

Complex Style Image Transformations for Domain Generalization in Medical Images

Nikolaos Spanos, Anastasios Arsenos, Paraskevi-Antonia Theofilou,
Paraskevi Tzouveli, Athanasios Voulodimos, Stefanos Kollias
National Technical University of Athens, Greece

{nspanos, partheofilou, anarsenos}@ails.ece.ntua.gr,
tpar@image.ece.ntua.gr, thanosv@mail.ntua.gr, stefanos@cs.ntua.gr

Abstract

The absence of well-structured large datasets in medical computer vision results in decreased performance of automated systems and, especially, of deep learning models. Domain generalization techniques aim to approach unknown domains from a single data source. In this paper we introduce a novel framework, named CompStyle, which leverages style transfer and adversarial training, along with high-level input complexity augmentation to effectively expand the domain space and address unknown distributions. State-of-the-art style transfer methods depend on the existence of sub-domains within the source dataset. However, this can lead to an inherent dataset bias in the image creation. Input-level augmentation can provide a solution to this problem by widening the domain space in the source dataset and boost performance on out-of-domain distributions. We provide results from experiments on semantic segmentation on prostate data and corruption robustness on cardiac data which demonstrate the effectiveness of our approach. Our method increases performance in both tasks, without added cost to training time or resources.

1. Introduction

Artificial intelligence and, more specifically, deep learning and neural networks, have been at the fore-

front of computer vision these past years. These models generally depend on large amounts of annotated data to be able to learn properly how to extract the required information. However, there are two main issues with sufficiently training deep neural networks. The first is that for several computer vision applications, such as medical image segmentation, available data is scarce or costly to acquire, thus leading to the creation of sub-optimally trained systems. The second is the domain shift problem [1]. Deep learning systems have been shown that despite their adequate performance on data similar to the source domain, performance drops drastically when tested on data from different distributions. The first problem exacerbates the second in a large degree, thus creating the need for advanced methods to help models generalize properly to unseen domains.

Single-source domain generalization aims to address this challenge by developing robust models capable of generalizing from a single source domain [2], [3], [4]. By learning representations that are invariant to domain shifts and effectively disentangling domain-specific variations from task-related features, these models hold promise for improving the adaptability and performance of computer vision systems across diverse domains.

Recent methods for medical image domain generalization, such as [5], propose that creating hard image examples from a single-source dataset can be

used to reliably approach domains that the model will under-perform and, thus, generalize better through failure. However, these methods requires a lot of time to train and resources, thus making it more difficult to optimize and the training method might suffer from dataset bias, since the generated images are based only on the single-source dataset.

Inspired by [5], we aim to raise performance by creating a composite data augmentation framework that leverages both feature-level space augmentation and input-level space augmentation. Through combined training and by leveraging the high-complexity input augmentation and style transfer techniques, we manage to increase the domain space so that the model can reliably generalize on unseen domain datasets.

The main contributions are:

- We propose a novel framework that uses adversarial data augmentation and style transfer for domain generalization, while optimizing model robustness and sufficient domain space generation through complex input-level image mixing.
- Our method can raise out-of-domain performance, without requiring additional training time or resources.
- We validate our framework’s performance on two different tasks for medical image segmentation and compare with the state-of-the-art.

2. Related Work

Deep Learning In Medical Analysis: In recent years, the field of deep learning in medical analysis has seen remarkable advancements, especially accelerated due to the COVID-19 pandemic [6], [7], [8], [9], [10], [11], [12], [13], [14], owing to the heightened global interest and the substantial accumulation of medical data. More recently, cutting-edge deep learning models are being leveraged to address a myriad of challenges in the medical domain, including image artifact correction [15], [16], [17] and text-guided applications [17], [18]. We mainly focus on alleviating the domain generalization problem in the medical field using models that can be used for practical applications in medical settings, as guidance or consultation tools for

doctors during clinical procedures.

