



UNIVERSITY OF LEEDS

This is a repository copy of *Style Curriculum Learning for Robust Medical Image Segmentation*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/187259/>

Version: Accepted Version

Proceedings Paper:

Liu, Z, Manh, V, Yang, X et al. (6 more authors) (2021) Style Curriculum Learning for Robust Medical Image Segmentation. In: MICCAI 2021: Medical Image Computing and Computer Assisted Intervention. MICCAI 2021, 27 Sep - 01 Oct 2021, Strasbourg, France. Springer Cham , pp. 451-460. ISBN 9783030871925

https://doi.org/10.1007/978-3-030-87193-2_43

© 2021 Springer Nature Switzerland AG. This is an author produced version of a conference paper, published in MICCAI 2021: Medical Image Computing and Computer Assisted Intervention. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Style Curriculum Learning for Robust Medical Image Segmentation

Zhendong Liu^{1,2,3*}, Van Manh^{1,2,3*}, Xin Yang^{1,2,3}, Xiaoqiong Huang^{1,2,3},
Karim Lekadir⁴, Víctor Campello⁴, Nishant Ravikumar^{5,6}, Alejandro F
Frangi^{1,5,6,7}, Dong Ni^{1,2,3(✉)}

¹National-Regional Key Technology Engineering Laboratory for Medical Ultrasound, School of Biomedical Engineering, Health Science Center, Shenzhen University, China
nidong@szu.edu.cn

²Medical Ultrasound Image Computing (MUSIC) Lab, Shenzhen University, China

³Marshall Laboratory of Biomedical Engineering, Shenzhen University, China

⁴Artificial Intelligence in Medicine Lab (BCN-AIM), Dept. de Matemàtiques i Informàtica, Universitat de Barcelona, Spain

⁵Centre for Computational Imaging and Simulation Technologies in Biomedicine (CISTIB), School of Computing and School of Medicine, University of Leeds, UK

⁶Leeds Institute of Cardiovascular and Metabolic Medicine, University of Leeds, UK

⁷Medical Imaging Research Center (MIRC), KU Leuven, Leuven, Belgium

Abstract. The performance of deep segmentation models often degrades due to distribution shifts in image intensities between the training and test data sets. This is particularly pronounced in multi-centre studies involving data acquired using multi-vendor scanners, with variations in acquisition protocols. It is challenging to address this degradation because the shift is often not known *a priori* and hence difficult to model. We propose a novel framework to ensure robust segmentation in the presence of such distribution shifts. Our contribution is three-fold. First, inspired by the spirit of curriculum learning, we design a novel style curriculum to train the segmentation models using an easy-to-hard mode. A style transfer model with style fusion is employed to generate the curriculum samples. Gradually focusing on complex and adversarial style samples can significantly boost the robustness of the models. Second, instead of subjectively defining the curriculum complexity, we adopt an automated gradient manipulation method to control the hard and adversarial sample generation process. Third, we propose the Local Gradient Sign strategy to aggregate the gradient locally and stabilise training during gradient manipulation. The proposed framework can generalise to unknown distribution without using any target data. Extensive experiments on the public M&Ms Challenge dataset demonstrate that our proposed framework can generalise deep models well to unknown distributions and achieve significant improvements in segmentation accuracy.

Keywords: Image segmentation · Style transfer · Curriculum learning

* Zhendong Liu and Van Manh contribute equally to this work.

1 Introduction

Recent studies have witnessed the great success of deep models in medical image segmentation [11]. However, these deep models often suffer from a drop in performance when applied to new data distributions, different from the training data (see Fig. 1). It is expensive and practically impossible to collect a large amount of manually annotated data from each new distribution to retrain the model in clinical practice. Therefore, a general and retraining-free framework that ensures model robustness to distribution shifts is highly desired in the clinic. To date, several approaches have been proposed to address this issue.

