ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 978-1-7281-7605-5/20/$31.00 ©2021 IEEE | DOI: 10.1109/ICASSP39728.2021.9414327

SUPER-RESOLUTION AND INFECTION EDGE DETECTION CO-GUIDED LEARNING FOR COVID-19 CT SEGMENTATION

Yu Sang [1], Jinguang Sun [1], Simiao Wang [1], Heng Qi [2], Keqiu Li [3]

[1] School of Electronic and Information Engineering, Liaoning Technical University, Huludao 125105, China
[2] School of Computer Science and Technology, Dalian University of Technology, Dalian 116023, China
[3] College of Intelligence and Computing, Tianjin University, Tianjin 300072, China

ABSTRACT

In this paper, we propose a novel super-resolution and infection edge detection *co*-guided learning network for COVID-19 CT *seg*mentation (*CogSeg*). Our *CogSeg* is a coherent framework consisting of two branches. Specifically, we use image super-resolution (SR) as an auxiliary task, which assist segmentation to recover high-resolution representations. Moreover, we propose an infection edge detection guided region mutual information (RMI) loss, which uses the edge detection results of segmentation to explicitly maintain the high order consistency between segmentation prediction and ground truth around infection edge pixels. Our *CogSeg* network can effectively maintain high-resolution representation and leverages edge details to improve the segmentation performance. When evaluated on two publicly available COVID-19 CT datasets, our *CogSeg* improves 10.63 and 13.02 points than the established baseline method (i.e. U-Net) w.t.r *mIoU*. Moreover, our *CogSeg* achieves more appealing results both quantitatively and qualitatively than the state-of-the-art methods.

Index Terms— COVID-19, Computed Tomography (CT); segmentation, Super-Resolution (SR), edge detection

1. INTRODUCTION

The rapid spread of COVID-19 has seriously threatened global health. Accurate and rapid testing is extremely vital for timely prevention of COVID-19 spread. As a basic but challenging task of the diagnostic framework, segmentation of the infection lesions from CT scans plays a crucial role for quantitative measurement of the disease progression in accurate diagnosis and follow-up assessment [1].

Recently, deep learning (DL) has penetrated into the field of medical imaging and has a good development [1-4], which brings unique benefits to the automated segmentation of medical images. U-Net [4] is the most widely used encoder-decoder deep network architecture and work decently well for medical image segmentation. Subsequently, some variants of U-Net [5-7] are also proposed to improve the performance for segmentation of different organs. For COVID-19, some DL-based methods [1,8,9] have been proposed to detect patients infected via radiological imaging to provide assistance in diagnosis, but there are only a few works [10-12] related COVID-19 CT infection segmentation with deep learning, which remains under explored and is still a challenging task. On the other hand, high resolution representations and edge information are two important factors for CT segmentation, but they are ignored in current methods.

In this paper, we propose a super-resolution and infection edge detection *co*-guided learning network for COVID-19 CT *seg*mentation (*CogSeg*). Our contributions are threefold. 1) We integrate SR learning into COVID-19 CT segmentation network based on U-Net encoder-decoder backbone, which makes full use of the fine-grained structural information recovered from SR learning to facilitate CT segmentation. 2) Considering the high correlations between segmentation and infection edge, we use the latter one as an auxiliary task, and allow it provide explicit guidance to assist CT segmentation. The guidance is built by developing an infection edge detection guided RMI loss to constrain segmentation predictions around infection edge regions. 3) The experiments on two open-access COVID-19 CT datasets show that our *CogSeg* outperforms other state-of-the-art segmentation methods and yields more precise predictions.

2. RELATED WORK

2.1. Medical Image Segmentation

Driven by the rapid development of convolutional neural networks (CNNs), FCNs [13,14], DeepLabs [15], DSRL [16] and HRNet [17] based methods have dominated majority of this field. For medical image segmentation, Ronneberger et al. [4] propose a well-known U-Net architecture. SegNet [18] makes careful designs to improve the U-Net. Li et al. [5] propose an H-DenseUNet to efficiently extract intra-slice features for liver and tumor segmentation. Attention U-Net [6,19] improves U-Net by using the attention mechanism to learn to focus on target structures. U-Net++ [7] improves U-Net by introducing a series of nested, dense skip connections between the encoder and decoder sub-networks.