Data Augmentation: Data augmentation is a classic method of augmenting the training dataset to help a model generalize better. Common methods include rotating, contrast changing and cropping input images [19], [20], [21], [22]. More advanced methods include mixing images of different labels or replacing parts of an image with another [23], [24]. More recent methods aim to appropriately enlarge the domain space by combining different augmentation techniques or by introducing more severe complexity in the input data, in the form of image mixing [25]. Despite their simplicity, they are still prominent in the field of computer vision and provide a useful way of alleviating low-data performance degradation and should always be considered when creating application systems.

Style Transfer: Style transfer is a prominent technique in computer vision where the style of one image is enforced on another, while preserving the semantic content [26], [27], [28]. State-of-the-art work mainly focuses on diffusion models and their ability to extract style information from text prompts or use probabilistic functions to create new styles [29], [30], [31]. There exists heavy correlation between domain shift and alterations in image style across various domains, which can be mitigated by augmenting the diversity of training image styles [32], [33]. Namely successful implementations are MixStyle [34], which generates ‘novel’ styles by linearly mixing style statistics from two arbitrary training instances from the same domain at feature level. DSU [35] perturbs feature styles with random noise to account for potential style shifts. MaxStyle [5] utilizes adversarial training to generate as many negative examples of images as possible to maximize the networks robustness through constant optimizing . We utilize the MaxStyle augmentation network, combined with input level augmentations, to approach a larger domain space and create more complex and varied images.

3. Methodology

The main objective of the data augmentation framework is to maximize robustness in a model M from a single-source domain D . During training, we want to create a wide domain space S so that we can approximate various distributions $\{S1, S2, \dots\}$ to simulate the domain shift. In order to achieve that, we create an augmentation pipeline that utilizes input and feature complexity to create new domains.

3.1. Image Mixing

We hypothesize that for enriching the domain space generated by the augmentation, not only do we require feature-level perturbations, but also input-level. Feature-level perturbations are highly reliant on the dataset used, so, despite the fact that

the model is trained so that the domain space is widely generalized, it is still restricted by the input domain. Introducing images of high enough complexity during training can help not only widen the domain space generated, but also retroactively help the feature-level augmentation network create more complex images [25]. For the image mixing process we use fractal images and standard augmentation techniques. It has been shown in extended research that utilizing complex, non-natural images for image mixing has been beneficial for improving model performance [36], [37]. We can also create fractal images using various algorithms, meaning that its relatively easy to enrich the mixing set at any time.

For our framework, we randomly mix an input batch b with fractal images, along with randomly augmenting the input batch with different, traditional augmentations t . So, a batch after mixing is characterized by: $b^+ = mix(t(b))$, where mix can be additive or multiplicative functions between the fractal image and the image slice. We empirically found through our experiments that mixing intensity should not be too high, so that semantic content can be preserved and that the augmentation should not be applied across the entire training and only through controlled settings. Otherwise, the adversarial training will lead to the creation of highly corrupted images and the backpropagation algorithm will echo the effect. We present some image examples in Figure 1.

3.2. Style Transfer Network

For the style transfer network we use the method referred to in [5]. The style transfer network uses an auxiliary image decoder to create adversarial samples from the latent representations of the source domain. Specifically, given feature f_i extracted at a certain CNN layer in the image decoder D_{ϕ_i} with x_i as input, the network augments f_i via:

$$\begin{aligned} \text{Net}(f_i) = & (\gamma_{\text{mix}} + \Sigma_{\gamma} \cdot \epsilon_{\gamma}) \odot f_i \\ & + (\beta_{\text{mix}} + \Sigma_{\beta} \cdot \epsilon_{\beta}), \end{aligned} \quad (1)$$

where \odot represents element-wise multiplication. In this equation, f_i undergoes augmentation with mixed styles $\gamma_{\text{mix}}, \beta_{\text{mix}}$, accompanied by additional




Figure 1. Examples of original and augmented images produced by the pipeline on cardiac data.