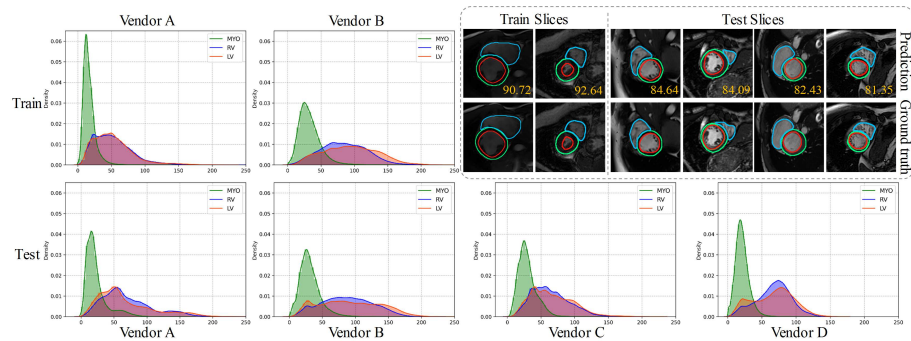


Fig. 1. Remarkable distribution shift between the training set and test set across different vendors. The results of the performance drop are shown in the upper right corner. The yellow digits denote the average DSC of the MYO (green), LV (red) and RV (blue).

Data Augmentation. Data augmentation (DA) using spatial/intensity transformations is commonly employed to expand the distribution of training data and improve model generalisation. Zhang et al. [19] proposed *Mixup*, where, two inputs and their corresponding labels are proportionally interpolated to augment the training data. Similarly, Yun et al. [18] cut and paste local patches from one image into another. Although DA methods can mitigate model overfitting to some degree, they cannot guarantee the ability of deep models to generalise to multi-centre, multi-vendor data, typically encountered in real clinical scenarios.

Domain Adaptation. GAN-based domain adaptation methods often focus on learning domain-invariant representations [4,8] or aligning the feature space of different distributions [17,21] and have shown promise for dealing with variations in data distributions. Curriculum-based domain adaptation methods typically define curriculum complexity subjectively, such as using the average distance of the domain feature space [13] and developing a ‘simpler’ task than semantic segmentation [20]. As these methods rely on the availability of unlabelled target data, they are restricted to the task of domain-mapping/adaptation. Hence, they may not generalise well to new ‘unseen’ data distributions.

Distribution Generalisation. Compared to the domain adaptation setting, distribution generalisation methods tested on unseen data distributions, are better suited to address the challenges encountered in real clinical scenarios. Some studies employ style transfer methods to remove the distribution shift in test data [12,14]. Although these approaches are novel, the selection of style data is subjective. Other studies have utilised adversarial training to augment the training set [2,16]. Forcing networks to learn from the most difficult samples can enhance model robustness to a certain degree. Still, it also makes the training more difficult to optimise, even causing catastrophic overfitting [9].

This work proposes a retraining-free and general framework to improve the ability of segmentation models to generalise to unknown distributions. Our contribution is three-fold. *First*, inspired by the advantages of curriculum learning in improving model generalisation, we design a novel style curriculum, structured in an easy-to-hard adversarial learning strategy, to train the model. Gradually forcing the model to learn from progressively harder adversarial samples allows the model to generalise well to new distributions. Specifically, a style transfer model with style fusion is adopted to generate samples for the curriculum. *Second*, instead of defining the complexity order subjectively, we employ a novel gradient manipulation method to automatically control the sample generation process following the increasingly harder direction. *Third*, since the gradient manipulation-based training is difficult to optimise, we propose Local Gradient Sign (LGS), which locally aggregates the gradient to increase the complexity of the generated samples gradually, thus making the training more stable. Extensive experiments on public data from the Multi-Centre, Multi-Vendor and Multi-Disease Cardiac Segmentation (M&Ms) Challenge [3] show that our method outperforms all 2D methods, and is on par with the top performing approaches of the challenge.

2 Methodology

Our framework is shown in Fig. 2. First, we adopt a style transfer model with a style fusion operation to generate the curriculum samples z_i . Then, we employ the gradient manipulation method to update the learning weight Γ , so the curriculum samples are arranged from easy to hard. To stabilise the training, we further propose Local Gradient Sign to operate the gradient locally.

2.1 Curriculum Learning for Robustness

Deep neural networks are typically trained using a sequence of unordered samples. Curriculum learning methods [1,6,20] can guide the model to learn better by organising the training samples in a meaningful order. A common curriculum needs to address three challenges: (1) Decide the curriculum samples. (2) Arrange the samples in an easy-to-hard order. (3) Ensure model stability during training. Using curriculum learning, deep models can leverage information learned from easy examples, to ease learning of new and harder samples. Gradually concentrating on learning the harder tasks can make the deep model more

