

ConDiSR: Contrastive Disentanglement and Style Regularization for Single Domain Generalization

Aleksandr Matsun
MBZUAI

United Arab Emirates

aleksandr.matsun@mbzuai.ac.ae

Numan Saeed
MBZUAI

United Arab Emirates

numan.saeed@mbzuai.ac.ae

Fadillah Adamsyah Maani
MBZUAI

United Arab Emirates

fadillah.maani@mbzuai.ac.ae

Mohammad Yaqub
MBZUAI

United Arab Emirates

mohammad.yaqub@mbzuai.ac.ae

Abstract

Medical data often exhibits distribution shifts, leading to performance degradation of deep learning models trained using standard supervised learning pipelines. Domain Generalization (DG) addresses this challenge, with Single-Domain Generalization (SDG) being notably relevant due to the privacy and logistical constraints often inherent in medical data. Existing disentanglement-based SDG methods heavily rely on structural information from segmentation masks, but classification labels do not offer similarly dense information. This work introduces a novel SDG method for medical image classification, utilizing channel-wise contrastive disentanglement. The method is further refined with reconstruction-based style regularization to ensure distinct style and structural feature representations are extracted. We evaluate our method on the complex tasks of multicenter histopathology image classification and Diabetic Retinopathy (DR) grading in fundus images, benchmarking it against state-of-the-art (SOTA) SDG baselines. Our results demonstrate that our method consistently outperforms the SOTA independently on the choice of the source domain while exhibiting greater performance stability. This study underscores the importance and challenges of exploring SDG frameworks for classification tasks. The code is publicly available at <https://github.com/BioMedIA-MBZUAI/ConDiSR>

1. Introduction

In recent years, medical imaging analysis has experienced significant advancements powered by deep learning [19]. These models are typically trained under the assumption

of independent and identically distributed (i.i.d.) data samples. However, this assumption can significantly hamper performance during inference in real-world scenarios due to distribution shifts. This issue is particularly significant in medical imaging data due to differences between medical centers, variations in equipment, and the inherent complexity and variability of biological structures. Using large and diverse datasets during training to address these distribution shifts is essential. However, obtaining such datasets is challenging due to privacy concerns [27].

To address these constraints, the training process can be arranged in a Domain Adaptation (DA) [11] or Domain Generalization (DG) [5] framework. DA focuses on adapting a model trained on one or more source domains to perform well on a target domain, which may have a different distribution. On the other hand, DG aims to develop models that generalize well across unseen domains without having access to data from those (target) domains during training. These approaches assume that data can be divided into distinct domains, within which samples follow the i.i.d. assumption. Applying these techniques makes it possible to develop machine-learning models that exhibit improved performance and generalizability, even with the limited availability of heterogeneous data.

In the setting of Unsupervised DA (UDA) [21,25], training is performed using data from the source domain(s) as well as unlabeled data from the target domain. Conversely, the Multi-Source DG (MSDG) [12, 27] setting involves training using data from multiple source domains and evaluating on unseen target domain(s). However, obtaining data from multiple sources and target domains is often infeasible, especially in medical imaging.

On the other hand, Single Domain Generalization (SDG)

[8, 14, 15, 17, 21] uses data from a single source domain during training, better mimicking real-world scenarios. Compared to UDA and MSDG, SDG presents a more challenging task, as using data from one source increases the likelihood of model overfitting and reduces its ability to generalize to other domains. Methods that can help learn a generalizable model using only one source domain include broadening the input space via image augmentations [14, 17] or performing feature-level augmentations that avoid semantic perturbations of the content [13, 26, 28]. Recently, [8] proposed an SDG method that achieves domain-independent feature representation extraction by disentangling style and structure-related components based on the assumption that semantic information is primarily encoded in an image’s high-frequency components. In contrast, style/texture-related information resides in its low-frequency components [20].

While [8] shows SOTA results in the segmentation task, it performs poorly when adapted to classification. This is because the segmentation model training pipeline involves dense ground truth labels that provide significant structural information, aiding the model’s disentanglement module. In contrast, a standard classification model is only given one numeric value per image during training - the class label, which provides significantly less structural information. This issue is especially noticeable in classifying histopathology images as well as fundus images. For the former it is caused by significant style variations due to different staining procedures [23] while for the latter by high impact of scanning equipment settings.

In this work, we propose a novel method, ConDiSR, that combines **Contrastive Disentanglement and Style Regularization** for robust medical image classification in the SDG setting. Our key contributions are as follows:

- We develop a new SDG method that utilizes reconstruction-based style regularization for improved structure/style disentanglement.
- We develop a criterion suitable for the classification task that integrates reconstruction loss, contrastive loss, and cross-entropy loss functions.
- We propose a solution to the medical images classification task that improves over the previous SOTA methods in SDG on two DG datasets representing different medical imaging modalities.

