Mixing-AdaSIN: Constructing a De-biased Dataset using Adaptive Structural Instance Normalization and Texture Mixing

Myeongkyun Kang[1], Philip Chikontwe[1], Miguel Luna[1], Kyung Soo Hong[2], June Hong Ahn[2]*, and Sang Hyun Park[1]*

[1] Robotics Engineering, Daegu Gyeongbuk Institute of Science and Technology (DGIST), Daegu, Korea
{mkkang,shpark13135}@dgist.ac.kr
[2] Division of Pulmonology and Allergy, Department of Internal Medicine, Regional Center for Respiratory Diseases, Yeungnam University Medical Center, College of Medicine, Yeungnam University, Daegu, Korea
fireajh@yu.ac.kr

Abstract. Following the pandemic outbreak, several works have proposed to diagnose COVID-19 with deep learning in computed tomography (CT); reporting performance on-par with experts. However, models trained/tested on the same in-distribution data may rely on the inherent data biases for successful prediction, failing to generalize on out-of-distribution samples or CT with different scanning protocols. Early attempts have partly addressed bias-mitigation and generalization through augmentation or re-sampling, but are still limited by collection costs and the difficulty of quantifying bias in medical images. In this work, we propose Mixing-AdaSIN; a bias mitigation method that uses a generative model to generate de-biased images by *mixing* texture information between different labeled CT scans with semantically similar features. Here, we use Adaptive Structural Instance Normalization (AdaSIN) to enhance de-biasing generation quality and guarantee structural consistency. Following, a classifier trained with the generated images learns to correctly predict the label without bias and generalizes better. To demonstrate the efficacy of our method, we construct a biased COVID-19 vs. bacterial pneumonia dataset based on CT protocols and compare with existing state-of-the-art de-biasing methods. Our experiments show that classifiers trained with de-biased generated images report improved in-distribution performance and generalization on an external COVID-19 dataset.

1 Introduction

Recently, several methods have been proposed to improve COVID-19 patient diagnosis or treatment planning using CT scans [12,16,24,25]. The most methods are often trained and evaluated on single source data; producing models that

---

* Corresponding author.

exploit underlying biases in the data for better predictions. Yet, they fail to generalize when the bias shifts in external data reporting lower performance. This has critical implications, especially in medical imaging where biases are hard to define or accurately quantify. To address this, extensive data augmentation [5] or re-sampling is often employed; though is still limited by collection (multi-institute data) and how to express the bias when it is unknown. Thus, there is a need for methods that mitigate bias in part or fully towards improved model performance and better generalization.

In general, models trained on biased data achieve high accuracy and despite their capacity lack the motivation to learn the complexity of the intended task. For instance, a model trained on our in-house biased dataset with COVID-19 and bacterial pneumonia reported 97.18% (f1-score) on the validation set, yet degrades to 33.55% when evaluated on an unbiased test-set. Here, we believe bias may originate from varied CT protocols based on exam purpose, scanners, and contrast delivery requirements [3]. Though contrast CT is a standard protocol, it is challenging for practitioners to meet requirements during the pandemic due to extra processes such as contrast agent injection and disinfection [21,11]. Further, protocols may also vary for other pneumonia that exhibit similar imaging characteristics with COVID-19 in CT. Consequently, we believe biased datasets are often constructed unexpectedly and sometimes unavoidably due to the aforementioned factors.

Among the existing techniques proposed to remove model dependence on bias, augmentation is a de-facto technique for medical images; with other methods pre-defining the bias the trained model should be independent of. This assumes bias is easily defined, but one has to take extra care in the medical setup where such assumptions do not hold. To address this, we propose to construct a de-biased dataset where spurious features based on texture information become uninformative for accurate prediction. A key motivation is that accurate prediction of COVID-19 from other pneumonia's is dependent on the CT protocols related to texture features and contrast. Thus, we propose to generate COVID-19 CTs with bacterial pneumonia protocol characteristics and vice versa for bacterial pneumonia with COVID-19, respectively.