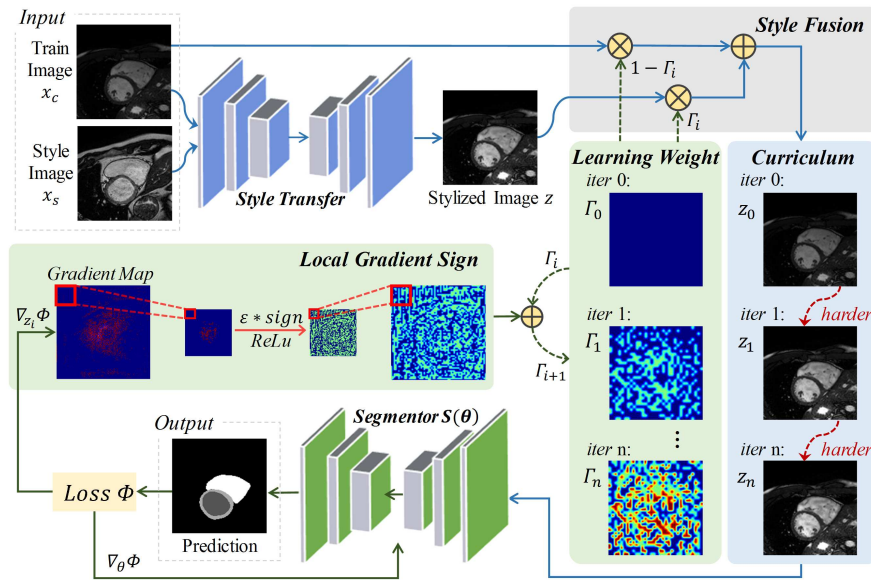


Fig. 2. Schematic view of our proposed framework.

robust. Creating an effective curriculum is critical to design a reasonable samples generator, a meaningful learning strategy, and a stable learning method.

2.2 Style Transfer based Sample Generation

Many works [7,10] argued that the images' style could be described as the distribution of colours, edges, smoothness, etc. Since style transfer (ST) methods are powerful tools to render one image's style onto another, we employ this approach to produce large sets of stylised images. This expands the training set and increases the richness of the information available for training the segmentation model. ST is only an optional module of our system to generate initial samples.

A reasonable samples generator of the curriculum requires the ST method to balance transfer quality and efficiency. Therefore, we follow WaveCT-AIN [12] and re-implement their high-quality ST. To better control the stylised degree, we propose a style fusion operation to modulate the content of training images and the style of stylised images. The style fusion weight is updated during back-propagation to control the degree of stylisation.

2.3 Gradient Manipulation based Learning Strategy

ST is a powerful technique that can expand style curriculum samples, but the random transformation lacks control for generating samples in a progressive, easy-to-hard fashion. Following the work from Goodfellow et al. [5], we adopt a novel gradient manipulation adversarial strategy to phase curriculum learning in

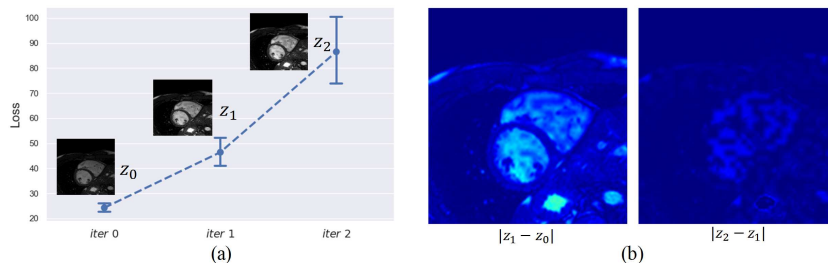


Fig. 3. (a) The loss statistics of curriculum samples at different iterations in training. (b) The absolute difference of samples z_i between iter 0 and 1, iter 1 and 2.

increasing order of complexity. They proposed the Fast Gradient Sign Method to generate adversarial examples through the following formula.

$$x' = x + \epsilon \text{sign}(\nabla_x J(f_\theta(x), y)) \quad (1)$$

Where θ are the parameters of the model, x is the input of the model, y is the label of x , ϵ is a small perturbation vector, x' is the generated adversarial samples, and $J(f_\theta(x), y)$ is the loss function used to train the model. The method adjusts the input data by a small step (ϵ) in the gradient direction of the cost function that maximises the loss. With this adversarial approach, the method can quickly generate more difficult samples for the deep models. Similarly, we iteratively accumulate the gradient of the cost function to produce increasingly harder adversarial samples (see step 7 in **Algorithm 1**). Moreover, in contrast with the typical gradient manipulation adversarial methods [5,9,16], which add the perturbation into the samples directly, we use it as the learning weight to modulate the training image and stylised image. This ensures control over the adversarial samples generated using our approach (see step 5 in **Algorithm 1**). To verify the effectiveness of our method, we visualise the loss of three curriculum samples iteratively generated by each training sample (Fig.3 (a)) and the image differences of the adjacent curriculum samples (we perform pseudo-color processing in Fig.3 (b) for better visualisation). The increase of the training loss shows that the method generates harder samples as the iterations progress.

