

Perspective+ Unet: Enhancing Segmentation with Bi-Path Fusion and Efficient Non-Local Attention for Superior Receptive Fields

Jintong Hu¹, Siyan Chen², Zhiyi Pan¹, Sen Zeng¹, and Wenming Yang¹(✉)

¹ Shenzhen International Graduate School, Tsinghua University, Beijing, China
yang.wenming@sz.tsinghua.edu.cn

² College of Electronics and Information Engineering, Shenzhen University, Shenzhen, China
<https://github.com/tljxyys/Perspective-Unet>

Abstract. Precise segmentation of medical images is fundamental for extracting critical clinical information, which plays a pivotal role in enhancing the accuracy of diagnoses, formulating effective treatment plans, and improving patient outcomes. Although Convolutional Neural Networks (CNNs) and non-local attention methods have achieved notable success in medical image segmentation, they either struggle to capture long-range spatial dependencies due to their reliance on local features, or face significant computational and feature integration challenges when attempting to address this issue with global attention mechanisms. To overcome existing limitations in medical image segmentation, we propose a novel architecture, Perspective+ Unet. This framework is characterized by three major innovations: (i) It introduces a dual-pathway strategy at the encoder stage that combines the outcomes of traditional and dilated convolutions. This not only maintains the local receptive field but also significantly expands it, enabling better comprehension of the global structure of images while retaining detail sensitivity. (ii) The framework incorporates an efficient non-local transformer block, named ENLTB, which utilizes kernel function approximation for effective long-range dependency capture with linear computational and spatial complexity. (iii) A Spatial Cross-Scale Integrator strategy is employed to merge global dependencies and local contextual cues across model stages, meticulously refining features from various levels to harmonize global and local information. Experimental results on the ACDC and Synapse datasets demonstrate the effectiveness of our proposed Perspective+ Unet. The code is available in the supplementary material.

Keywords: Segmentation · Dual-pathway strategy · Efficient Non-local Transformer.

1 Introduction

3D Medical image segmentation plays a crucial role in the realm of diagnostic radiology and surgical planning. Accurate segmentation of anatomical structures from volumetric datasets, such as magnetic resonance imaging (MRI) or

CT, enables the extraction of critical clinical information, facilitating precise interventions, monitoring of disease progression and personalization of treatment strategies [2, 4].

Recently, Convolutional Neural Networks (CNNs) have achieved tremendous success across various domains [19, 18, 3]. Their 3D-focused variants, like V-Net [11] and 3D U-Net [13], play a prominent role in local feature extraction for 3D medical image segmentation. These models, including enhanced versions such as Res-UNet [20] with its residual connections, and U-Net++ [24] along with UNet3+ [5], with their advanced skip pathways and nested designs, excel in detecting detailed local patterns. Despite their proficiency, the inherent limitation of CNNs’ local receptive fields restricts their ability to capture broader spatial relationships essential for a comprehensive understanding of images. To compensate for this, multiple layers are required to expand the coverage, potentially leading to increased model complexity without effectively integrating global context.

Contrastingly, non-local methods expand the receptive fields of segmentation models, harnessing global context over mere local details in image analysis. These strategies focus on spatial relationships of the entire image to improve segmentation efficacy. Emblems of this approach are transformer-based [12, 2] or hybrid transformer models [6, 7, 15], which apply global attention to image regions. Despite their objective to synthesize distant spatial information, these modules risk integrating noise and face quadratic computation costs tied to input size. While containing non-local interactions offers computational relief, it paradoxically trims crucial global insights. Efforts to balance computational efficiency and global context awareness in image segmentation have led to innovations like the Swin-Unet [1, 9] and non-local Attention variants [10, 17], which partition images for blockwise attention calculation before global synthesis. Although this reduces complexity, it also limits the capture of wider spatial information.

