine normalization: Using cosine similarity instead of dot product in neural networks. In *Artificial Neural Networks and Machine Learning—ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part I 27*, pages 382–391. Springer, 2018. 3
- [19] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020. 3
- [20] Cheng Ouyang, Chen Chen, Surui Li, Zeju Li, Chen Qin, Wenjia Bai, and Daniel Rueckert. Causality-inspired single-source domain generalization for medical image segmentation. *IEEE Transactions on Medical Imaging*, 2022. 2
- [21] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 2, 5
- [22] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29, 2016. 3
- [23] Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 68–83. Springer, 2020. 3
- [24] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring ro-

- bustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems*, 33:18583–18599, 2020. [1](#)
- [25] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. [3](#)
- [26] 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*, 31, 2018. [2](#), [5](#)
- [27] Chaoqun Wan, Xu Shen, Yonggang Zhang, Zhiheng Yin, Xinmei Tian, Feng Gao, Jianqiang Huang, and Xian-Sheng Hua. Meta convolutional neural networks for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4682–4691, 2022. [2](#), [5](#)
- [28] Chen Wang, Jianfei Yang, Lihua Xie, and Junsong Yuan. Kervolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 31–40, 2019. [3](#)
- [29] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh. Learning to diversify for single domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 834–843, 2021. [2](#), [5](#)
- [30] Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020. [2](#)
- [31] Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020. [1](#)
- [32] Skyler Wu, Fred Lu, Edward Raff, and James Holt. Exploring the sharpened cosine similarity. In *I Can’t Believe It’s Not Better Workshop: Understanding Deep Learning Through Empirical Falsification*, 2022. [5](#)
- [33] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. [3](#)
- [34] Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer. Robust and generalizable visual representation learning via random convolutions. In *International Conference on Learning Representations*, 2021. [2](#), [6](#), [8](#)
- [35] Long Zhao, Ting Liu, Xi Peng, and Dimitris Metaxas. Maximum-entropy adversarial data augmentation for improved generalization and robustness. *Advances in Neural Information Processing Systems*, 33:14435–14447, 2020. [2](#)
- [36] 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](#)
- [37] Georgios Zoumpourlis, Alexandros Doumanoglou, Nicholas Vretos, and Petros Daras. Non-linear convolution filters for cnn-based learning. In *Proceedings of the IEEE international conference on computer vision*, pages 4761–4769, 2017. [3](#)