Specifically, we propose Mixing-AdaSIN; a generative model based bias removal framework that leverages Adaptive Structural Instance Normalization (AdaSIN) and texture *mixing* to generate de-biased images used to train classifiers robust to the original bias. For image generation, we employ two main components: (a) texture *mixing*, which enables realistic image generation, and (b) AdaSIN, which guarantees structural consistency and prevents bias retainment in the input image via modifying the distribution of the structure feature. To prevent incorrect image generation, we first pre-train a contrastive encoder [8] to learn key CT features and later use it to search similar image pairs for the texture *mixing* step in the proposed generative framework. For evaluation, we construct biased train/validation sets based on the CT protocol and an unbiased test set from the non-overlapping CT protocols of the train/validation sets, respectively. The proposed method reports high bias mitigation performance

(66.97% to 80.97%) and shows improved generalization performance when verified on an external dataset. The main contributions of this work are summarized as follows:

- We propose a generative model that can sufficiently mitigate the bias present in the training dataset using AdaSIN and texture *mixing*.
- We show that the use of contrastive learning for texture transfer pair selection prevents incorrect image generation.
- We constructed a biased COVID-19 vs. bacterial pneumonia dataset to verify bias mitigation performance. Our approach not only enabled improvements for standard classification models such as ResNet18 but also current state-of-the-art COVID-19 classification models.
- We also demonstrate the generalization performance of our classifier trained with the de-biased data on an external public dataset.

2 Related Works

CT based COVID-19 Classification. Several methods have been proposed to address this task since the inception of the pandemic [12,16,24,25]. For instance, Li et al. [16] encode CT slices using a 2D CNN and aggregate slice predictions via max-pooling to obtain patient-level diagnosis. Wang et al. [24] proposed COVID-Net, an approach that utilizes the long-range connectivity among slices to increase diagnostic performance. Later, Wang et al. [25] further improve COVID-Net by employing batch normalization and contrastive learning to make the model robust to multi-center training. However, these models did not address the bias in the dataset and thus may fail to generalize on other datasets.

Bias Mitigation. To mitigate bias, Alvi et al. [1] employed a bias classifier and a confusion loss to regularize the extracted features to make them indistinguishable from a bias classifier. Kim et al. [15] proposed to mitigate bias through mutual information minimization and the use of a gradient reversal layer. Though these methods can mitigate distinct biases such as color, they fail to mitigate bias in the medical domain since the bias from CT protocols is subtle and hard to distinguish even for humans [2].

Another line of work is the augmentation based models that utilize techniques such as arbitrary style transfer. Geirhos et al. [5] proposed shape-ResNet, an approach that finetunes a model pre-trained on AdaIN [10] generated images. Li et al. [18] proposed a multitask learning approach that enables accurate prediction by either using shape, texture, or both types of features. Though successful, a key drawback is the heavy reliance on artistic image generation techniques that may be detrimental for medical images and subsequent diagnoses. To address this, our approach is able to capture subtle differences in the image to generate consistent texture updated images. Thus, classifiers trained on the generated images can avoid subtle bias information.

Generative Texture Transfer. Methods [10] and [6] both proposed to generate texture updated images based on arbitrary style transfer, with adaptive
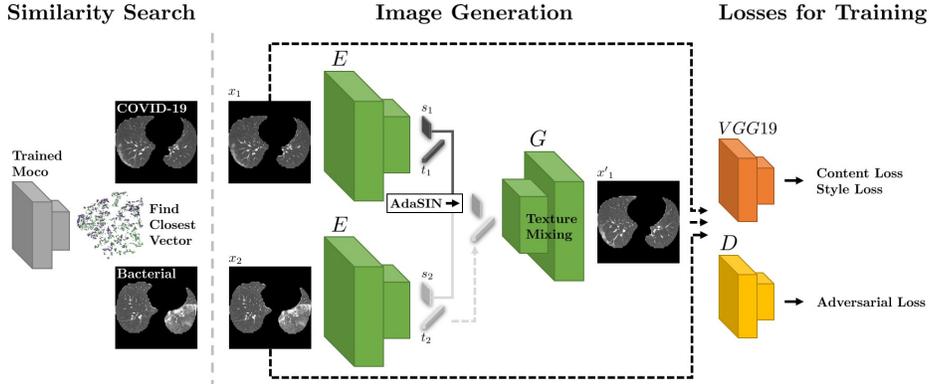
Fig. 1. Diagram of the proposed method.