In this paper, we propose Perspective+ Unet, a pioneering framework designed to augment the receptive field for enhanced 3D medical image segmentation. Our model incorporates an innovative Bi-Path encoder, utilizing two pathways: one utilizing standard convolutions to capture high-resolution local details, and the other employing dilated convolutions to capture broader contextual information. This configuration addresses issues commonly associated with dilated convolutions, allowing for the synthesis of a rich set of features with varying receptive field sizes, fostering a deeper and wider understanding of the imagery. Then, the extracted features are fed into an Efficient Non-Local Transformer Block (ENLTB), which builds upon the standard non-local attention mechanisms to provide a global perspective with a fraction of the usual computational overhead. The ENLTB effectively condenses long-range spatial dependencies into a compact feature set, enriching subsequent feature maps with globally attentive characteristics. In addition, to enhance the synergy between local and global feature representations, we design the Spatial Cross-Scale Integrator (SCSI) to encourage a coherent integration of information across different stages, ensuring that fine-grained details are preserved while benefiting from macroscopic context-




Fig. 1. The pipeline of the proposed Perspective+ Unet. The model consists of (i) a bi-path CNN based encoder to effectively capture local details and broad contextual information, (ii) a bottleneck composed of Efficient Non-Local Transformer Blocks and Spatial Cross-Scale Integrator for enhanced global perspective, and (iii) a decoder to incorporate both global and local information for generating segmentation results.

tual insights. The output of the SCSI module is incorporated into the decoder via skip connections, ultimately facilitating the generation of segmented images from the decoder output. Our model has been trained and tested on the ACDC and Synapse datasets, achieving competitive levels of performance in terms of both the Dice Similarity Coefficient (DSC) and Hausdorff Distance (HD) metrics.

2 Methods

Our proposed Perspective+ Unet adopts an encoder-bottleneck-decoder configuration, with the specific network workflow illustrated in Figure 1. In this chapter, we will provide a detailed introduction to the three modules within the network, starting with the Bi-Path Residual Block.

2.1 Bi-Path Residual Block (BPRB)

Increasing the size of the receptive field to improve accuracy has been a research focus in the advancement of 3D medical image segmentation. Traditionally, researchers sought to achieve this by stacking global modules, which, while expanding the view, often led to the loss of local region information: the broader receptive field to capture global information risked overlooking crucial local details.

To address this challenge, we have innovated the BPRB, a dual-pathway design that balances local and global information processing to enhance segmentation. One pathway uses dilated convolution to broaden the receptive field,

understanding wider spatial information, but introducing feature discontinuity due to its spacing method. To solve this, BPRB incorporates another pathway with traditional convolution, focusing on capturing detailed features and maintaining local information continuity, thus achieving an optimal balance between global and local information processing. The proposed BPRB can be formulated as:

$$\mathbf{x}_{s+1}^{loc} = \text{BatchNorm}(f_s(\text{BatchNorm}(f_s(\mathbf{x}_s)))) \quad (1)$$

$$\mathbf{x}_{s+1}^{glob} = \text{BatchNorm}(f_s^k(\text{BatchNorm}(f_s^k(\mathbf{x}_s)))) \quad (2)$$

$$\mathbf{x}_{s+1} = \text{Conv}(\text{Concat}[\mathbf{x}_{s+1}^{loc}, \mathbf{x}_{s+1}^{glob}]) \quad (3)$$

where $\mathbf{x}_{s+1}^{loc}$, $\mathbf{x}_{s+1}^{glob}$, $\mathbf{x}_{s+1}$ are local, global, and final features in stage $s+1$. f_s and f_s^k are the convolution and dilated convolution with dilation rate k in stage s .

2.2 Efficient Non-Local Transformer Block (ENLTB)

Capturing global information in input images is crucial for enhancing feature representation, which directly contributes to the effectiveness of segmentation. In response to this, we integrate the ENLTB as a strategic means to reconfigure the feature maps produced by the encoder, aiming to achieve enhanced contextual understanding and representational diversity.