2.4 Local Gradient Smoothing for Stability

Although the proposed gradient manipulation method guarantees the curriculum’s learning order, it is not easy to optimise the model by introducing the whole adversarial perturbations. Therefore, we further propose Local-Gradient-Sign (LGS), to reduce the influence of adversarial perturbations on the model. The LGS formula is defined as follows.

$$LGS(\Phi, \epsilon, size) = US(ReLU(\epsilon * \text{sign}(AP(\nabla_{z_i} \Phi, size)))), \quad i = 0, 1, \dots, n \quad (2)$$

Here, $US(\cdot)$ is an UpSample operation, $AP(\cdot)$ is an AvgPooling operation, z_i is curriculum sample and ϵ is a decimal from 0 to 1. In our curriculum setting, we set the max learning iter n to 3, learning step ϵ to 0.25 and the pooling size to 4×4 . The LGS operation first uses the AvgPooling operation to average the gradient matrix locally for smoothing the perturbations and then performs partial truncation using the ReLU function. Finally, it uses the UpSample operation to restore the original size of the gradient matrix. The proposed LGS operation can alleviate the perturbations to make the training more stable. In our design, the ST and LGS operations are only required during training for data augmentation and attack mitigation of hard samples, respectively. During testing, the trained model is lightweight and efficient without using these two modules.

Algorithm 1 Local Style Curriculum Learning (LSCL)

Notation: Segmentation model $f(\cdot)$; Style transfer model $S(\cdot)$; max iter n ; learning rate α ; segmentation model parameters θ ; full zeros matrix O and full ones matrix I
Input: Content data $X_c \in \mathbb{R}^{b \times h \times w}$ and label $Y_c \in \mathbb{R}^{b \times h \times w}$; style data $X_s \in \mathbb{R}^{b \times h \times w}$
Output: Optimal θ^*

- 1: **for** $(x_c, y_c) \in (X_c, Y_c)$, $x_s \in X_s$ **do**
- 2: Initialization: $z = S(x_c, x_s)$, $\Gamma_0 = O \in \mathbb{R}^{h \times w}$
- 3: **for** $i = 0, \dots, n$ **do**
- 4: Update stylised learning sample: $z_i = \Gamma_i * z + (I - \Gamma_i) * x_c$
- 5: Compute loss function: $\Phi = J(f(z_i, \theta), y_c)$
- 6: Accumulate learning weight: $\Gamma_{i+1} = \Gamma_i + LGS(\nabla_{z_i} \Phi)$
- 7: Update model parameters: $\theta \leftarrow \theta - \alpha \nabla_{\theta} \Phi$
- 8: **return** θ as θ^*

3 Experimental Results

Our experimental data came from the M&Ms Challenge. This challenge cohort had 375 patients scanned in clinical centres using four different magnetic resonance scanner vendors (A, B, C and D). The training set contained 150 images from two different vendors (75 each of vendor A and B), in which only the end-diastolic (ED) and end-systolic (ES) phases are annotated. The images have been segmented by experienced clinicians from the respective institutions, including contours for the left (LV), right ventricle (RV) blood pools, and the left ventricular myocardium (MYO). All experiments are evaluated on 50 new studies from vendor A, B, C and 50 additional studies from the unseen vendor D.

We totally obtained 3,284 training slices from the annotated ED & ES phases and 10,607 style slices cross the short-axis view of vendor A and B. We first used the training slices to train the segmentation model [15], employing Adam optimiser (learning rate of $10^{-3} \sim 10^{-5}$). Based on the well-trained segmentation model and pre-trained style transfer model, we randomly sampled the slices of vendors A, B as content-style input. Then we adopted the LSCL algorithm to

finetune the segmentation model for 10 epochs using SGD-momentum optimiser (learning rate of 10^{-5}). To further improve the model robustness, we finally adopted the Test-time Augmentation (TTA), including three rotation operations (90° , 180° , 270°) to aggregate the multiple predictions. We used a segmentation loss comprising cross-entropy loss (weight of 0.6) and dice loss (weight of 0.4).