and conditional instance normalization employed to transfer texture information. CycleGAN [27] is another popular method for texture transfer that uses the idea of consistency across several model outputs. However, these techniques not only change the texture but also induce structural distortion to the outputs which may introduce new forms of bias. Recently, Park et al. [20] achieved better results by using an autoencoder with texture swapping as well as a co-occurrence patch discriminator to capture high-level detail. In this method, the discriminator model may often change the original structural characteristics which is undesirable for medical images. Since our main objective is to maintain the structural information, we avoid techniques such as cycle consistency and patch discriminators that often produce structurally distorted images.

3   Methods

Given a chest CT dataset $\mathcal{D} = \{X^1, ..., X^n\}$ where X^i is a set of 2D CT slices $X^i = \{x_1^i, ..., x_m^i\}$ each with its label $y^i \in \{0, 1\}$ denoting bacterial pneumonia and COVID-19 samples, respectively. Our goal is to generate a dataset $\mathcal{D}' = \{X^{1\prime}, ..., X^{n\prime}\}$ where each $X^{i\prime}$ is a set of texture updated CT slices that may contain the bias information of the other label CT protocol. To achieve this, we first pre-train a contrastive encoder using $\mathcal{D}$ to learn representative slice features and then use it for slice similarity search i.e. $x_1^i, x_2^j \sim \mathcal{D}, y^i \neq y^j$ with semantically similar image structures. Second, to generate $x_1^{i\prime}$ we feed searched pairs x_1^i, x_2^j to an encoder network E that outputs structural and texture features used as input for a generator network G. Here, AdaSIN and texture *mixing* is employed on the respective features for improved generation. Lastly, following standard practice for adversarial-based methods, we also employ a discriminator network D to improve the quality of generation. To enforce style and content similarity, a pre-trained VGG19 [23] is used to optimize content loss between x_1' and x_1, and a style loss between x_1' and x_2, respectively. Through this process,

we can generate images to construct $\mathcal{D}'$ and then combine $\mathcal{D}$ and $\mathcal{D}'$ to train a classifier that will learn to ignore bias information. The overall framework is presented in Figure 1 with specific details of each step categorized below.

Slice Similarity Search. As opposed to arbitrarily sampling pairs for texture transfer in our generative framework, we employ similarity based sampling of image pairs with similar structural features between the classes. Here, we pre-train a momentum contrastive network (MoCo) [8] using the training data and find the closest slice based on the $L1$ distance. This is crucial for generation since using arbitrary image pairs for texture transfer can produce artificial images without any clinical significance. Following, we construct image pairs for the entire dataset for image generation.

Image Generation with AdaSIN and Texture *Mixing*. To generate a de-biased and texture mixed image, an encoder network E takes as input the sampled image pairs to produce texture and structure features that are first modified via AdaSIN before being feed to G. To retain texture and structure, both are required as inputs for G. Specifically, the structure feature is a feature with spatial dimensions pooled at an earlier stage of E, whereas the texture feature is an embedding vector obtained following 1×1 convolution in the later stages of E. In addition, we believe that directly employing these features for generation can still retain the inherent bias contained in the features, thus we propose to modify the structural features via AdaIN [10]. To achieve this, we use the mean($\mu(\cdot)$) and standard deviation ($\sigma(\cdot)$) of the features before passing to G. Formally,

$$AdaSIN(s_1, s_2) = \sigma(s_2) \left(\frac{s_1 - \mu(s_1)}{\sigma(s_1)} \right) + \mu(s_2), \tag{1}$$

where s_1 and s_2 denotes the extracted structure features of the input image pairs i.e. $x_1, x_2 \sim \mathcal{D}$. Next, G takes both features for an image generation and texture transfer. Texture transfer is achieved via convolution weight modulation and demodulation introduced in [14]. Herein, texture information is delivered to each convolutional layer in G except for the last layer.