The ENLTB employs the Efficient Non-Local Self-Attention (ENLSA) mechanism, which substantially accelerates processing by substituting the conventional exponential kernel $\exp(\mathbf{Q}^\top \mathbf{K})$, referenced in non-local neural network [16], with a more computationally efficient unbiased estimate σ . Specifically, $\exp(\mathbf{Q}^\top \mathbf{K}) = \sigma(\mathbf{Q})^\top \sigma(\mathbf{K})$, where $\sigma(\cdot)$ can be reparameterized into a linear mapping. The detailed proofs are given below (constants unrelated to variables are omitted for enhanced readability):

$$\begin{aligned} & \exp(\mathbf{Q}^\top \mathbf{K}) \\ &= \exp\left(\|\mathbf{Q} + \mathbf{K}\|^2 / 2\right) \cdot \exp\left(-\left(\|\mathbf{Q}\|^2 + \|\mathbf{K}\|^2\right) / 2\right) \\ &= \exp\left(-\left(\|\mathbf{Q}\|^2 + \|\mathbf{K}\|^2\right) / 2\right) \int \exp\left(\|\mathbf{Q} + \mathbf{K}\|^2 / 2 - \|\phi - (\mathbf{Q} + \mathbf{K})\|^2 / 2\right) d\phi \\ &= \exp\left(-\left(\|\mathbf{Q}\|^2 + \|\mathbf{K}\|^2\right) / 2\right) \int \exp\left(-\|\phi\|^2 / 2 + \phi^\top (\mathbf{Q} + \mathbf{K})\right) d\phi \\ &= \mathbb{E}_{\phi \sim \mathcal{N}(\mathbf{0}_c, \mathbf{I}_c)} \exp\left(\phi^\top (\mathbf{Q} + \mathbf{K}) - \left(\|\mathbf{Q}\|^2 + \|\mathbf{K}\|^2\right) / 2\right) \\ &= \mathbb{E}_{\phi \sim \mathcal{N}(\mathbf{0}_c, \mathbf{I}_c)} \exp\left(-\|\mathbf{Q} - \phi\|^2 / 2\right) \exp\left(-\|\mathbf{K} - \phi\|^2 / 2\right) \\ &= \sigma(\mathbf{Q})^\top \sigma(\mathbf{K}) \end{aligned} \quad (4)$$

where $\mathbf{Q}$ and $\mathbf{K}$ stand for query and key in the attention. $\sigma(\mathbf{x})$ can be reparameterized as product of a multivariate normal variable Φ and $\mathbf{x}$.

Relative to the quadratic computational complexity of the traditional transformer, the ENLSA module only has a complexity of $O(2N)$ for the matrix

projection computation of $\mathbf{Q}$ and $\mathbf{K}$, followed by $O(N)$ for the multiplication between $\sigma(\mathbf{Q})^T \sigma(\mathbf{K})$ and $\mathbf{V}$. This design ensures that the overall computational process is linearly related to the input size N , significantly reducing computational costs.

We have re-engineered the Transformer by replacing its attention mechanism with the ENLSA to construct the ENLTB. This modification streamlines the enhancement of feature representation. The first layer of ENLTB exclusively processes features received from the encoder. Subsequent ENLTB layers, however, aggregate both the features directly relayed from the encoder and those transmitted from the preceding ENLTB layer through patch merging, thereby intricately blending fine-grained details with coarse semantic information. The ENLTB process can be represented as:

$$\mathbf{x}_s^{in} = \begin{cases} \mathbf{x}_s, & \text{if } s = 3 \\ \text{concat}[\mathbf{x}_s, \text{PatchMerging}(\mathbf{x}_{s-1}^{ENLTB})], & \text{else} \end{cases} \quad (5)$$

$$\mathbf{y} = \mathbf{x}_s^{in} + \text{ENLSA}(\text{LN}(\mathbf{x}_s^{in})), \quad \mathbf{x}_s^{ENLTB} = \mathbf{y} + \text{MLP}(\text{LN}(\mathbf{y})) \quad (6)$$

where the definition of $\mathbf{x}_s$ remains consistent with that presented in Section 2.1. $\mathbf{x}_s^{in}$, $\mathbf{x}_s^{ENLTB}$ denote the input and output of ENLTB at stage s . LN and MLP stand for the layer normalization and Multilayer Perceptron, respectively.