2. Methodology

2.1. Preliminaries

Problem Definition: We first define $x \in \mathcal{X}$ as a medical image and $y \in \mathcal{Y}$ as its corresponding clinical label. Suppose that we have medical data from multiple domains or

institutions, denoted as $\mathcal{D} = \{\mathcal{D}^n\}_{n=1}^{N_{\mathcal{D}}}$, where each institution $\mathcal{D}^n = \{x_i^n, y_i^n\}_{i=1}^{N_{\mathcal{D}^n}}$ typically represents different distributions due to various factors such as different machines and demographics. Our main aim is to learn a function that can accurately map $\mathcal{X} \rightarrow \mathcal{Y}$ using the available data while being robust to distribution shifts.

Domain Generalization Setting: In real-world practice, medical institutions are often reluctant to share their data due to privacy concerns. By assuming that we have access to data from only some institutions, we can decompose \mathcal{D} into two sets:

$$\mathcal{D} = \mathcal{S} \cup \mathcal{T} \quad (1)$$

where \mathcal{S} represents data from source institutions $\mathcal{S} = \{\mathcal{S}^n\}_{n=1}^{N_{\mathcal{S}}}$ and \mathcal{T} represents data from target institutions $\mathcal{T} = \{\mathcal{T}^n\}_{n=1}^{N_{\mathcal{T}}}$. Consequently, we must rely on the source data to train a model that performs well on data from both seen and unseen distributions. Domain generalization aims to address this challenge.

Single Domain Generalization (SDG): This setting assumes that there is only a single domain available for training, i.e. $N_{\mathcal{S}} = 1$. Developing a robust model in SDG is challenging due to the limited diversity of a single training data distribution.

2.2. Proposed Method

In order to overcome the challenges that the classification task in SDG setting poses, we propose **Contrastive Disentanglement and Style Regularization (ConDiSR)**, a reconstruction-based style regularization in conjunction with the technique of channel-wise contrastive disentanglement. Our proposed method is illustrated in Figure 1. Our model takes a medical image $x \in \mathbb{R}^{3 \times W \times H}$ as input and produces two outputs: **1)** \hat{y} , a prediction of its true clinical label $y \in \{0, 1, \dots, C-1\}$, and **2)** a reconstructed image at a lower scale $\tilde{x} \in \mathbb{R}^{3 \times h \times w}$, used for reconstruction-based style regularization. Mathematically, the model can be represented as

$$(\hat{y}, \tilde{x}) = \mathcal{F}(x) \quad (2)$$

2.2.1 Model Design

We design our model \mathcal{F} to mitigate severe performance degradation caused by such distribution shifts. The model \mathcal{F} disentangles the features into style-rich features (f_{sty}) and structure-rich features (f_{str}). The style-rich features (f_{sty}) are used to reconstruct the input image, while the structure-rich features (f_{str}) are used for the classification task. This approach reduces the model’s reliance on image style that varies across different institutions or domains, thereby improving robustness to distribution shifts. The model \mathcal{F} comprises a *feature disentanglement module* ($\mathcal{F}_{\mathcal{D}, \mathcal{M}}$), a *classification network* (\mathcal{F}_C), and a *reconstruction network* (\mathcal{F}_R).