To train the entire framework, we follow the arbitrary style transfer training procedure described in [10]. A VGG19 pre-trained model is used to extract features from the input pairs which is then used in the style L_{style} and content $L_{content}$ losses, respectively. The style loss minimizing the mean squared error (MSE) between the generated output and the texture image, whereas the content loss minizing the MSE between generated output and the structure image [17]. Further, we use an adversarial loss to improve the image generation quality via $L_{GAN}(G, D) = -\mathbb{E}_{x_1, x_2 \sim \mathcal{D}}[D(G(AdaSIN(s_1, s_2), t_2))]$, with regularization (non-saturation loss) omitted for simplicity [7,14]. The final loss function is defined as:

$$L_{total} = L_{content} + \lambda L_{style} + L_{GAN}. \tag{2}$$

Implementation Details. A Mask-cascade-RCNN-ResNeSt-200 [26] with deformable convolution neural network (DCN) [4] was employed to extract the lung and lesion regions in the CT scans to mask out the non-lung regions and

Table 1. The statistics and bias information (CT protocol) of the train/validation/test dataset.

Split	COVID	CT protocol	Bacterial	CT protocol
Train	42	Non-Contrast (kVp: 120).	60	Contrast (kVp: 100).
Val	10	Non-Contrast (kVp: 120).	14	Contrast (kVp: 100).
Test	21	Various CT protocols from different scanners: Contrast (18), etc (3).	23	Various CT protocols from different scanners: Non-Contrast (18), etc (5).

excluding slices that do not contain a lesion. The contrastive slice encoder was trained for 200 epochs whereas the generative model was trained with a batch size of 8 for 40,000 steps with Adam optimizer and learning rate of 0.002. Further, the discriminator D follows Karras et al.'s [13,14] implementation and we empirically set $\lambda = 10$.

4   Experiments and Results

Experiments Settings For evaluation, we constructed an in-house biased dataset for training/validation and an unbiased testing dataset. First, we extracted the CT protocol per scan using metadata in the DICOM files and create splits COVID-19 and bacterial pneumonia based on the protocol. The dataset and bias information are shown in Table 1.

To evaluate the effect of bias mitigation using the generated images for classification, a pre-trained ResNet18 and recent COVID-19 classification models [9,16,24,25] i.e. COVNet, COVID-Net, Contrastive-COVIDNet were compared. The models were trained for 100 epochs with a batch size 64 and a learning rate 0.001. We also applied random crops, horizontal flips and intensity augmentations i.e. (brightness and contrast) as a baseline augmentation technique. Performance comparison of our approach against recent state-of-the-art non-generation based bias mitigation methods [15,1] applied for natural image classification is also reported. To verify the effectiveness of the proposed method, we include comparisons against a commonly used arbitrary style transfer model i.e. AdaIN, and the current state-of-the-art generation method i.e. swapping-autoencoder [10,20]. For a fair comparison of the generation based methods, the same texture pairs were utilized for a generation. Also, training and validation were performed three times for all methods with final average performance reported.

Results on Internal Dataset In Table 2, we present the evaluation results on the biased COVID-19 vs. bacterial pneumonia dataset. Initially, the model shows high f1-score on the validation dataset i.e. 97.18%, yet significantly drops to 33.55% on the unbiased test set. This shows that the classifier makes predictions based on the bias information. The results of the learning-based models

Table 2. The test results (f1-score) on biased COVID-19 vs. bacterial pneumonia dataset. The results of validation dataset are written as Val. The model without intensity augmentation is denoted w/o aug.

	ResNet18 [9]	COVNet [16]	COVID-Net [24]	Contrastive COVID-Net [25]
Val w/o aug	97.18	89.33	78.30	93.31
Base	66.97	63.41	41.89	53.37
Base w/o aug	33.55	59.52	43.99	35.64
LNTL [15]	31.87	-	-	-
Blind eye [1]	32.72	-	-	-
AdaIN [10]	75.71	66.35	51.35	65.76
Swap [20]	76.84	67.64	56.28	62.60
Mixing-AdaSIN	**80.97**	**74.61**	**61.57**	**66.16**

i.e. Learning not to learn, and Blind eye show no considerable performance improvements and highlight the failure to mitigate the bias, especially for medical domain images as reported in [2]. Further, these methods were proposed to remove bias that is distinctly recognizable in the image such as color. Capturing and mitigating the subtle bias difference in the medical image is considerably harder for such techniques.