The main idea of ENLTB is to employ linear approximations for reducing the computational complexity of non-local attention modules. Under nonlinear transformations, salient regions in feature space exhibit robustness to minor perturbations, retaining their saliency in the approximate representation. This approach enables emphasizing feature points with the most substantial impact on the final task while minimizing the effect on segmentation quality.

2.3 Spatial Cross-Scale Integrator (SCSI)

Collaborative feature interplay is instrumental in the enrichment of the network's interpretative performance, as a result, we design the SCSI module to ensure that detailed image complexities are segmented with both precision and finesse.

The SCSI starts with patch flattening, where feature maps produced by every ENLTB are gathered and merged into a unified sequence. Following this, a transformer processes the sequence, which supports the learning of connections between different features. The refined feature sequences are then carefully mapped back onto their original feature maps for each scale, maintaining their original order of connection. The process of SCSI can be formulated as:

$$\mathbf{x}_{\text{compound}} = \text{concat}[\text{flatten}(\mathbf{x}_s^{ENLTB}), \dots, \text{flatten}(\mathbf{x}_{s+t}^{ENLTB})] \quad (7)$$

$$\mathbf{x}_s^{SCSI}, \dots, \mathbf{x}_{s+t}^{SCSI} = \text{unflatten}(\text{transformer}(\mathbf{x}_{\text{compound}})) \quad (8)$$

where $\mathbf{x}_s^{ENLTB}$ and $\mathbf{x}_s^{SCSI}$ denote the output of ENLTB and SCSI at stage s , respectively. concat represents the concatenation operation.

Table 1. Segmentation accuracy of different methods on the Synapse multi-organ CT dataset. The best results are shown in **bold**.

Methods	Average		Aorta	Gallbladder	Kidney(L)	Kidney(R)	Liver	Pancreas	Spleen	Stomach
	DSC $\uparrow$	HD $\downarrow$								
U-Net [13]	76.85	39.70	89.07	69.72	77.77	68.60	93.43	53.98	86.67	75.58
R50 Att-UNet [2]	75.57	36.97	55.92	63.91	79.20	72.71	93.56	49.37	87.19	74.95
Att-UNet [12]	77.77	36.02	89.55	68.88	77.98	71.11	93.57	58.04	87.30	75.75
R50 ViT [2]	71.29	32.87	73.73	55.13	75.80	72.20	91.51	45.99	81.99	73.95
TransUnet [2]	77.48	31.69	87.23	63.13	81.87	77.02	94.08	55.86	85.08	75.62
SwinUNet [1]	79.12	21.55	85.47	66.53	83.28	79.61	94.29	56.58	90.66	76.60
AFTer-UNet [21]	81.02	-	90.91	64.81	87.90	85.30	92.20	63.54	90.99	72.48
ScaleFormer [6]	82.86	16.81	88.73	74.97	86.36	83.31	95.12	64.85	89.40	80.14
MISSFormer [7]	81.96	18.20	86.99	68.65	85.21	82.00	94.41	65.67	91.92	80.81
FCT [14]	83.53	-	89.85	72.73	88.45	86.60	95.62	66.25	89.77	79.42
MSAANet [22]	82.85	18.54	89.40	73.20	84.31	78.53	95.10	68.85	91.60	81.78
Ours	84.63	11.74	89.38	70.80	87.57	85.78	95.30	70.71	94.41	83.06

The SCSI module equips the Perspective+ Unet with the capability to assimilate complementary information extending across scales, thereby simplifying resolution of the ambiguities linked with complex scale variations.