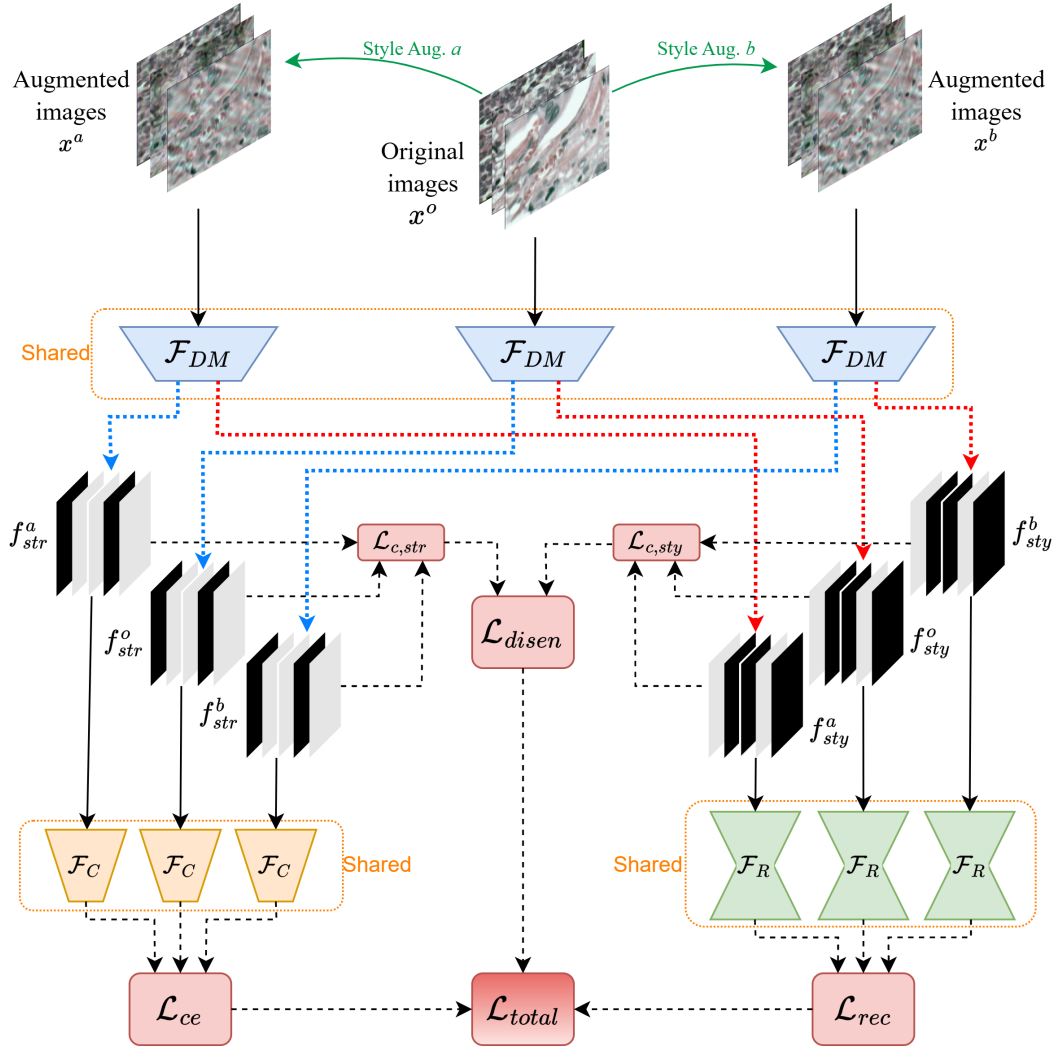


Figure 1. Overview of the proposed method. The feature disentanglement module (\mathcal{F}_{DM}) processes the original and augmented images to produce structure-related features (f_{str}) and style-related features (f_{sty}). The structure-related features (f_{str}) are passed further to the classification network (\mathcal{F}_C) for computing the classification loss, while the style-related ones (f_{sty}) go through the reconstruction network (\mathcal{F}_R) for further computation of the reconstruction loss. Additionally a contrastive loss is applied to minimize the distance between similar structure-related components and maximize the distance between different style-related components.

Feature Disentanglement Module: We design \mathcal{F}_{DM} to effectively provide us with structure-rich features f_{str} and style-rich features f_{sty} from an input image x . The image x is passed to the stem layer of the backbone convolutional network ($g = Conv2D(k = 7, s = 2, p = 3)$) with C_s output channels, to extract the shallow-level feature map, making the output more compact and enriched with distinct features. Let the output of the stem layer be denoted as $f_s \in \mathbb{R}^{C_s \times H/2 \times W/2}$. The channels of the feature map f_s are believed to carry distinct style- and structure-related features [22], so we perform the disentanglement by splitting these shallow level channels into structure- and style-related ones, thus comprising f_{str} and f_{sty} . We implement this by

introducing a learnable parameter $\theta_d \in \mathbb{R}^{2 \times C_s}$ that represents the model’s knowledge of which channels go to f_{str} and which to f_{sty} as shown on Figure 2. Each row of θ_d corresponds to one channel of f_d . In order to perform the disentanglement we separately apply softmax to each row of θ_d which gives us two weight values per channel of f_d . We use low temperature parameter in softmax to make sure that one of these weights closely approaches 0 while the other one - 1. By multiplying each of the channels with these weights, we facilitate the disentanglement of style and structure features. This approach allows the model to emphasize the contribution of each channel in providing style or structure information. Mathematically, the operations performed in

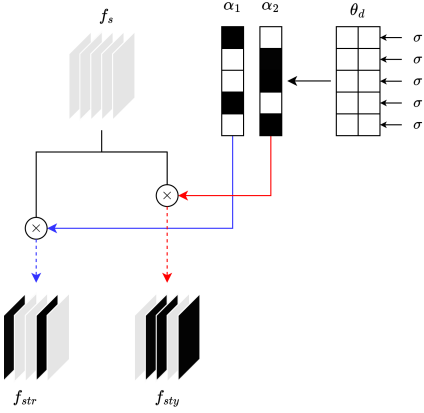


Figure 2. Overview of the feature disentanglement module of the network. Here σ represents the softmax operator, applied across the channel-related dimension of the parameter θ_d

$(f_{sty}, f_{str}) = \mathcal{F}_{DM}(x)$ can be written as follows:

$$f_s = g(x), \quad f_s \in \mathbb{R}^{C_s \times H/2 \times W/2} \quad (3)$$

$$\alpha_1, \alpha_2 = \sigma_c(\theta_d/\tau), \quad \alpha \in [0, 1]^{C_s} \quad (4)$$

$$f_{sty} = \alpha_1 * f_s, \quad f_{sty} \in \mathbb{R}^{C_s \times H/2 \times W/2} \quad (5)$$

$$f_{str} = \alpha_2 * f_s, \quad f_{str} \in \mathbb{R}^{C_s \times H/2 \times W/2} \quad (6)$$

where σ_c is the channel-wise softmax and $\tau \in \mathbb{R}$ is a temperature constant. By default, we set C_s to 64.