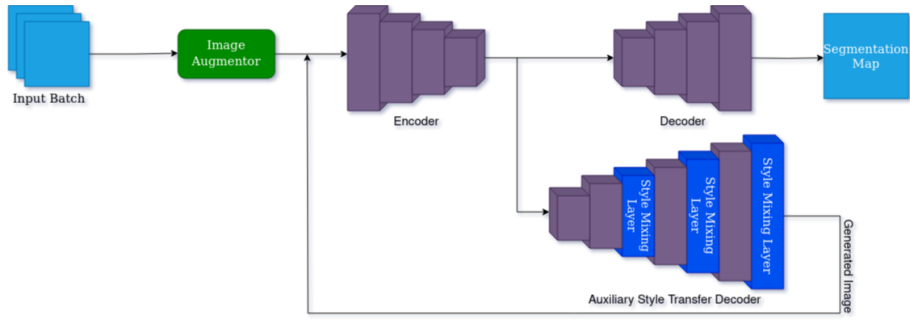


Figure 2. The complete network architecture. Images are augmented before they enter the encoder portion of the network and the latent representations are used to create the adversarial images. Since the entire batch is mixed, the fractal images are mixed together through style mixing in the added decoder to create more varied styles during generation. Adversarial images are fed back to the training loop, while closing the auxiliary decoder, to enhance training.

style noise $\Sigma_\gamma \cdot \epsilon_\gamma, \Sigma_\beta \cdot \epsilon_\beta \in \mathbb{R}^c$. These noise terms are introduced to facilitate exploration of potential unknown domain shifts. Style noises are sampled from a re-scaled Gaussian distribution with variances $\Sigma_\gamma, \Sigma_\beta \in \mathbb{R}^c$, estimated from style statistics of a batch of B instances (including x_i):

$$\Sigma_\gamma = \sigma^2 (\{\sigma(f_j)\}_{j=1,i,\dots,B}), \quad (2)$$

$$\Sigma_\beta = \sigma^2 (\{\mu(f_j)\}_{j=1,i,\dots,B}), \quad (3)$$

where $\epsilon_\gamma, \epsilon_\beta \sim \mathcal{N}(0, 1)$.

Because the style transfer network depends on the existence of sub-domains inside of a single-source domain and their subsequent mixing [34], this may lead to an unintended bias to the source domain dataset. Enhancing the images with high-complexity augmentation can help create more mixed domains during training at an input level, thus enhancing the model’s generalization.

3.3. Network Architecture

We present our complete network architecture in Figure 2. For our experiments, we utilized two different Fully Convolutional Networks (FCNs). The architecture features an encoder-decoder structure, wherein the encoder section comprises multiple convolutional layers responsible for hierarchical feature extraction from input images. These layers progressively capture both low-level and high-level representations, culminating in a condensed,

lower-dimensional representation of the input features. The decoder section, on the other hand, reconstructs the spatial information from the encoded features. This allows for the generation of pixel-wise segmentation maps corresponding to the input images.

The style transfer network functions as an additional decoder that receives the latent representations as input. The noise augmentation layers are placed within the decoder layers to facilitate the style transfer process. We utilize two architectures: an FCN-16 and an FCN-64 [38] for prostate data image segmentation. This investigation aims to determine the potential benefits of input augmentation techniques on the complexity of the model. Additionally, we perform evaluation on the model’s robustness against image corruptions, specifically focusing on the FCN-16 architecture applied to cardiac data.

4. Experiments and Results

4.1. Datasets

We evaluate the effect of out-of-domain performance on prostate data and corruption robustness on cardiac data.




Figure 3. *Artifacted cardiac images. From left to right: Bias Artifact (showing as difference in brightness on different images areas), Ghosting Artifact (showing as idols due to movement of the organ or patient), Motion Artifact (showing as displacement), Spike Artifact (showing as stripes due to electromagnetic interference)*

4.1.1 Prostate Data

For our experiments on prostate data, we utilize the publicly available prostate MRI segmentation dataset from the Medical Decathlon Challenge, which includes, in total, 148 patients. This includes the following 7 different datasets:

NCI-ISBI-2013: Two different datasets from the 2013 competition, conducted by the National Cancer Institute (NCI) in collaboration with the International Society for Biomedical Imaging (ISBI). Half of the images were acquired on a 1.5T (Philips Achieva) with an endo-rectal coil (from Boston Medical Center), and the other half on a 3T (Siemens TIM) with a surface coil (from Radboud University Medical Center, Nijmegen, the Netherlands) [39], [40], [41]. Thus, the images are divided according to their origin for testing on out-of-domain performance. The datasets notated on the results tables are **A**, **B**.