3 Experiments

3.1 Datasets and Evaluation

Synapse: The Synapse dataset, from MICCAI 2015 Multi-Atlas Abdomen Labeling Challenge, contains 30 abdominal 3D CT scans. Eighteen cases are randomly selected for training purposes and the remaining 12 are allocated for testing. Our methodology is evaluated using the average Dice Similarity Coefficient (DSC) and the average Hausdorff Distance (HD) for eight abdominal organs.

ACDC: The ACDC Challenge assembled MRI scan results from various patients. Each patient’s scan is manually annotated with three labels: Left Ventri-

Table 2. Segmentation accuracy of different methods on the ACDC dataset. The best results are shown in **bold**.

Methods	DSC $\uparrow$	RV	Myo	LV
R50 U-Net [2]	87.55	87.10	80.63	94.92
R50 Att-UNet [2]	86.75	87.58	79.20	93.47
R50 ViT [2]	87.57	86.07	81.88	94.75
TransUnet [2]	89.71	88.86	84.53	95.73
SwinUNet [1]	90.00	88.55	85.62	95.83
ScaleFormer [6]	90.17	87.33	88.16	95.04
UNETR [4]	88.61	85.29	86.52	94.02
MCTE [8]	91.31	89.14	89.51	95.27
MISSFormer [7]	91.19	89.85	88.38	95.34
nnFormer [23]	92.06	90.94	89.58	95.65
Ours	92.54	90.92	90.49	96.20




Fig. 2. Visualized segmentation results of different methods on the Synapse multi-organ CT dataset. Our method (the last column) exhibits the smoothest boundaries and the most accurate segmentation outcomes.

cle (LV), Right Ventricle (RV), and Myocardium (MYO). We randomly select 70 cases for training, 10 cases for validation, and 20 cases for testing. The methodology is evaluated using the average DSC.

3.2 Experiment Details

In our Perspective+ Unet, the encoder or decoder requires no pre-training and is entirely trained from scratch. To mitigate the risk of overfitting, a comprehensive suite of data augmentation techniques is employed, including probabilistic image flipping, additive Gaussian noise, and a variety of affine transformations such as scaling, rotation, shearing, and translation. Training is conducted on the Synapse and ACDC datasets, with a regimen of 600 and 1000 epochs, respectively, and a batch size of 12. Images are resized to 224×224 pixels. To optimize the learning process, a learning rate of 0.05 is implemented and the weight decay is set at 0.0001 using the SGD optimizer. All training and testing procedures are performed on an NVIDIA RTX A5000 GPU.

3.3 Quantitative and qualitative segmentation results

To validate the efficacy of Perspective+ Unet, we compare it against 11 advanced methods on the Synapse dataset. In the comparison between 2D and 3D methodologies, our analysis employs a slice-wise evaluation approach as previous works. Highlighted in Table 1, our method excels in DSC at 84.63% and HD at 11.74%, showing marked improvements over MSAANet, of 1.78% in DSC and




Fig. 3. Visualization of attention heat maps from the intermediate layers of the network. Highlighting areas are closely aligned with segmentation labels, demonstrating our Perspective+ Unet’s accuracy in feature identification and localization.

a reduction in HD by 6.8mm. This advancement is credited to the BPRB and ENLTB modules, enhancing the model’s receptive field. Further validation on the ACDC dataset confirmed Perspective+ Unet’s superior versatility and robustness in capturing critical representations, as shown in Table 2. The segmentation results of different methods on the Synapse dataset are shown in Figure 2.

In our comparative analysis, ENLSA is evaluated against another non-local approach in terms of parameter count and number of floating point operations (FLOPs). The results presented in Table 3 demonstrate that ENLSA significantly outperforms NLSA in efficiency.

Furthermore, to further validate the effectiveness of the model, this study employs visualization of the outputs from the intermediate layers of the network, as well as the creation of attention heat maps. As demonstrated in Figure 3, the highlighted areas within the heat map closely align with the segmentation labels. This not only proves the model’s capability to efficiently identify key features from the input images, but also showcases its accuracy in feature localization.