Classification Network: We harness the information embedded in structure-rich features f_{str} to perform image classification, i.e. $\hat{y} = \mathcal{F}_C(f_{str})$. These features encapsulate essential structural information that is robust to distribution shifts and also potentially contain less but sufficient style information to aid in classification. By leveraging these comprehensive features, our classifier \mathcal{F}_C is capable of effectively generalizing across different domains, thus maintaining high accuracy.

Reconstruction Network: In order to enhance performance of the disentanglement module and make sure that f_{sty} indeed carries style-related features we impose an additional constraint on f_{sty} : it must carry enough information to reconstruct the lower-resolution version of the input image. We denote the reconstructed image as $\tilde{x} = \mathcal{F}_R(f_{sty})$. It is crucial that downsampled version of the x is utilized as the reconstruction target in order to eliminate possible high-frequency (structural) features from it and thus avoid the possibility of f_{sty} carrying structural information that could be necessary for the reconstruction of full-scale image. The style augmentations are also applied to the images that are used as the reconstruction targets, since we want f_{sty} to carry the information that depends on the style-altering modifications.

2.2.2 Training

Contrastive Disentanglement: We utilize the concept of contrastive learning to train \mathcal{F}_{DM} to effectively disentangle f_{sty} and f_{str} from f_s . For every original image x , we generate two augmented images using Bezier curve transformation from SLAug [18] and low-frequency component replacement [24]. These two augmentation techniques alter the image style while preserving the original image structure. Let x^o denote the original image, and (x^a, x^b) denote the augmented images. We then pass these images to the feature disentanglement module $\mathcal{F}_{DM}(\cdot)$ to produce their corresponding style-rich and structure-rich features, i.e. (f_{sty}^o, f_{str}^o) , (f_{sty}^a, f_{str}^a) , and (f_{sty}^b, f_{str}^b) . The contrastive loss components are computed following the principle that, after disentanglement, f_{str}^o, f_{str}^a and f_{str}^b have to be the same as they represent the same structural pattern, while f_{sty}^o, f_{sty}^a and f_{sty}^b should be different due to the style alterations caused by augmentations. The contrastive disentanglement loss is then expressed as follows:

$$\mathcal{L}_{c, str} = \frac{1}{2} \cdot \sum_{\substack{i, j \in \{o, a, b\} \\ i \neq j}} \left| p(f_{str}^i) - p(f_{str}^j) \right| \quad (7)$$

$$\mathcal{L}_{c, sty} = -\frac{1}{2} \cdot \sum_{\substack{i, j \in \{o, a, b\} \\ i \neq j}} \left| p(f_{sty}^i) - p(f_{sty}^j) \right| \quad (8)$$

$$\mathcal{L}_{disen} = \mathcal{L}_{c, str} + \mathcal{L}_{c, sty} \quad (9)$$

where the $p : \mathbb{R}^{C_s \times H/2 \times W/2} \rightarrow \mathbb{R}^{C_c}$ is a learnable operator, performing projection of the feature maps into a latent space of smaller dimension C_c . The projector is a small network, that consists of a convolution followed by a fully-connected layer. By default C_c is set to 1024.

Style Regularization: The next step involves computation of the style regularization loss components, which have been specifically designed to impose additional constraints on the style-dependent features $(f_{sty}^o, f_{sty}^a, f_{sty}^b)$ and improve the overall disentanglement capability of the architecture. To serve as the ground truth for the reconstruction, we utilize a downscaled version of the original images, which typically preserve style-related information and lose their inherent structural features due to resizing. We apply a reconstruction loss to compare the reconstructed images $(\tilde{x}^o, \tilde{x}^a, \tilde{x}^b)$ with the downsampled images:

$$\mathcal{L}_{rec} = \sum_{i \in \{o, a, b\}} \|\tilde{x}^i - rs(x^i)\|^2, \quad (10)$$

where rs is the resizing operator.

Final Objective Function: Ultimately, the entire model \mathcal{F} is trained by utilizing the total loss:

$$\mathcal{L}_{total} = \mathcal{L}_{ce} + \mathcal{L}_{disen} + \mathcal{L}_{rec} \quad (11)$$

where \mathcal{L}_{ce} denotes the cross-entropy loss for supervised classification training.