I2CVB: A dataset from the Initiative for Collaborative Computer Vision Benchmarking (I2CVB) [39], [40], [42]. The images were taken by a 3T scanner Siemens, with the following methods for better representation: Magnetic Resonance Imaging with T2 weighting (T2-W), Dynamic Contrast-Enhanced Magnetic Resonance Imaging (DCE), Diffusion-Weighted Magnetic Resonance Imaging (DWI), Magnetic Resonance Spectroscopic Imaging (MRSI). The dataset notated on the results tables is **C**.

PROMISE12: Three datasets from different medical centers and different acquisition methods from the PROMISE12 competition for prostate

medical image segmentation [39], [40], [43]. The datasets notated on the results tables are **D**, **E** and **F**.

Medical Decathlon Dataset: A new dataset provided by the challenge competition. This specific dataset was used for training and validation in within-distribution applications [44]. The dataset notated on the results tables is **G**.

4.1.2 Cardiac Data

For our experiments we evaluate corruption robustness on the publicly available dataset the M&M (Multi-Centre, Multi-Vendor, and Multi-Disease) cardiac segmentation challenge dataset [45], consisting of a total of 100 heart MRI scans. Of the three-dimensional images, we make use of the slices corresponding to end-diastole and end-systole periods. The dataset is multi-labelled, consisting of annotations for left ventricle, myocardium and right ventricle. We evaluate on the following 4 types of corruptions created by TorchIO software [46]:

- **Motion:** Random "motion", which usually appears physiologically in MRI images due to organ motion.
- **Spike:** Random spike artifacts, also known as Herringbone artifacts, create stripes in different directions in the image space due to spikes generated in the machine's electromagnetic field.
- **Ghosting:** Random displacement of an image ghost, usually due to cardiac motion, patient motion during the examination, or blood flow.

Methods	OOD							IID
	A	B	C	D	E	F	Mean	G
Baseline	87.9	75.5	77.7	76.8	29.5	24.2	58.6	83.8
RandConv	89	75.7	<u>82.3</u>	85.1	41.4	52.7	71.03	<u>86.8</u>
AdvNoise	87.3	66.4	76.5	72.8	44.6	50.3	66.3	82.1
AdvBias	90.2	84.3	81.1	87.2	39.7	51.6	72.35	87.4
MixStyle	87.5	69.3	80.0	73.4	43.1	40.8	65.86	84.5
DSU	88.2	54.4	81.5	84.1	<u>58.6</u>	64.0	71.8	83.1
MaxStyle	<u>89.7</u>	81.9	82.2	83.4	52.8	71.5	<u>76.92</u>	85.5
CompStyle (Ours)	89.6	<u>83.9</u>	82.4	<u>85.4</u>	64.6	<u>64.5</u>	78.4	86.4

Table 1. Comparison of methods for domain generalization on prostate data with FCN-16 architecture

Methods	OOD							IID
	A	B	C	D	E	F	Mean	G
Baseline	90.4	67.9	79.3	79.1	45.7	60.9	70.5	83.8
RandConv	90.7	80.7	77.1	<u>88.0</u>	27.5	42.2	67.7	87.4
AdvNoise	90.0	76.0	81.2	80.5	41.2	47.6	69.42	86.9
AdvBias	91.3	<u>81.6</u>	82.3	87.2	45.1	57.5	<u>74.0</u>	89.3
MixStyle	<u>91.1</u>	68.9	82.6	82.4	43.8	55.4	70.7	85.6
DSU	90.4	63.6	79.9	78.2	<u>57.2</u>	<u>61.2</u>	71.75	87.3
MaxStyle	90.5	80.4	<u>84.5</u>	86.0	45.8	56.0	73.87	<u>88.0</u>
CompStyle (Ours)	90.7	86.2	85.7	88.5	75.2	77.2	83.92	86.4