Table 3. Efficiency comparison of ENLSA. -m is the number of input channel. The input resolution is set as 224×224 .

Module	FLOPs↓	Params↓
NLSA-64 [10]	2.06G	0.04M
ENLSA-64	0.62G	0.01M
NLSA-256	32.88G	0.66M
ENLSA-256	9.87G	0.20M

Table 4. Ablation study on the impact of modules. BPRB: Bi-Path Residual Block, ENLTB: Efficient Non-Local Transformer Block, SCSi: Spatial Cross-Scale Integrator.

Model	BPRB	SCSi	ENLTB	DSC↑	HD↓
Setting 1	✗	✗	✗	84.04	16.63
Setting 2	✓	✗	✗	83.36	14.70
Setting 3	✓	✓	✗	83.92	13.94
Setting 4	✓	✓	✓	84.63	11.74

3.4 Ablation Study

In our proposed Perspective+ Unet architecture, we perform an in-depth ablation study on the added modules: BPRB, SCSi, and ENLTB, to assess their impact on model performance. Data from Table 4 reveal that the incorporation of the BPRB module alone slightly reduces DSC by -0.68%, but significantly

improves precision by reducing HD by 1.93mm. Further integration of the SCS module with BPRB enhances both DSC by 0.56% and reduces HD by 0.76mm, evidencing the SCS module’s role in increasing reconstruction accuracy and precision. Employing all three modules achieves optimal outcomes, with a DSC of 84.63% and an HD of 11.74mm, highlighting each module’s critical contribution to enhancing model performance.

4 Conclusion

In conclusion, Perspective+ Unet is distinguished by its innovative approach to enhancing spatial perception in 3D medical image segmentation. We design a dual-channel strategy that combines traditional and dilated convolutions with a novel transformer block and spatial cross-scale integrator. This architecture ensures comprehensive integration of image features across different stages and effectively handles both global structures and fine-grained parts, significantly enhancing the receptive field. Our experimental results demonstrate that Perspective+ Unet exhibits good segmentation accuracy in the ACDC and Synapse datasets, highlighting its significant potential in medical analysis.

Acknowledgments. This project was supported by National Key Research and Development Program of China(Grant No.2023YFB4302200).

Disclosure of Interests. The authors declare no competing interests in the paper.

References

1. Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation. In: Karlinsky, L., Michaeli, T., Nishino, K. (eds.) *Computer Vision – ECCV 2022 Workshops*. pp. 205–218. Springer Nature Switzerland, Cham (2023)
2. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: Transunet: Transformers make strong encoders for medical image segmentation (2021)
3. Ding, X., Zhang, X., Ma, N., Han, J., Ding, G., Sun, J.: Repvgg: Making vgg-style convnets great again. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 13728–13737 (2021). <https://doi.org/10.1109/CVPR46437.2021.01352>
4. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 1748–1758 (2022)
5. Huang, H., Lin, L., Tong, R., Hu, H., Zhang, Q., Iwamoto, Y., Han, X., Chen, Y.W., Wu, J.: Unet 3+: A full-scale connected unet for medical image segmentation. In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1055–1059 (2020)