3. Experimental Setup

Datasets. The main experiments are conducted using the binary classification dataset Camelyon17-WILDS, which is a patchified version of the original Camelyon17 [1] breast cancer whole-slide images dataset. The data was collected from five medical centers each representing a separate domain. All input images are sized 96×96 , and the labels indicate whether the central 32×32 area contains any tumor tissue. Additionally a composition of DR grading datasets Aptos [9], EyePACS [4], Messidor and Messidor2 [2] is used for testing the models’ generalizability.

As depicted in Figures 3 and 4, sample images from the dataset showcase significant variations in texture throughout the different domains, making it a challenging task to develop a generalizable model.

Evaluation Methods. We use the DomainBed [5] evaluation protocol with training-domain validation model selection strategy. In order to properly estimate the generalization ability of the models in the SDG setting, we iteratively repeat training with each of the centers being set as the source domain while the union of the remaining four serve as the target domain.

Networks architecture. As the main backbone all the models utilize ResNet18 [7] pre-trained on ImageNet [3]. Specifically the shallow feature map extraction operator g of the feature disentanglement module F_{DM} is implemented with the stem layer of ResNet18 while the classification network F_C is composed of its remaining layers. The reconstruction network F_R is implemented with a light-weight encoder-decoder model that involves 3 downsampling blocks and variable number of upsampling blocks depending on the reconstruction resolution.

Implementation Details. The experiments were run on 24GB Quadro RTX 6000 GPU. Each experiment was run for 50 epochs with the batch size of 256 using Adam optimizer [10] with the learning rate of $1e-3$. The temperature parameter τ was fixed at 0.1 similar to [8].

We apply the same data augmentation strategy that we use in our method to all the other methods for a fair comparison. By doing so we eliminate the possibility of data augmentations being the deciding factor in the model’s performance and equalize the total number images each model is fed during training.

Ablation studies. We conduct a study to investigate the impact of target resolution on the reconstructed image, thereby exploring the effects of varying downsampling factors on the original image. We hypothesize that a higher downsampling factor will result in the loss of firstly structure-related information from the original image and therefore stronger style dependency of the feature map uti-

Src	Method	C_0	C_1	C_2	C_3	C_4
C_0	ERM	–	91.7 \pm 0.1	89.8 \pm 1.0	92.3\pm0.9	77.0 \pm 2.0
	C ² SDG	–	91.8 \pm 0.1	90.0 \pm 0.8	92.1 \pm 2.1	76.6 \pm 2.1
	Ours	–	92.1\pm0.1	90.8\pm0.6	92.2 \pm 0.6	80.9\pm0.7
C_1	ERM	86.7 \pm 1.7	–	83.3 \pm 1.4	88.3\pm0.8	71.0 \pm 5.3
	C ² SDG	85.4 \pm 2.7	–	82.9 \pm 2.2	87.1 \pm 4.0	71.8 \pm 7.2
	Ours	87.2\pm0.3	–	84.3\pm0.5	88.1 \pm 0.8	81.6\pm1.3
C_2	ERM	94.2 \pm 0.1	81.1 \pm 1.0	–	87.3 \pm 0.3	93.4 \pm 0.1
	C ² SDG	93.9 \pm 0.3	80.5 \pm 0.5	–	86.6 \pm 1.3	91.8 \pm 1.4
	Ours	94.3\pm0.2	84.3\pm0.3	–	88.5\pm0.4	94.5\pm0.6
C_3	ERM	87.4 \pm 0.8	75.9 \pm 0.6	88.9 \pm 0.7	–	62.5 \pm 3.6
	C ² SDG	85.3 \pm 1.4	75.2 \pm 0.8	88.4 \pm 0.6	–	59.8 \pm 3.9
	Ours	88.6\pm0.4	79.1\pm0.8	89.9\pm0.5	–	71.2\pm0.7
C_4	ERM	90.0 \pm 0.6	84.4 \pm 0.7	87.8 \pm 0.6	91.4 \pm 0.6	–
	C ² SDG	90.7 \pm 1.2	84.9 \pm 0.9	88.3 \pm 1.1	89.9 \pm 2.6	–
	Ours	91.5\pm0.5	86.6\pm0.6	89.4\pm0.3	92.2\pm0.5	–

Table 1. Performance comparison on Camelyon17-WILDS in SDG setting. Each medical center number i is denoted as C_i . The column Src represents the source domain and the columns $C_0 - C_4$ - respective target domains. Each value shows accuracy (%) averaged over three runs with the standard deviation in the subscript

lized for reconstruction. However, we also acknowledge the possibility of a threshold beyond which excessively high downsampling factors may lead to the loss of crucial style-related information. In this work we check the reconstruction resolutions 96, 48 and 24.