Table 2. Comparison of methods for domain generalization on prostate data with FCN-64 architecture

Methods	Corruption											
	RandBias			RandSpike			RandMotion			RandGhosting		
	LV	MYO	RV	LV	MYO	RV	LV	MYO	RV	LV	MYO	RV
Baseline	91.9	85.5	83.2	1	1	1	91.6	83.6	82.8	91.1	83.7	81.6
MaxStyle	91.5	84.9	83.2	57.9	48.4	44.9	91.8	84.2	83.9	91.0	83.6	82.7
CompStyle (Ours)	91.9	84.8	83.4	62.5	49.5	51.0	91.6	83.9	83.0	90.9	83.4	80.3

Table 3. Comparison table of methods for domain generalization on Cardiac Data with FCN-16 Architecture

- Bias Field: Random fluctuations in image intensity usually due to the in-homogeneity of the MRI machine’s field.

Examples of the images are shown in Figure 3.

4.2. Implementation and experiment set-up

We follow the same training schema and dataset preprocessing as in [5]. For the style transfer layers of the dual-branch network, we raise the number

of iterations required for optimizing the adversarial images. For the input-level augmentation, we lower the intensity from the default parameters to ensure preservation of semantic content in the input images. Each batch in training is either mixed by the augmentation with a probability of 0.4 or noise is added to prevent overfitting. We evaluate the composite augmentation frameworks on two different architectures, FCN-16 and FCN-64. We retain the Adam optimizer and multi-task loss from [5] and train for 600 epochs with batch size 20 on the FCN-16 architecture for both tasks and for 600 epochs with batch size 15 on the FCN-64 architecture. All experiments were conducted on an NVIDIA V100 GPU and an NVIDIA GeForce 4060Ti GPU. We compare the augmentation frameworks with competitive methods such as: RandConv [47], adversarial noise [48], adversarial bias field [49], MixStyle [34], DSU [35] and MaxStyle [5]. We use the DICE score metric for evaluation, averaged on the entirety of each dataset [50]. All experiments were re-done locally due to an error in the evaluation of [5], which was noted in their code implementation.

4.3. Results

Results for the prostate segmentation task for the FCN-16 and FCN-64 architectures can be viewed on tables 1 and 2 respectively. We present the best result for evaluation from our framework. Input-level methods generally achieved better intra-domain performance than the style-transfer feature-level methods, especially AdvBias. However, both MaxStyle and DSU outperformed the rest on the **E** and **F** datasets. We can see from the results that the composite augmentation framework achieves the best out-of-domain performance in both architectures. Increasing the amount of features within the architecture help boost performance for all datasets on the composite framework, while the other methods suffered heavily from the added complexity. Intra-domain performance for our framework remains competitive with the rest of the methods, albeit with a small drop in scoring. The small drop may be impacted by the large number of unknown domains introduced during training. Other meth-

ods can be more favorable towards the source domain, thus appearing in evaluation as performing better. However, favoring the source domain means that out-of-domain performance will severely drop, which is detrimental towards developing practical systems. In general, the added input complexity combined with a larger network helped create a more stable out-of-domain performance on all datasets.

We present results for corruption robustness on table 3. We compare with MaxStyle since it showed the best performance on the FCN-16 architecture. Of all four artifacts, RandSpike appeared to be the most detrimental towards the model's performance, which is in-line with how much it influences the entire image. The composite augmentation helped boost performance towards the most detrimental artifact and remained competitive on all labels with the rest of the artifacts.