Due to the fast progression and infectious ability of COVID-19 that have never appeared before, there has only a few works related CT segmentation with deep learning. Currently, Zhou et al. [10] propose a machine-agnostic method that can segment and quantify the infection regions

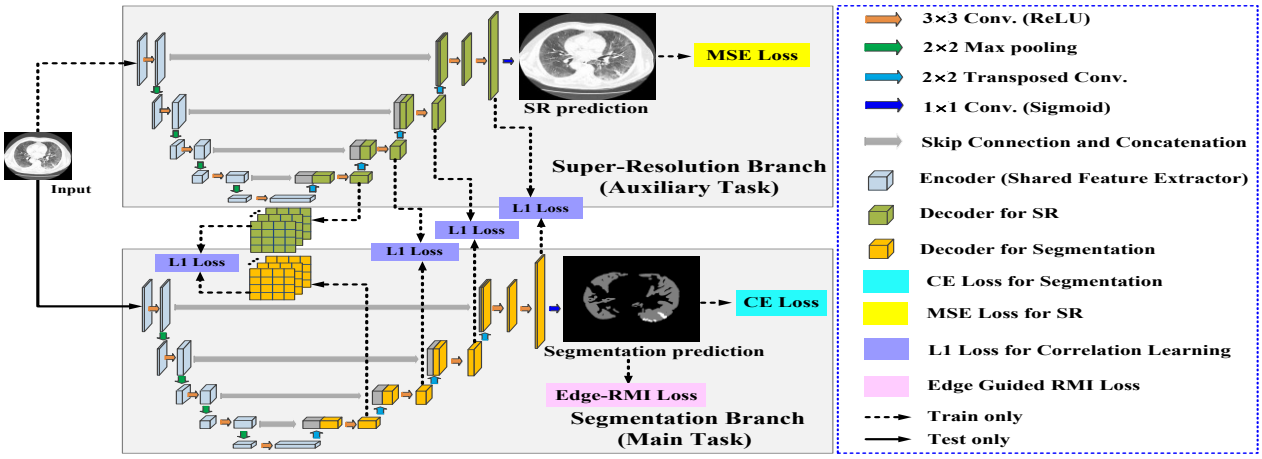

Figure 1: The architecture of our *CogSeg*. It consists of two branches, which are used in training phase, and only segmentation is conducted in test phase. The encoders are shared between SR and segmentation. The architecture will be optimized with four terms: MSE loss for SR, CE loss for segmentation, L1 loss for correlation learning of spatial dimension between output of decoder for segmentation and output of decoder for SR, and edge detection guided RMI loss.

on CT scans. Fan et al. [11] propose an Inf-Net, which utilizes an implicit reverse attention and a semi-supervised solution to improve the identification of infected regions. Chen et al. [12] propose a residual attention U-Net for automated segmentation of COVID-19 chest CT images.

Actually, edge information can help improve the segmentation performance, which has not been paid attention to in previous works. To this end, we exploit an edge guided loss function to address this issue in this work.

2.2. Image Super-Resolution

Recently, CNNs have indicated that they can provide remarkable performance in the SR problem. Dong et al. [20] proposed the first CNNs-based SRCNN. Since then, many CNNs-based SR versions are proposed; the networks are tend to be deeper and deeper from the SRCNN to deeper VDSR [21] and Memnet [22], etc., and then to the very deep RCAN [23]. Furthermore, other effective methods construct the entire network by connecting a series of identical feature extraction modules such as RDN [24] and MSRN [25], indicating the capability of each block plays a crucial role. In addition, some deep learning based medical image SR methods [26-28] are well developed. However, they only reconstruct high-resolution images, and do not pay attention to whether the reconstructed medical images are useful for high-level tasks such as semantic segmentation.

In this work, we employ U-Net as SR backbone, which aims to be the same as network structure of segmentation to guide the correlation learning of spatial dimension, thus benefiting the segmentation task.