Multi-Source Domain Generalization. In addition to the experiments carried out in the SDG setting we also test our method against the baselines in the MSDG setting. The training is performed in the leave-one-domain-out manner where each domain is iteratively set out as the target domain and the rest serve as the source domains. Due to the fact that all of the methods involved are originally designed for SDG, no domain labels are utilized during the training and all of the source domains are effectively merged into one. We conduct MSDG experiments on Camelyon17-WILDS and DR datasets.

Style augmentation. We also study the impact of combining our method with various feature-map style-augmentation techniques. In the scope of this work we evaluate such methods as MixStyle [28], Domain Shift with Uncertainty (DSU) [13] and Correlated Style Uncertainty (CSU) [26].

4. Results and Discussion

Table 1 compares performance of our method versus ERM and C²SDG [8] in SDG setting for each separate medical center being set as the source domain. Each value is achieved by averaging three experimental results with different seeds.

As it can be seen from the tables, our method consistently outperforms baselines by 1-2% on average. Moreover




Figure 3. Sample images from the domains of the five centers of Camelyon17-WILDS.




Figure 4. Sample images from the DR datasets.

Src	Method	Aptos	EyePACS	Messidor	Messidor2
Aptos	ERM	–	$65.2_{\pm 0.5}$	$46.3_{\pm 0.2}$	$58.2_{\pm 0.3}$
	C ² SDG	–	$63.3_{\pm 1.8}$	$46.1_{\pm 0.2}$	$58.4_{\pm 0.2}$
	Ours	–	$65.7_{\pm 0.6}$	$47.9_{\pm 0.3}$	$59.3_{\pm 0.4}$
EyePACS	ERM	$60.2_{\pm 0.4}$	–	$49.6_{\pm 0.3}$	$60.3_{\pm 0.4}$
	C ² SDG	$60.6_{\pm 0.7}$	–	$49.3_{\pm 0.7}$	$60.0_{\pm 0.4}$
	Ours	$61.9_{\pm 0.4}$	–	$51.2_{\pm 0.7}$	$60.9_{\pm 0.3}$
Messidor	ERM	$64.0_{\pm 0.7}$	$42.0_{\pm 2.8}$	–	$59.6_{\pm 1.5}$
	C ² SDG	$62.4_{\pm 1.8}$	$39.7_{\pm 4.2}$	–	$59.3_{\pm 0.4}$
	Ours	$64.0_{\pm 0.4}$	$44.6_{\pm 0.3}$	–	$59.8_{\pm 1.5}$
Messidor2	ERM	$62.2_{\pm 2.0}$	$38.9_{\pm 2.5}$	$62.1_{\pm 2.0}$	–
	C ² SDG	$64.2_{\pm 1.7}$	$41.8_{\pm 3.4}$	$64.0_{\pm 1.1}$	–
	Ours	$67.7_{\pm 1.0}$	$46.6_{\pm 1.3}$	$65.0_{\pm 1.0}$	–

Table 2. Performance comparison on DR datasets in SDG setting with the column Src representing source domain and the columns from 3rd to 6th representing target domains. Each value shows accuracy (%) averaged over three runs with the standard deviation in the subscript

	C_0	C_1	C_2	C_3	C_4	Average
ERM	$96.5_{\pm 0.2}$	$89.1_{\pm 1.2}$	$94.8_{\pm 0.5}$	$92.8_{\pm 0.4}$	$78.7_{\pm 2.3}$	<u>90.4</u>
C ² SDG	$96.3_{\pm 0.5}$	$89.4_{\pm 2.0}$	$94.3_{\pm 1.0}$	$90.3_{\pm 2.6}$	$79.7_{\pm 2.6}$	90.0
Ours	$97.3_{\pm 0.5}$	$92.1_{\pm 0.7}$	$96.1_{\pm 0.3}$	$94.2_{\pm 0.3}$	$84.1_{\pm 0.3}$	92.8

Table 3. Performance comparison in MSDG setting on Camelyon17-WILDS. Each column represents accuracy on domain C_i after training on the other 4 domains

it shows more stable performance with significantly smaller deviation across the runs. We argue that reconstruction-

	Aptos	EyePACS	Messidor	Messidor2	Average
ERM	$64.4_{\pm 1.4}$	$50.1_{\pm 0.9}$	$62.7_{\pm 1.8}$	$62.2_{\pm 0.5}$	<u>59.9</u>
C ² SDG	$60.2_{\pm 2.3}$	$48.5_{\pm 1.8}$	$61.7_{\pm 2.5}$	$60.3_{\pm 1.6}$	57.7
Ours	$66.1_{\pm 1.0}$	$51.8_{\pm 0.4}$	$63.0_{\pm 0.8}$	$63.3_{\pm 0.8}$	61.1

Table 4. Performance comparison in MSDG setting on the composition of DR datasets. Each column represents accuracy on the respective domain after training on the other 3 domains

based regularization aids the disentanglement module of the framework by enforcing stronger dependency between the input images and style-dependent channels of the shallow-level feature map. Due to direct interconnection this improves extraction of the structure-related channels as well.