Among generation based methods, the proposed method reports the best performance. Even though AdaIN can transfer texture well, the quality of the generated image is extremely low. Consequently, this inhibits classifier training as shown by the limited performance improvements. Though swapping-autoencoder updates the texture with high quality image generation results, two major drawbacks were noted: (i) the generated image still retains bias, and (ii) it distorts key representative characteristics by artificial translation of lesions i.e. changing the lesion of COVID-19 to appear as bacterial pneumonia. Such phenomena may be due to the direct usage of structure features, and use of the co-occurrence discriminator which leads to structural information deformation. On the other hand, our model employs two modules i.e. boosting a high-quality image gen-
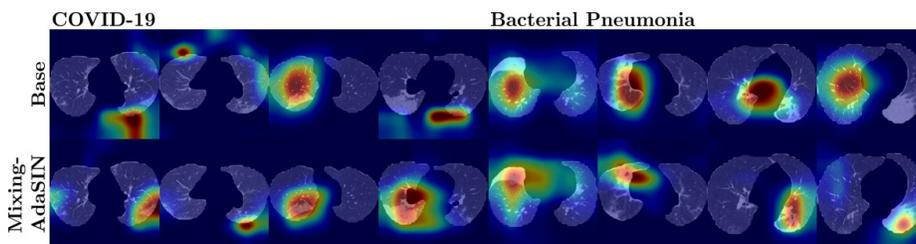

Fig. 2. Grad-CAM [22] visualizations of the ResNet18 [9] classifier. Herein, Grad-CAM of the base classifier pointed the normal lung area. On the other hand, the classifier trained with Mixing-AdaSIN pointed the lesion correctly. Hence, the model with debiasing can be more generalized.

Table 3. The test results (accuracy) of the ResNet18 [9] on the external COVID-19 dataset.

	Base	Base w/o aug	AdaIN [10]	Swap [20]	Mixing-AdaSIN
Accuracy	24.11	34.04	50.35	76.60	**77.30**

eration via texture *mixing*; and minimizing the bias artifact transfer through an AdaSIN. We consider these techniques as instrumental in mitigating bias. In addition, the classifier with a de-biased method employed a more generalized feature as shown in Figure 2. The Grad-CAM [22] of the trained classifier pointed to the lesion more correctly, thus a better generalization performance is expected.

Our proposed method can be easily applied to existing CT based COVID-19 classification models. In particular, COVNet reported 74.61% f1-score which represents successful mitigation of bias artifacts. However, COVID-Net and Contrastive COVID-Net showed relatively low accuracy, mainly due to slight differences in training details and architectures. Also, due to the long-range connectivity in the models, reliance on bias information is heavily induced during training.

Results on External Dataset To verify the generalization efficacy of our trained classifier on external data, we employ the publicly available MosMed dataset [19]. It consists of 1110 CT scans from COVID-19 positive patients. However, the dataset contains CT scans that are not consistent with pneumonia observed in our original dataset. Thus, we selected scans of severe and critical patients only to evaluate the trained models. In addition, since we trained three classifiers from an internal experiment, we tested each classifier three times and final average performance reported. In Table 3, results are fairly consistent with improvements shown on the internal dataset evaluation. Our model shows a significant improvement over the baseline with +1% gain over swapping-autoencoder. More importantly, even though the classifier has not observed the CT samples with a different protocol, performance was still consistent verifying the utility of the proposed de-biasing technique.

5   Conclusion

In this work, we have proposed a novel methodology to train a COVID-19 vs. bacterial pneumonia classifier that is robust to bias information present on training data. We constructed an in-house biased training dataset in conjunction with an unbiased testing dataset and proved that our method allowed the classifier to learn the appropriate features to correctly predict the labels without considering bias information and achieved better generalization. All of this was possible thanks to an adequate image generation design that relies on two major components: (a) texture *mixing*, which enables realistic image generation, and (b)

AdaSIN, which prevents bias flow from the input to the output image in the generation stage, while maintaining structural consistency. We proved the benefits of our pipeline by achieving the best bias mitigation performance when compared to other related methods in both our in-house dataset as well as in an external dataset. Considering that biases can be easily included when constructing datasets, we hope that our findings help to improve performance in various medical tasks.