6. Huang, H., Xie, S., Lin, L., Iwamoto, Y., Han, X.H., Chen, Y.W., Tong, R.: Scaleformer: Revisiting the transformer-based backbones from a scale-wise perspective for medical image segmentation. In: Raedt, L.D. (ed.) Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22. pp. 964–971. International Joint Conferences on Artificial Intelligence Organization (7 2022). <https://doi.org/10.24963/ijcai.2022/135>, <https://doi.org/10.24963/ijcai.2022/135>, main Track
7. Huang, X., Deng, Z., Li, D., Yuan, X., Fu, Y.: Missformer: An effective transformer for 2d medical image segmentation. *IEEE Transactions on Medical Imaging* **42**(5), 1484–1494 (2023). <https://doi.org/10.1109/TMI.2022.3230943>
8. Li, J.: Mcte: Marrying convolution and transformer efficiently for end-to-end medical image segmentation. In: 2023 IEEE International Conference on Image Processing (ICIP). pp. 1100–1104 (2023). <https://doi.org/10.1109/ICIP49359.2023.10222041>
9. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9992–10002 (2021). <https://doi.org/10.1109/ICCV48922.2021.00986>
10. Mei, Y., Fan, Y., Zhou, Y.: Image super-resolution with non-local sparse attention. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3516–3525 (2021). <https://doi.org/10.1109/CVPR46437.2021.00352>
11. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 565–571 (2016). <https://doi.org/10.1109/3DV.2016.79>
12. Oktay, O., Schlemper, J., Folgoc, L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N., Kainz, B., Glocker, B., Rueckert, D.: Attention u-net: Learning where to look for the pancreas (04 2018)
13. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. pp. 234–241. Springer International Publishing, Cham (2015)
14. Tragakis, A., Kaul, C., Murray-Smith, R., Husmeier, D.: The fully convolutional transformer for medical image segmentation. In: 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 3649–3658 (2023). <https://doi.org/10.1109/WACV56688.2023.00365>
15. Wang, C., Zhang, D., Ge, R.: Eye-guided dual-path network for multi-organ segmentation of abdomen. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. pp. 23–32. Springer Nature Switzerland, Cham (2023)
16. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7794–7803. IEEE Computer Society, Los Alamitos, CA, USA (jun 2018). <https://doi.org/10.1109/CVPR.2018.00813>, <https://doi.ieeecomputersociety.org/10.1109/CVPR.2018.00813>
17. Xia, B., Hang, Y., Tian, Y., Wenming, Y., Liao, Q., Zhou, J.: Efficient non-local contrastive attention for image super-resolution. *Proceedings of the AAAI Conference on Artificial Intelligence* **36**, 2759–2767 (06 2022). <https://doi.org/10.1609/aaai.v36i3.20179>

18. Xia, B., Tian, Y., Zhang, Y., Hang, Y., Yang, W., Liao, Q.: Meta-learning-based degradation representation for blind super-resolution. *IEEE Transactions on Image Processing* **32**, 3383–3396 (2023). <https://doi.org/10.1109/TIP.2023.3283922>
19. Xia, B., Zhang, Y., Wang, Y., Tian, Y., Yang, W., Timofte, R., Gool, L.V.: Basic binary convolution unit for binarized image restoration network. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=h8T5dZWTZ-Z>
20. Xiao, X., Lian, S., Luo, Z., Li, S.: Weighted res-unet for high-quality retina vessel segmentation. In: *2018 9th International Conference on Information Technology in Medicine and Education (ITME)*. pp. 327–331 (2018)
21. Yan, X., Tang, H., Sun, S., Ma, H., Kong, D., Xie, X.: After-unet: Axial fusion transformer unet for medical image segmentation. In: *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 3270–3280. IEEE Computer Society, Los Alamitos, CA, USA (jan 2022). <https://doi.org/10.1109/WACV51458.2022.00333>, <https://doi.ieeecomputersociety.org/10.1109/WACV51458.2022.00333>
22. Zeng, H., Shan, X., Feng, Y., Wen, Y.: Msaanet: Multi-scale axial attention network for medical image segmentation. In: *2023 IEEE International Conference on Multimedia and Expo (ICME)*. pp. 2291–2296 (2023). <https://doi.org/10.1109/ICME55011.2023.00391>
23. Zhou, H.Y., Guo, J., Zhang, Y., Han, X., Yu, L., Wang, L., Yu, Y.: nnformer: Volumetric medical image segmentation via a 3d transformer. *IEEE Transactions on Image Processing* **32**, 4036–4045 (2023). <https://doi.org/10.1109/TIP.2023.3293771>
24. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. *IEEE Transactions on Medical Imaging* **39**(6), 1856–1867 (2020)