5. Conclusion

We introduced a new augmentation framework for domain generalization from a single-source domain. It combines complex image augmentations and a style transfer adversarial network to enrich the domain space during the training process. By creating images of diverse and hard styles, we raise network robustness towards the domain shift problem and assist in boosting the model's generalization. Higher complexity networks can better utilize the added input complexity, in comparison with networks that have fewer features in their layers. We also highlight the importance of tuning the augmentation in a way that the mixing process does not alter the semantic content, especially for medical images where it is vital that the correct information is fed to the model. We validated the effect of the augmentation on two different tasks, for prostate segmentation and cardiac data corruption robustness. From our experiments, we conclude that certain methods favor certain datasets. Degradation can be task-dependent or even dataset-dependent, but utilizing enough image complexity during training can help leverage an even increase for all out-of-domain datasets. Our method works without a need for more resources or more train-

ing time, proving the importance of standard input augmentation techniques when tackling model generalization. For future work, it would be interesting to see the effect of combining other techniques, such as methods of adding noise in the encoder portion of the network, or these methods validated on more recent image segmentation networks, such as vision transformers, or different tasks, such as autonomous driving or agricultural applications.

6. Acknowledgements

This research work is Co-funded from the European Union’s Horizon Europe Research and Innovation programme under Grant Agreement No 101119714 — dAIry 4.0.

References

- [1] M. Zhang, H. Marklund, N. Dhawan, A. Gupta, S. Levine, and C. Finn, “Adaptive risk minimization: Learning to adapt to domain shift,” 2021. [1](#)
- [2] C. Ouyang, C. Chen, S. Li, Z. Li, C. Qin, W. Bai, and D. Rueckert, “Causality-inspired single-source domain generalization for medical image segmentation,” *IEEE Transactions on Medical Imaging*, vol. 42, no. 4, pp. 1095–1106, 2022. [1](#)
- [3] A. Arsenos, D. Kollias, E. Petrongonas, C. Skliros, and S. Kollias, “Uncertainty-guided contrastive learning for single source domain generalisation,” 2024. [1](#)
- [4] S. Lee, H. Seong, S. Lee, and E. Kim, “Wildnet: Learning domain generalized semantic segmentation from the wild,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9936–9946, June 2022. [1](#)
- [5] C. Chen, Z. Li, C. Ouyang, M. Sinclair, W. Bai, and D. Rueckert, “Maxstyle: Adversarial style composition for robust medical image segmentation,” 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [6] D. Kollias, A. Arsenos, and S. Kollias, “Ai-enabled analysis of 3-d ct scans for diagnosis of covid-19 & its severity,” in *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, p. 1–5, IEEE, 2023. [2](#)
- [7] A. Arsenos, A. Davidhi, D. Kollias, P. Prassopoulos, and S. Kollias, “Data-driven covid-19 detection through medical imaging,” in *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, p. 1–5, IEEE, 2023. [2](#)
- [8] D. Kollias, A. Arsenos, and S. Kollias, “A deep neural architecture for harmonizing 3-d input data analysis and decision making in medical imaging,” *Neurocomputing*, vol. 542, p. 126244, 2023. [2](#)
- [9] D. Kollias, A. Arsenos, and S. Kollias, “Ai-mia: Covid-19 detection and severity analysis through medical imaging,” in *European Conference on Computer Vision*, p. 677–690, Springer, 2022. [2](#)
- [10] A. Arsenos, D. Kollias, and S. Kollias, “A large imaging database and novel deep neural architecture for covid-19 diagnosis,” in *2022 IEEE 14th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP)*, p. 1–5, IEEE, 2022. [2](#)
- [11] D. Kollias, A. Arsenos, L. Soukissian, and S. Kollias, “Mia-cov19d: Covid-19 detection through 3-d chest ct image analysis,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, p. 537–544, 2021. [2](#)
- [12] D. Kollias, N. Bouas, Y. Vlaxos, V. Brillakis, M. Seferis, I. Kollia, L. Sukissian, J. Wingate, and S. Kollias, “Deep transparent prediction through latent representation analysis,” *arXiv preprint arXiv:2009.07044*, 2020. [2](#)
- [13] D. Kollias, Y. Vlaxos, M. Seferis, I. Kollia, L. Sukissian, J. Wingate, and S. D. Kollias, “Transparent adaptation in deep medical image diagnosis,” in *TAILOR*, p. 251–267, 2020. [2](#)
- [14] D. Gerogiannis, A. Arsenos, D. Kollias, D. Nikitopoulos, and S. Kollias, “Covid-19 computer-aided diagnosis through ai-assisted ct imaging analysis: Deploying a medical ai system,” 2024. [2](#)
- [15] W. Xie and M. B. Blaschko, “Dense transformer based enhanced coding network for unsupervised metal artifact reduction,” 2023. [2](#)
- [16] A. Jiang, C. Huang, Q. Cao, S. Wu, Z. Zeng, K. Chen, Y. Zhang, and Y. Wang, “Multi-scale cross-restoration framework for electrocardiogram anomaly detection,” 2023. [2](#)
- [17] J. Jing, T. Wang, H. Yu, Z. Lu, and Y. Zhang, “Interslice consistency for unpaired low-dose ct denoising using boosted contrastive learning,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023* (H. Greenspan, A. Madabhushi, P. Mousavi, S. Salcudean, J. Duncan, T. Syeda-Mahmood, and R. Taylor, eds.), (Cham), pp. 238–247, Springer Nature Switzerland, 2023. [2](#)
- [18] S. Zhang, J. Zhang, Y. Xie, and Y. Xia, “Tpro: Text-prompting-based weakly supervised histopathol-