References

1. Alvi, M., Zisserman, A., Nellåker, C.: Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In: Proceedings of the European Conference on Computer Vision Workshops. pp. 0–0 (2018)
2. Bissoto, A., Valle, E., Avila, S.: Debiasing skin lesion datasets and models? not so fast. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 740–741 (2020)
3. Caschera, L., Lazzara, A., Piergallini, L., Ricci, D., Tuscano, B., Vanzulli, A.: Contrast agents in diagnostic imaging: Present and future. Pharmacological research **110**, 65–75 (2016)
4. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 764–773 (2017)
5. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In: International Conference on Learning Representations (2019)
6. Ghiasi, G., Lee, H., Kudlur, M., Dumoulin, V., Shlens, J.: Exploring the structure of a real-time, arbitrary neural artistic stylization network. In: British Machine Vision Conference (2017)
7. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of wasserstein gans. In: Advances in Neural Information Processing Systems. p. 5769–5779. Curran Associates Inc. (2017)
8. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
10. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1501–1510 (2017)
11. Kalra, M.K., Homayounieh, F., Arru, C., Holmberg, O., Vassileva, J.: Chest ct practice and protocols for covid-19 from radiation dose management perspective. European Radiology pp. 1–7 (2020)

12. Kang, M., Hong, K.S., Chikontwe, P., Luna, M., Jang, J.G., Park, J., Shin, K.C., Park, S.H., Ahn, J.H.: Quantitative assessment of chest ct patterns in covid-19 and bacterial pneumonia patients: a deep learning perspective. Journal of Korean medical science **36**(5) (2021)
13. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: International Conference on Learning Representations (2018)
14. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8110–8119 (2020)
15. Kim, B., Kim, H., Kim, K., Kim, S., Kim, J.: Learning not to learn: Training deep neural networks with biased data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9012–9020 (2019)
16. Li, L., Qin, L., Xu, Z., Yin, Y., Wang, X., Kong, B., Bai, J., Lu, Y., Fang, Z., Song, Q., et al.: Using artificial intelligence to detect covid-19 and community-acquired pneumonia based on pulmonary ct: evaluation of the diagnostic accuracy. Radiology **296**(2), E65–E71 (2020)
17. Li, Y., Wang, N., Liu, J., Hou, X.: Demystifying neural style transfer. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence. p. 2230–2236. IJCAI'17, AAAI Press (2017)
18. Li, Y., Yu, Q., Tan, M., Mei, J., Tang, P., Shen, W., Yuille, A., Xie, C.: Shape-texture debiased neural network training. In: International Conference on Learning Representations (2021)
19. Morozov, S.P., Andreychenko, A.E., Blokhin, I.A., Gelezhe, P.B., Gonchar, A.P., Nikolaev, A.E., Pavlov, N.A., Chernina, V.Y., Gombolevskiy, V.A.: Mosmeddata: data set of 1110 chest ct scans performed during the covid-19 epidemic. Digital Diagnostics **1**(1), 49–59 (2020)
20. Park, T., Zhu, J.Y., Wang, O., Lu, J., Shechtman, E., Efros, A., Zhang, R.: Swapping autoencoder for deep image manipulation. In: Advances in Neural Information Processing Systems. vol. 33, pp. 7198–7211. Curran Associates, Inc. (2020)
21. Pontone, G., Scafuri, S., Mancini, M.E., Agalbato, C., Guglielmo, M., Baggiano, A., Muscogiuri, G., Fusini, L., Andreini, D., Mushtaq, S., et al.: Role of computed tomography in covid-19. Journal of cardiovascular computed tomography (2020)
22. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
23. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (2014)
24. Wang, L., Lin, Z.Q., Wong, A.: Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. Scientific Reports **10**(1), 1–12 (2020)
25. Wang, Z., Liu, Q., Dou, Q.: Contrastive cross-site learning with redesigned net for covid-19 ct classification. IEEE Journal of Biomedical and Health Informatics **24**(10), 2806–2813 (2020)
26. Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Zhang, Z., Lin, H., Sun, Y., He, T., Muller, J., Manmatha, R., Li, M., Smola, A.: Resnest: Split-attention networks. arXiv preprint arXiv:2004.08955 (2020)
27. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)