3. PROPOSED METHOD

3.1. Overview

As shown in Figure 1, our *CogSeg* framework consists of two branches: SR and segmentation. We treat segmentation as the main task, and SR learning as auxiliary task. SR branch is optimized with reconstruction supervision during training, and it will be freely removed from the network in the inference stage, thus causing cost-free overhead. We define the overall learning objective function as:

$$L_{tatal} = \lambda_1 \cdot L_{CE} + \lambda_2 \cdot L_{MSE} + \lambda_3 \cdot \sum_{i=1}^{4} L1(i) + \lambda_4 \cdot L_{RMI}^{Egde}, \quad (1)$$

where λ_1, λ_2, λ_3 and λ_4 are weighting parameters to balance four losses. L_{CE} and L_{MSE} are the conventional multi-class Cross-Entropy (CE) loss for segmentation [4] and Mean Squared Error (MSE) loss for SR [20], respectively. We choose L1 loss function [25] for correlation learning of spatial dimension between output of decoder for segmentation and output of decoder for SR. In Eq. (1), the weights of segmentation task have greater values and bring larger gradients when training; while L1 loss has a smaller gradient, which is equivalent to providing high-resolution clues for segmentation branch. $L1(i)$ is the i-th L1 loss of the four L1 losses, defined as follows:

$$L1(i) = \frac{1}{C_i \cdot N} \sum_{c_i=1}^{C_i} \sum_{j=1}^{N} \left\| G\left(X_{c_i,j}\right) - G\left(Y_{c_i,j}\right) \right\|_1, \quad (2)$$

where N means the pixel number, C_i denotes the number of feature maps generated by the i-th decoder. $G\left(X_{c_i,j}\right)$ and $G\left(Y_{c_i,j}\right)$ refer the C_i feature maps generated by the i-th decoders of SR and segmentation, respectively. Moreover, L_{RMI}^{Edge} indicates our infection edge detection guided RMI loss, which will be further described in Section 3.3. We minimize the whole objective function end-to-end.

3.2. SR Guided Learning for CT Segmentation

For semantic segmentation, current most of state-of-the-art methods [4,15,29] often subsample the original image and the corresponding ground truth, and then rescaling to the size of the original image in the post-processing stage. This may result in the loss of effective label information due to

subsampled ground truth. Besides, they also suffer from another common disadvantage; that is, recovering high-resolution features via the decoder module is difficult to recover the original details. Because the decoder is either a bilinear upsampling layer or a simple sub-network, it will not bring any additional information since the input is in a subsampled low-resolution. So, only relying on the decoder module is not enough to recover high-resolution feature representation.

There are still such problems for COVID-19 CT segmentation. To improve the performance of CT segmentation, we integrate SR learning with segmentation, which apply the high-resolution features recovered from SR to guide the learning of high-resolution representations of segmentation and enable the segmentation framework to achieve comparable performance with low computational cost. Specifically, we employ the same U-Net encoder-decoder architectures as SR backbone and segmentation backbone as shown in Figure 1, which aims to guide the correlation learning of spatial dimension between output of decoder for segmentation and output of decoder for SR as shown the L1 loss optimization between the green feature maps and orange feature maps in Figure 1. The details can be modeled by the relationships between internal pixels, making up for the simple design of the decoder. Our network starts with two encoders with shared weights, which aims at learning the task-independent representations for SR and segmentation, simultaneously. And then, we use two decoders to learn task-dependent representations for two tasks, respectively. At the end of SR decoder, we apply one 1×1 convolutional layer to reconstruct high-resolution image, while at the end of segmentation decoder, we apply one 1×1 convolutional layer and one soft max layer to predict a probability distribution for each pixel. Specifically, we use additional supervision at different scales of decoder to produce an amplified output for two tasks, e.g., taking an input of 256×256, and generating an output of 512×512. In brief, we leverage SR representation to assist segmentation.