3.1 Result Details

We performed the comparative experiment and ablation experiment using indicators of Dice similarity coefficient (DSC), Jaccard index (JAC), Hausdorff distance (HD, [mm]) and Average symmetric surface distance (ASSD, [mm]). In all experiments, we adopted the same evaluation criteria and ranking method as the M&Ms Challenge performs [3], including the *Min-max Score*. Due to the space limitations, we computed the averaged metrics over the LV, RV and MYO.

Table 1. The mean (std) results of DSC and HD on different vendors of all patients.

Methods	Vendor A		Vendor B		Vendor C		Vendor D		DSC Score	HD Score	Min-max Score
	DSC \uparrow	HD \downarrow	DSC \uparrow	HD \downarrow	DSC \uparrow	HD \downarrow	DSC \uparrow	HD \downarrow			
U-net	0.858 (0.034)	17.113 (13.218)	0.835 (0.056)	13.392 (7.541)	0.822 (0.067)	17.616 (13.544)	0.835 (0.045)	17.153 (14.328)	0.8345	16.674	0.000
Mixup	0.872 (0.034)	10.441 (3.315)	0.859 (0.047)	11.483 (3.605)	0.837 (0.065)	12.596 (6.010)	0.837 (0.040)	11.904 (4.506)	0.8465	11.821	0.528
P3	0.878 (0.042)	12.587 (12.279)	0.887 (0.053)	9.872 (3.607)	0.865 (0.055)	10.461 (6.499)	0.864 (0.050)	13.969 (16.845)	0.8705	11.887	0.779
LSCL	0.877 (0.033)	9.888 (3.042)	0.865 (0.051)	10.746 (3.315)	0.857 (0.063)	11.224 (5.445)	0.866 (0.036)	11.479 (5.027)	0.8647	11.007	0.789
P2	0.884 (0.039)	12.461 (12.237)	0.892 (0.051)	9.796 (3.369)	0.870 (0.046)	9.548 (3.294)	0.867 (0.049)	13.432 (14.845)	0.8750	11.370	0.869
LSCL-TTA	0.884 (0.032)	9.802 (3.529)	0.876 (0.047)	10.238 (2.979)	0.865 (0.059)	10.592 (5.192)	0.868 (0.034)	11.195 (4.890)	0.8710	10.602	0.890
P1	0.889 (0.042)	12.072 (12.641)	0.893 (0.046)	9.482 (3.343)	0.876 (0.042)	9.465 (3.562)	0.877 (0.042)	13.091 (14.838)	0.8813	11.111	0.958

Table 2. Training and inference time of different methods.


Method	Training time	Inference time	GPU Device
LSCL-TTA	5 h	0.2 s	GTX 1080 Ti
P1	60 h	$\approx 1s$	Titan XP
P2	48 h	4.8 s	Tesla V100
P3	4-5 days	N/A	TITAN V100

Method Comparison. The quantitative comparisons among our methods and others on M&Ms data are shown in Tab.1. U-net is the baseline and Mixup is a common data augmentation method. P3-P1 are the top three methods on the leaderboard of the M&Ms Challenge. Tab.1 shows our method helps the U-net

Table 3. Ablation experiment on ED and ES volumes.

Methods	ED				ES			
	DSC \uparrow	JAC \uparrow	HD \downarrow	ASSD \downarrow	DSC \uparrow	JAC \uparrow	HD \downarrow	ASSD \downarrow
WaveCT-AIN	0.847	0.747	11.469	1.146	0.827	0.713	12.386	1.413
SCL	0.861	0.766	10.743	1.038	0.843	0.736	11.042	1.159
LSCL	0.876	0.787	10.543	0.920	0.853	0.751	11.459	1.146
LSCL-TTA	0.882	0.797	10.246	0.870	0.861	0.762	10.898	1.077

significantly improve the performance by 3.65% and 6.072 mm in terms of the DSC and HD metrics. Smaller standard deviations show that our method is more robust and stable against distribution shifts. Tab.1 and Tab. 2 show that our proposed method (LSCL-TTA) achieves competitive performance among the top ranking methods, while presents the lowest computational costs.