The comparison between ERM and C²SDG shows that both methods achieve very close performance scores, which leads us to believe that when it comes to the binary classification task, the primary driver of good performance of C²SDG might be extensive augmentations as well as increased number of training iterations resulting from the augmentation strategy in [8] rather than the idea of channel-wise contrastive disentanglement in its default version.

The qualitative comparison of performance between C²SDG and ConDiSR is shown on the Figure 5. It is obtained by visualizing gradient activation maps [16] on samples with positive tumor label. Comparison with segmentation masks from original Camelyon17 challenge dataset shows that our method pays more attention to areas, containing semantically important information.

Table 5. Evaluation of the effect of changing the reconstruction resolution in ConDiSR is SDG setting. Each column represents performance after training on the domain C_i averaged across the remaining 4 domains.

	C_0	C_1	C_2	C_3	C_4	Average
96×96	86.8 \pm 0.9	83.0\pm0.8	88.9 \pm 1.1	73.4 \pm 0.4	88.6 \pm 0.5	84.1
48×48	87.4\pm0.6	81.8 \pm 1.0	90.4\pm0.7	75.0\pm0.8	89.7\pm0.6	84.9
24×24	83.4 \pm 0.7	81.3 \pm 0.3	89.1 \pm 0.3	73.1 \pm 0.3	87.5 \pm 1.5	82.9

Table 6. Investigation of the effect of combining various style augmentation techniques with ConDiSR. Each column represents performance after training on the domain C_i averaged across the remaining 4 domains.

	C_0	C_1	C_2	C_3	C_4	Average
MixStyle	85.2 \pm 0.4	80.9 \pm 0.8	92.1\pm0.1	73.7 \pm 0.5	91.1\pm0.3	84.6
+ Ours	87.2\pm0.1	82.8\pm0.5	92.1\pm0.6	75.6\pm0.6	90.6 \pm 0.5	85.7
DSU	85.0\pm0.7	81.4 \pm 0.8	91.1\pm0.2	76.4\pm1.0	90.2\pm0.7	84.8
+ Ours	84.5 \pm 0.7	84.2\pm0.7	90.7 \pm 0.4	75.8 \pm 0.8	89.5 \pm 0.6	84.9
CSU	86.0\pm0.5	79.6 \pm 2.1	90.7 \pm 0.2	76.4 \pm 0.6	90.4\pm0.4	84.6
+ Ours	86.0\pm0.4	85.7\pm0.5	90.8\pm0.7	78.1\pm0.6	87.7 \pm 0.9	85.7




Figure 5. Qualitative performance comparison between ConDiSR and C^2SDG [8] via Grad-CAM [16] with the tumor presence masks taken from the original Camelyon17 challenge dataset [1].

Table 2 compares performance of our method versus ERM and C^2SDG in SDG setting on the combination of the Diabetic Retinopathy datasets. Each of the datasets is one-by-one set as the source domain. As it can be seen from the tables, our method consistently outperforms the baselines while showing higher stability with lower deviation in results. The improvement is especially noticeable with Messidor2 as the source domain.

Multi-Source Domain Generalization. The results of

the experiments in MSDG setting are shown in the Tables 3 and 4. As it can be seen from the Table 3 when applied to Camelyon17-WILDS our method noticeably outperforms baselines on each of the target domains. The difference is especially noticeable on the data from the medical center 4, which overall is the hardest to generalize to for all of the methods. The average accuracy improvement is from 90.4% to 92.8% which corresponds to approximately 20% classification error reduction. When it comes to DR grading, our method again outperforms the baselines by the margin of more than 1%.

Reconstruction resolution. From the Table 5 we can see that the optimal reconstruction resolution for our method is 48×48 . This supports our preliminary assumption, that reconstruction of the input images of the original size (96×96) forces unnecessary structural features into f_{sty} and hampers the work of disentanglement module. On the other hand, reconstruction of the images resized to the resolutions that are too small (24×24) does not help the model’s performance. Similarly to [6], the gap seems to arise between the tasks of recognition and reconstruction. When the target image resolution is too low, the decoder of the reconstruction model is too shallow (as it consists of fewer upsampling blocks) which makes the input of the reconstruction network, f_{sty} , too specialized for reconstruction but not particularly informative from the point of view of recognition.

It is important to clarify, that Table 5 shows perfor-

mance in the SDG setting with each column representing the source domain and the accuracy being computed over the rest of the domains unlike it is presented in the Tables 1 and 2.