- ogy tissue segmentation,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023* (H. Greenspan, A. Madabhushi, P. Mousavi, S. Salcudean, J. Duncan, T. Syeda-Mahmood, and R. Taylor, eds.), (Cham), pp. 109–118, Springer Nature Switzerland, 2023. [2](#)
- [19] L. Zhang, X. Wang, D. Yang, T. Sanford, S. Harmon, B. Turkbey, B. J. Wood, H. Roth, A. Myronenko, D. Xu, and Z. Xu, “Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation,” *IEEE Transactions on Medical Imaging*, vol. 39, no. 7, pp. 2531–2540, 2020. [2](#)
- [20] C. Chen, W. Bai, R. H. Davies, A. N. Bhuva, C. H. Manisty, J. B. Augusto, J. C. Moon, N. Aung, A. M. Lee, M. M. Sanghvi, K. Fung, J. M. Paiva, S. E. Petersen, E. Lukaschuk, S. K. Piechnik, S. Neubauer, and D. Rueckert, “Improving the generalizability of convolutional neural network-based segmentation on cmr images,” *Frontiers in Cardiovascular Medicine*, vol. 7, June 2020. [2](#)
- [21] S. Otálora, M. Atzori, V. Andrearczyk, A. R. Khan, and H. Müller, “Staining invariant features for improving generalization of deep convolutional neural networks in computational pathology,” *Frontiers in Bioengineering and Biotechnology*, vol. 7, 2019. [2](#)
- [22] R. Volpi and V. Murino, “Addressing model vulnerability to distributional shifts over image transformation sets,” 2019. [2](#)
- [23] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” 2018. [2](#)
- [24] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, “Cutmix: Regularization strategy to train strong classifiers with localizable features,” 2019. [2](#)
- [25] D. Hendrycks, A. Zou, M. Mazeika, L. Tang, B. Li, D. Song, and J. Steinhardt, “Pixmix: Dream-like pictures comprehensively improve safety measures,” 2022. [2](#), [3](#)
- [26] L. A. Gatys, A. S. Ecker, and M. Bethge, “A neural algorithm of artistic style,” 2015. [2](#)
- [27] T. Park, A. A. Efros, R. Zhang, and J.-Y. Zhu, “Contrastive learning for unpaired image-to-image translation,” 2020. [2](#)
- [28] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” 2019. [2](#)
- [29] Y. Zhang, N. Huang, F. Tang, H. Huang, C. Ma, W. Dong, and C. Xu, “Inversion-based style transfer with diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10146–10156, June 2023. [2](#)
- [30] Z. Wang, L. Zhao, and W. Xing, “StyLEDif-fusion: Controllable disentangled style transfer via diffusion models,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7677–7689, October 2023. [2](#)
- [31] H. Lu, H. Tunanyan, K. Wang, S. Navasardyan, Z. Wang, and H. Shi, “Specialist diffusion: Plug-and-play sample-efficient fine-tuning of text-to-image diffusion models to learn any unseen style,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14267–14276, June 2023. [2](#)
- [32] V. Dumoulin, J. Shlens, and M. Kudlur, “A learned representation for artistic style,” 2017. [2](#)
- [33] X. Huang and S. Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” 2017. [2](#)
- [34] K. Zhou, Y. Yang, Y. Qiao, and T. Xiang, “Domain generalization with mixstyle,” 2021. [2](#), [4](#), [7](#)
- [35] X. Li, Y. Dai, Y. Ge, J. Liu, Y. Shan, and L.-Y. Duan, “Uncertainty modeling for out-of-distribution generalization,” 2022. [2](#), [7](#)
- [36] K. Nakashima, H. Kataoka, A. Matsumoto, K. Iwata, and N. Inoue, “Can vision transformers learn without natural images?,” 2021. [3](#)
- [37] H. Kataoka, K. Okayasu, A. Matsumoto, E. Yamagata, R. Yamada, N. Inoue, A. Nakamura, and Y. Satoh, “Pre-training without natural images,” 2021. [3](#)
- [38] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” 2015. [4](#)
- [39] Q. Liu *et al.*, “Ms-net: Multi-site network for improving prostate segmentation with heterogeneous mri data,” *IEEE Transactions on Medical Imaging (TMI)*, 2020. [5](#)
- [40] Q. Liu, Q. Dou, and P. A. Heng, “Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2020. [5](#)
- [41] “NCI-ISBI 2013 Challenge - Automated Segmentation of Prostate Structures - The Cancer Imaging Archive (TCIA) Public Access - Cancer Imaging Archive Wiki — wiki.cancerimagingarchive.net.” <https://wiki.cancerimagingarchive.net>.