**Fig. 4.** Visualisation of ablation methods at different vendors.

Ablation Experiment. To thoroughly evaluate our proposed framework, we conducted two ablation experiments. The first experiment uses only the style transfer method *WaveCT-AIN* [12] to generate random style transformations and then feeds the stylised outputs to train the segmentation model. The second experiment does not introduce the LGS method, which we name SCL. As shown in Tab. 3, the WaveCT-AIN shows the worst performance since the model learns the stylised samples arbitrarily, causing the model to be trapped in local minima. SCL, instead, can make the model learn increasingly harder tasks. We further proposed LSCL to stabilise the adversarial training. Fig.4 visualises the segmen-

tation results of test samples corresponding to different ablation methods. The proposed LSCL-TTA shows the best segmentation results on each vendor.

4 Conclusion

This paper proposes a novel style curriculum learning framework to ensure segmentation models are robust against the distribution shift. Extensive experiments show that our proposed framework can significantly improve the generalisation. The proposed framework is universal and retraining-free, which makes it compelling for use in clinical practice.

Acknowledgement. This work was supported by the National Key R&D Program of China (No. 2019YFC0118300), Shenzhen Peacock Plan (No. KQTD20160-53112051497, KQJSCX20180328095606003), Royal Academy of Engineering under the RAEng Chair in Emerging Technologies (CiET1919/19) scheme, EP-SRC TUSCA (EP/V04799X/1), the Royal Society CROSSLINK Exchange Programme (IES/NSFC/201380), European Union’s Horizon 2020 research and innovation program under grant agreement number 825903 (euCanSHare project), Spanish Ministry of Science, Innovation and Universities under grant agreement RTI2018-099898-B-I00.

References

1. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the 26th annual international conference on machine learning. pp. 41–48 (2009)
2. Cai, Q.Z., Liu, C., Song, D.: Curriculum adversarial training. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence. pp. 3740–3747 (2018)
3. Campello, V.M., Gkontra, P., Izquierdo, C., Martín-Isla, C., Sojoudi, A., Full, P.M., Maier-Hein, K., Zhang, Y., He, Z., Ma, J., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: The m&ms challenge. *IEEE Transactions on Medical Imaging* (2021)
4. Chen, C., Ouyang, C., Tarroni, G., Schlemper, J., Qiu, H., Bai, W., Rueckert, D.: Unsupervised multi-modal style transfer for cardiac mr segmentation. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 209–219. Springer (2019)
5. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014)
6. Hacoheh, G., Weinshall, D.: On the power of curriculum learning in training deep networks. In: International Conference on Machine Learning. pp. 2535–2544. PMLR (2019)
7. Huang, E., Gupta, S.: Style is a distribution of features. arXiv preprint arXiv:2007.13010 (2020)
8. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proceedings of the European conference on computer vision (ECCV). pp. 172–189 (2018)

9. Li, B., Wang, S., Jana, S., Carin, L.: Towards understanding fast adversarial training. arXiv preprint arXiv:2006.03089 (2020)
10. Li, Y., Wang, N., Liu, J., Hou, X.: Demystifying neural style transfer. arXiv preprint arXiv:1701.01036 (2017)
11. Liu, S., Wang, Y., Yang, X., Lei, B., Liu, L., Li, S.X., Ni, D., Wang, T.: Deep learning in medical ultrasound analysis: a review. *Engineering* **5**(2), 261–275 (2019)
12. Liu, Z., Yang, X., Gao, R., Liu, S., Dou, H., He, S., Huang, Y., Huang, Y., Luo, H., Zhang, Y., et al.: Remove appearance shift for ultrasound image segmentation via fast and universal style transfer. In: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). pp. 1824–1828. IEEE (2020)
13. Liu, Z., Miao, Z., Pan, X., Zhan, X., Lin, D., Yu, S.X., Gong, B.: Open compound domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12406–12415 (2020)
14. Ma, C., Ji, Z., Gao, M.: Neural style transfer improves 3d cardiovascular mr image segmentation on inconsistent data. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 128–136. Springer (2019)
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
16. Volpi, R., Namkoong, H., Sener, O., Duchi, J., Murino, V., Savarese, S.: Generalizing to unseen domains via adversarial data augmentation. arXiv preprint arXiv:1805.12018 (2018)
17. Yan, W., Wang, Y., Gu, S., Huang, L., Yan, F., Xia, L., Tao, Q.: The domain shift problem of medical image segmentation and vendor-adaptation by unet-gan. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 623–631. Springer (2019)
18. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6023–6032 (2019)
19. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
20. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2020–2030 (2017)
21. Zhang, Z., Yang, L., Zheng, Y.: Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In: Proceedings of the IEEE Conference On Computer Vision and Pattern Recognition. pp. 9242–9251 (2018)