3.3. Infection Edge Detection Guided RMI Loss

The CE loss is most commonly used in semantic segmentation learning. But it ignores the relationship among pixels in a local region of the image. Actually, there are strong dependencies among pixels and these dependencies represent the structure of objects. To solve this issue, we introduce the region mutual information (RMI) loss [30] to model the dependencies among pixels efficiently. When observing the CT segmentation maps, we find that most infection regions exhibit flat texture, and only those regions around the boundaries show distinct structures. Based on this observation, we propose an infection edge detection guided RMI loss for COVID-19 CT segmentation. Intuitively, as show in Figure 2, given the infection edge

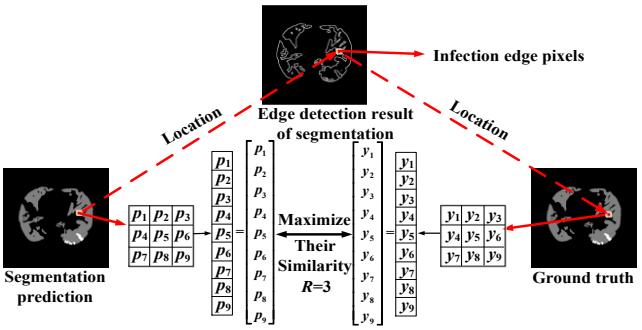

Figure 2 Illustration of the infection edge detection guided RMI loss (R=3).

detection maps that always obtained by segmentation prediction itself, for each boundary pixel, we define a square window centered with it; we map this local window to the segmentation prediction and ground truth maps, and collect the pixel values inside the window, generating two distributions; finally, we maximize the similarity between the two distributions. Specifically, if the size of square region is $R\times R$, then $d=R\times R$, we can get multivariate random variables $P=[p_1, p_2, \dots, p_d]$ from CT segmentation predictions and $Y=[y_1, y_2, \dots, y_d]$ from ground truth. Following [30], we maximize the lower bound of mutual information between their multivariate distributions. Suppose the variables of arbitrary pixels and boundary pixels are P_1 and P_2, respectively. Then, the infection edge detection guided RMI (Edge-RMI) loss can be formulated as

$$L_{RMI}^{Egde}=-\frac{1}{B}\sum_{b=1}^{B}\sum_{e=1}^{E}\frac{1}{2}\left(\log\left(\det\left(\Sigma_{Y|P_1}^{b,e}\right)\right)_{R\in\{3\}}+\log\left(\det\left(\Sigma_{Y|P_2}^{b,e}\right)\right)_{R\in\{5\}}\right), \quad (3)$$

where B denotes the number of images in a mini-batch, E is the number of semantic categories, $\det(\cdot)$ is the determinant of the matrix, and $\Sigma_{Y|P}$ is the posterior covariance matrix of Y, given P. Since the number of edge pixels is small, we add a full image RMI loss to provide a more comprehensive constraint over local structures.

4. EXPERIMENTS

4.1. Datasets

We evaluate our *CogSeg* on two open-access COVID-19 CT segmentation datasets [31], referred to as COVID-19_1/2. COVID-19_1 includes 100 axial CT images from 60 patients, including three labels: ground-glass opacities, consolidation and pleural effusion [32]. 70% samples are selected as training set, and the rest is as test set. COVID-19_2 contains 9 volumes, total 829 slices, where 373 slices have lesions and been labeled by radiologists. The first 6 volumes are defined as the training set, and the rest is as the test set.

4.2. Implementation details

We implement *CogSeg* in PyTorch. All models are trained from scratch on two NVIDIA 2080Ti GPUs. We use ADAM optimizer to train our model and follow the training details

Table 3 Comparison with state-of-the-art methods. Bold indicates the best results; underline indicates the second best results.

Datasets	Metrics	Methods							
		U-Net	U-Net++	Attention U-Net	H-DenseUNet	FCN-8s	DeepLabv2	HRNet	*CogSeg* (Ours)
COVID-19_1	*mIoU*	79.12	84.21	83.72	83.94	83.12	85.98	<u>87.56</u>	**89.75**
	DSC	73.35	77.24	75.45	74.78	75.15	77.55	<u>80.35</u>	**83.04**
	Sens.	80.36	83.53	83.84	82.65	83.41	83.52	<u>85.77</u>	**86.98**
	Spec.	95.83	96.56	96.86	96.29	96.66	96.46	<u>97.18</u>	**98.09**
COVID-19_2	*mIoU*	60.44	68.35	68.98	67.93	68.24	68.45	<u>71.54</u>	**73.46**
	DSC	57.35	66.72	67.85	63.32	67.47	67.69	<u>70.02</u>	**72.43**
	Sens.	66.98	75.779	76.36	74.66	75.15	76.28	<u>77.23</u>	**79.32**
	Spec.	94.67	96.78	97.05	96.75	96.59	97.44	<u>97.65</u>	**98.16**

Table 1 The effectiveness of SR.