Style augmentation. As it can be seen in the Table 6, application of various style-augmentation methods [28] to the feature maps after each of the first three layers of the network does not provide any significant improvement compared to the baselines. However in conjunction with ConDiSR, it results in the highest accuracy reached, which raises an assumption that despite the usage of the disentanglement technique, structure-related feature maps extracted by our method still carry various important style-dependent features that are correlated with the classification label and benefit from augmentations. Similarly to Table 5, Table 6 shows performance in the SDG setting with each column representing the source domain and the accuracy being computed over the rest of the domains.

5. Conclusion

This paper presents a novel method, ConDiSR, designed for single domain generalization in the classification of images. Leveraging contrastive channel-wise disentanglement alongside style regularization through low-resolution reconstruction, ConDiSR demonstrates an improvement over baseline methods, achieving an increase in accuracy in both SDG and MSDG settings on various medical imaging modalities. Successful application of ConDiSR shows the potential of disentanglement-based methodologies in enhancing model performance and generalizability, particularly in classification tasks within complex and heterogeneous datasets such as histopathology images. Our results underscore the importance of further exploration of disentanglement-based approaches and encourage further research in this domain. Potential future work may involve evaluation of the proposed method's performance on other classification datasets in order to test its generalizability.

References

- [1] Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 2018. 5, 7
- [2] Etienne Decencière, Xiwei Zhang, Guy Cazuguel, Bruno Lay, Béatrice Cochener, Caroline Trone, Philippe Gain, Richard Ordonez, Pascale Massin, Ali Erginay, Béatrice Charton, and Jean-Claude Klein. Feedback on a publicly distributed image database: The messidor database. *Image Analysis & Stereology*, 0, 07 2014. 5
- [3] 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. 5
- [4] Emma Dugas, Jared, Jorge, and Will Cukierski. Diabetic retinopathy detection, 2015. 5
- [5] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization, 2020. 1, 5
- [6] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 7
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [8] Shishuai Hu, Zehui Liao, and Yong Xia. Devil is in channels: Contrastive single domain generalization for medical image segmentation. *arXiv preprint arXiv:2306.05254*, 2023. 2, 5, 6, 7
- [9] Sohier Dane Karthik, Maggie. Aptos 2019 blindness detection, 2019. 5
- [10] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [11] Wouter M. Kouw and Marco Loog. An introduction to domain adaptation and transfer learning, 2019. 1
- [12] Haoliang Li, YuFei Wang, Renjie Wan, Shiqi Wang, Tie-Qiang Li, and Alex Kot. Domain generalization for medical imaging classification with linear-dependency regularization. *Advances in neural information processing systems*, 33:3118–3129, 2020. 1
- [13] Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, and Ling-Yu Duan. Uncertainty modeling for out-of-distribution generalization. *arXiv preprint arXiv:2202.03958*, 2022. 2, 5
- [14] 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*, 42(4):1095–1106, 2022. 2
- [15] 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
- [16] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128(2):336–359, Oct. 2019. 6, 7
- [17] Zixian Su, Kai Yao, Xi Yang, Kaizhu Huang, Qiufeng Wang, and Jie Sun. Rethinking data augmentation for single-source domain generalization in medical image segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 2366–2374, 2023. 2
- [18] Zixian Su, Kai Yao, Xi Yang, Kaizhu Huang, Qiufeng Wang, and Jie Sun. Rethinking data augmentation for single-source

- domain generalization in medical image segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37:2366–2374, 06 2023. 4
- [19] Gaël Varoquaux and Veronika Cheplygina. Machine learning for medical imaging: methodological failures and recommendations for the future. *NPJ digital medicine*, 5(1):48, 2022. 1
- [20] Jingye Wang, Ruoyi Du, Dongliang Chang, Kongming Liang, and Zhanyu Ma. Domain generalization via frequency-domain-based feature disentanglement and interaction. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM '22. ACM, Oct. 2022. 2
- [21] 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, 2
- [22] Xiaozheng Xie, Jianwei Niu, Xuefeng Liu, Zhengsu Chen, Shaojie Tang, and Shui Yu. A survey on incorporating domain knowledge into deep learning for medical image analysis. *Medical Image Analysis*, 69:101985, Apr. 2021. 3
- [23] Chundan Xu, Ziqi Wen, Zhiwen Liu, and Chuyang Ye. Improved domain generalization for cell detection in histopathology images via test-time stain augmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 150–159. Springer, 2022. 2
- [24] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4084–4094, 2020. 4
- [25] Yifan Zhang, Ying Wei, Qingyao Wu, Peilin Zhao, Shuaicheng Niu, Junzhou Huang, and Mingkui Tan. Collaborative unsupervised domain adaptation for medical image diagnosis. *IEEE Transactions on Image Processing*, 29:7834–7844, 2020. 1
- [26] Zheyuan Zhang, Bin Wang, Debesh Jha, Ugur Demir, and Ulas Bagci. Domain generalization with correlated style uncertainty. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2000–2009, 2024. 2, 5
- [27] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–20, 2022. 1
- [28] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle, 2021. 2, 5, 8