net / display / public / nci - isbi + 2013 + challenge+ - +automated + segmentation + of + prostate + structures. [Access datee 15-09-2023]. 5

- [42] G. Lemaitre, R. Marti, J. Freixenet, J. C. Vilanova, P. M. Walker, and F. Meriaudeau, "Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric mri: A review," *Computers in Biology and Medicine*, vol. 60, pp. 8–31, 2015. 5
- [43] "PROMISE12 - Grand Challenge — promise12.grand-challenge.org." [https : // promise12 . grand - challenge . org](https://promise12.grand-challenge.org). [Access Date 15-09-2023]. 5
- [44] M. Antonelli *et al.*, "The medical segmentation de-cathlon," *arXiv:2106.05735*, 2021. 5
- [45] V. M. Campello *et al.*, "Multi-centre, multi-vendor and multi-disease cardiac segmentation: The m&ms challenge," *IEEE Transactions on Medical Imaging*, 2021. 5
- [46] F. Pérez-García, R. Sparks, and S. Ourselin, "Torchio: A python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning," *Computer Methods and Programs in Biomedicine*, vol. 208, p. 106236, Sept. 2021. 5
- [47] Z. Xu, D. Liu, J. Yang, C. Raffel, and M. Niethammer, "Robust and generalizable visual representation learning via random convolutions," 2021. 7
- [48] T. Miyato, S. ichi Maeda, M. Koyama, and S. Ishii, "Virtual adversarial training: A regularization method for supervised and semi-supervised learning," 2018. 7
- [49] C. Chen, C. Qin, H. Qiu, C. Ouyang, S. Wang, L. Chen, G. Tarroni, W. Bai, and D. Rueckert, "Realistic adversarial data augmentation for mr image segmentation," 2020. 7
- [50] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," 2015. 7