Datasets	Methods	Metrics (%)			
		mIoU	*DSC*	*Sens.*	*Spec.*
COVID-19_1	U-Net (Baseline)	79.12	73.35	80.36	95.83
	+SR with 1 L1 loss	85.03	77.46	83.78	96.14
	+SR with 4 L1 losses	**86.24**	**78.89**	**84.63**	**96.65**
COVID-19_2	U-Net (Baseline)	60.44	57.35	66.98	94.67
	+SR with 1 L1 loss	67.81	66.84	73.84	96.16
	+SR with 4 L1 losses	**69.96**	**68.52**	**75.47**	**96.88**

Table 2 Influence of loss functions.

Datasets	Loss	Metrics (%)			
		mIoU	*DSC*	*Sens.*	*Spec.*
COVID-19_1	CE (Baseline)	86.24	78.89	84.63	96.65
	+RMI	87.45	80.67	85.16	97.33
	+Edge-RMI (ours)	**89.75**	**83.04**	**86.98**	**98.09**
COVID-19_2	CE (Baseline)	69.96	68.52	75.47	96.88
	+RMI	71.02	70.15	77.21	97.43
	+Edge-RMI (ours)	**73.46**	**72.43**	**79.32**	**98.16**

of U-Net [4]. The initial learning rate is set to 0.0001 for all layers and decreased by half after every 50 epochs. Our model converges after 250 epochs. For our joint loss function, we set the weight of each term as λ_1=1, λ_2=0.5, λ_3=0.5 and λ_4=0.5 in Eq. (1). We set $R \in \{3,5\}$ as shown in Eq. (3). For COVID-19_1, consider about the size of the dataset is small, data augmentation is necessary for training the neural network to achieve high generalizability [21,22]. We rotate the existing images and the corresponding masks 45◦, 90◦, 135◦, 180◦, 270◦ and 360◦ to generate another 600 examples. Besides, we scale the image to 0.5, 0.75, 1.25 and 1.5 separately to generate another 400 images and its corresponding masks.

4.3. Ablation Study

Effectiveness of SR. Following [33], we use four widely adopted metrics, i.e., Mean Intersection over Union (*mIoU*), Dice Similarity Coefficient (*DSC*), Sensitivity (*Sens.*), and Specificity (*Spec.*). As shown in Table 1, the performance can be improved while adding the SR with 1 L1 loss and 4 L1 losses of Eq. (1) into segmentation, thus indicating that transferring the structure information between SR and segmentation is necessary.

Influence of Learning Objectives. The evaluation of different learning objectives is shown in Table 2. Compared to the baseline only using CE loss and additional RMI loss, we obtain improvement by adding our Edge-RMI loss. This gain is mainly due to more accurate predictions around the boundary regions between semantic parts.

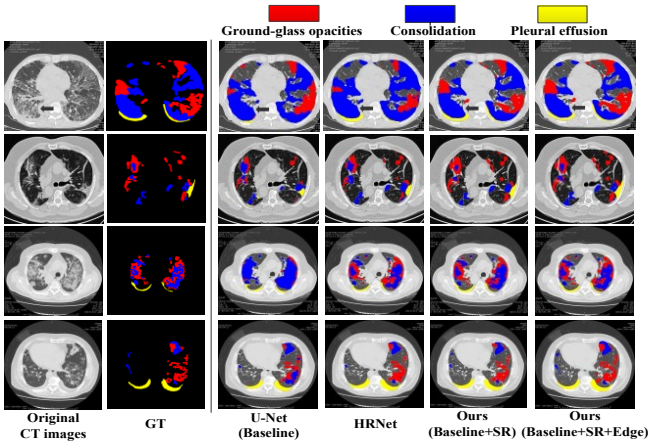

Figure 3 Visualization of infection segmentation results produced by state-of-the-art methods and our method.

4.4. Comparisons with the State-of-the-arts

We compare our *CogSeg* with U-Net [4], U-Net++ [7], Attention U-Net [6], H-DenseUNet [5], FCN-8s [14], DeepLabv2 [15] and HRNet [17]. Quantitative results are shown in Table 3. As can be seen, the proposed *CogSeg* outperforms other methods in terms of four metrics. We attribute this improvement to the integration of SR in segmentation and the attention of edge regions. In addition, we show some qualitative results in Figure 3, in which our method yields more precise predictions, while other two methods, especially U-Net, loss components and lead to wrong classification.

5. CONCLUSION

In this work, we investigate the task of COVID-19 CT segmentation from two aspects: high-resolution representations and infection edge detection guidance. Specifically, we first propose to integrate SR learning into segmentation to guide the learning of high-resolution representations of segmentation. We further propose an infection edge detection guided RMI loss that leverages the edge detection results to progressively refine the predictions of CT segmentation. The experimental results demonstrate the effectiveness of the proposed *CogSeg*.

6. REFERENCES

[1] F. Shi, J. Wang, J. Shi, *et al*, "Review of Artificial Intelligence Techniques in Imaging Data Acquisition, Segmentation and Diagnosis for COVID-19," *IEEE Reviews in Biomedical Engineering*, 3333(c), pp. 1–13, 2020.

[2] D. Nie, Y. Gao, L. Wang, and D. Shen, "ASDNet: Attention Based Semi-supervised Deep Networks for Medical Image Segmentation," In *MICCAI*, pp. 370–378, 2018.

[3] N. Tajbakhsh, L. Jeyaseelan, Q. Li, *et al*, "Embracing Imperfect Datasets: A Review of Deep Learning Solutions for Medical Image Segmentation," *Medical Image Analysis*, 63: 101693, 2020.

[4] O. Ronneberger, F. Philipp, and B. Thomas, "U-Net: Convolutional Networks for Biomedical Image Segmentation," In *MICCAI*, pp. 234–241, 2015.

[5] X.M. Li, H. Chen, X.J. Qi, D. Qi, C.W. Fu, and P.A. Heng, "H-DenseUNet: Hybrid Densely Connected U-Net for Liver and Tumor Segmentation from CT Volumes," *IEEE Transactions on Medical Imaging*, 37(12), pp. 2663–2674, 2018.

[6] O. Ozan, J. Schlemper, L.L. Folgoc, *et al*, "Attention U-Net: Learning Where to Look for the Pancreas," In *International Conference on Medical Imaging with Deep Learning*, 2018.

[7] Z. Zhou, M.M.R. Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: Redesigning Skip Connections to Exploit Multiscale Features in Image Segmentation," *IEEE Transactions on Medical Imaging*, pp. 3–11, 2019.

[8] L. Wang and A. Wong, "COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest Radiography Images," *arXiv:2003.09871*, 2020.

[9] J. Zhang, Y. Xie, Y. Li, C. Shen, and Y. Xia, "COVID-19 Screening on Chest X-ray Images Using Deep Learning Based Anomaly Detection," *arXiv:2003.12338*, 2020.

[10] L.X. Zhou, Z.X. Li, J.X. Zhou, *et al*, "A Rapid, Accurate and Machine-agnostic Segmentation and Quantification Method for CT-based COVID-19 Diagnosis," *IEEE Transactions on Medical Imaging*, PP(99), 2020.

[11] D.P. Fan, T. Zhou, G.P. Ji, *et al*, "Inf-net: Automatic COVID-19 Lung Infection Segmentation from CT Images," *IEEE Transactions on Medical Imaging*, 39(8), pp. 2626–2637, 2020.

[12] X.C. Chen, L.N. Yao, and Y. Zhang, "Residual Attention U-Net for Automated Multi-Class Segmentation of COVID-19 Chest CT Images," *arXiv:2004.05645*, 2020.

[13] J. Long, E. Shelhamer, and T. Darrell, "Fully Convolutional Networks for Semantic Segmentation," In *CVPR*, pp. 3431–3440, 2015.

[14] E. Shelhamer, J. Long, and T. Darrell, "Fully Convolutional Networks for Semantic Segmentation," In *T-PAMI*, 39(4), pp. 640–651, 2017.

[15] L.C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. Yuille, "Deeplab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs," In *T-PAMI*, 40(4), pp. 834–848, 2017.

[16] L. Wang, D. Li, Y.S. Zhu, L. Tian, and Y. Shan, "Dual Super-Resolution Learning for Semantic Segmentation," In *CVPR*, pp. 3773–3782, 2020.

[17] K. Sun, B. Xiao, D. Liu, and J.D. Wang, "Deep High-Resolution Representation Learning for Human Pose Estimation," In *CVPR*, pp. 5693–5703, 2019.

[18] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation," In *T-PAMI*, 39(12), pp. 2481–2495, 2017.

[19] G. Gaál, B. Maga, and A. Lukács, "Attention U-net based Adversarial Architectures for Chest X-ray Lung Segmentation," *arXiv preprint arXiv:2003.10304*, 2020.

[20] C. Dong, C.C. Loy, K.M. He, and X.O. Tang, "Learning a Deep Convolutional Network for Image Super-Resolution," In *ECCV*, pp. 184–199, 2014.

[21] J. Kim, J.K. Lee, and K.M. Lee, "Accurate Image Super-Resolution Using Very Deep Convolutional Networks," In *CVPR*, pp. 1646–1654, 2016.

[22] Y. Tai, J. Yang, X.M. Liu, and C.Y. Xu, "MemNet: A Persistent Memory Network for Image Restoration," In *CVPR*, pp. 4539–4547, 2017.

[23] Y.L. Zhang, K.P. Li, K. Li, L.C. Wang, B.N. Zhong, and Y. Fu, "Image Super-Resolution Using Very Deep Residual Channel Attention Networks," In *ECCV*, pp. 294–310, 2018.

[24] Y.L. Zhang, Y.P. Tian, Y. Kong, and B.N. Zhong, "Residual Dense Network for Image Super-Resolution," In *CVPR*, pp. 2472-2481, 2018.

[25] J.C. Li, F. Fang, K. Mei, *et al*, "Multi-Scale Residual Network for Image Super-Resolution," In *ECCV*, pp. 527-542, 2018.

[26] A.S. Chaudhari, Z. Fang, F. Kogan, *et al*, "Super-Resolution Musculoskeletal MRI Using Deep Learning," *Magnetic Resonance in Medicine*, 80(5), pp. 2139–2154, 2018.

[27] Y. Sang, J.G. Sun, S.M. Wang, H. Qi, and K.Q. Li, "Medical Image Super-Resolution via Granular Multi-Scale Network in NSCT Domain," In *ICME*, pp. 1–6, 2020.

[28] C.P. Wang, S.M. Wang, B. Ma, *et al*, "Transform Domain Based Medical Image Super-Resolution via Deep Multi-Scale Network," In *ICASSP*, pp. 2387–2391, 2019.

[29] H.S. Zhao, J.P. Shi, X.J. Qi, X.G. Wang, and J.Y. Jia, "Pyramid Scene Parsing Network," In *CVPR*, 2017.

[30] S. Zhao, Y. Wang, Z. Yang, and D. Cai, "Region Mutual Information Loss for Semantic Segmentation," In *NeurIPS*, pp. 11115-11125, 2019.

[31] MedSeg, "COVID-19 CT segmentation dataset," URL https://medicalsegmentation.com/covid19/, 2020.

[32] C. Huang, Y. Wang, X. Li, *et al*, "Clinical Features of Patients Infected with 2019 Novel Coronavirus in Wuhan, China," *The Lancet*, 395(10223), pp. 497–506, 2020.

[33] F. Shan, Y. Gao, J. Wang, *et al*, "Lung Infection Quantification of COVID-19 in CT Images with Deep Learning," *arXiv:2003.04655*, 2020.