



# HHS Public Access

Author manuscript

*Med Image Comput Assist Interv.* Author manuscript; available in PMC 2024 January 31.

Published in final edited form as:

*Med Image Comput Assist Interv.* 2023 October ; 14222: 736–746.

doi:10.1007/978-3-031-43898-1\_70.

## Laplacian-Former: Overcoming the Limitations of Vision Transformers in Local Texture Detection

Reza Azad<sup>1</sup>, Amirhossein Kazerouni<sup>2</sup>, Babak Azad<sup>3</sup>, Ehsan Khodapanah Aghdam<sup>4</sup>, Yury Velichko<sup>5</sup>, Ulas Bagci<sup>5</sup>, Dorit Merhof<sup>6</sup>

<sup>1</sup>Faculty of Electrical Engineering and Information Technology, RWTH Aachen University, Germany

<sup>2</sup>School of Electrical Engineering, Iran University of Science and Technology, Iran

<sup>3</sup>South Dakota State University, Brookings, USA

<sup>4</sup>Department of Electrical Engineering, Shahid Beheshti University, Iran

<sup>5</sup>Department of Radiology, Northwestern University, Chicago, USA

<sup>6</sup>Faculty of Informatics and Data Science, University of Regensburg, Germany

### Abstract

Vision Transformer (ViT) models have demonstrated a breakthrough in a wide range of computer vision tasks. However, compared to the Convolutional Neural Network (CNN) models, it has been observed that the ViT models struggle to capture high-frequency components of images, which can limit their ability to detect local textures and edge information. As abnormalities in human tissue, such as tumors and lesions, may greatly vary in structure, texture, and shape, high-frequency information such as texture is crucial for effective semantic segmentation tasks. To address this limitation in ViT models, we propose a new technique, Laplacian-Former, that enhances the self-attention map by adaptively re-calibrating the frequency information in a Laplacian pyramid. More specifically, our proposed method utilizes a dual attention mechanism via efficient attention and frequency attention while the efficient attention mechanism reduces the complexity of self-attention to linear while producing the same output, selectively intensifying the contribution of shape and texture features. Furthermore, we introduce a novel efficient enhancement multi-scale bridge that effectively transfers spatial information from the encoder to the decoder while preserving the fundamental features. We demonstrate the efficacy of Laplacian-former on multi-organ and skin lesion segmentation tasks with +1.87% and +0.76% dice scores compared to SOTA approaches, respectively. Our implementation is publically available at GitHub.

### Keywords

Deep Learning; Texture; Segmentation; Laplacian Transformer

## 1 Introduction

The recent advancements in Transformer-based models have revolutionized the field of natural language processing and have also shown great promise in a wide range of computer vision tasks [5,15]. As a notable example, the Vision Transformer (ViT) model utilizes Multi-head Self-Attention (MSA) blocks to globally model the interactions between semantic tokens created by treating local image patches as individual elements [7]. This approach stands in contrast to CNNs, which hierarchically increase their receptive field from local to global to capture a global semantic representation. Nevertheless, recent studies [25,3] have shown that ViT models struggle to capture high-frequency components of images, which can limit their ability to detect local textures and it is vital for many diagnostic and prognostic tasks. This weakness in local representation can be attributed to the way in which ViT models process images. ViT models split an image into a sequence of patches and model their dependencies using a self-attention mechanism, which may not be as effective as the convolution operation used in CNN models in extracting local features within receptive fields. This difference in how ViT and CNN process images may explain the superior performance of CNN models in local feature extraction [9,1]. Innovative approaches have been proposed in recent years to address the insufficient local texture representation within Transformer models. One such approach is the integration of CNN and ViT features through complementary methods, aimed at seamlessly blending the strengths of both in order to compensate for any shortcomings in local representation [5,16].

### Transformers as a Complement to CNNs:

TransUNet [5] is one of the earliest approaches incorporating the Transformer layers into the CNN bottle-neck to model both local and global dependency using the combination of CNN and ViT models. Heidari et al. [12] proposed a novel solution called HiFormer, which leverages a Swin Transformer module and a CNN-based encoder to generate two multi-scale feature representations, which are then integrated via a Double-Level Fusion module. UNETR [11] used a Transformer to create a powerful encoder with a CNN decoder for 3D medical image segmentation. By bridging the CNN-based encoder and decoder with the Transformer, CoTr [31], and TransBTS [27], the segmentation performance in low-resolution stages was improved. Despite these advances, there remain some limitations in these methods such as computationally inefficiency (e.g., TransUNet model), the requirement of a heavy CNN backbone (e.g., HiFormer), and the lack of consideration for multi-scale information. These limitations have resulted in less effective network learning results in the field of medical image segmentation.

### New Attention Models:

The redesign of the self-attention mechanism within pure Transformer models is another method aiming to augment feature representation to enhance the local feature representation ultimately [14]. In this direction, Swin-Unet [4] utilizes a linear computational complexity Swin Transformer [18] block in a U-shaped structure as a multi-scale backbone. MISS-Former [13] besides exploring the Efficient Transformer [30] counterpart to diminish the parameter overflow of vision transformers, applies a non-invertible down-sampling operation on input blocks transformer to reduce the parameters. D-Former [29] is a pure

transformer-based pipeline that comprises a double attention module to capture locally fine-grained attention and interaction with different units in a dilated manner through its mechanism.

### Drawbacks of Transformers:

Recent research has revealed that traditional self-attention mechanisms, while effective in addressing local feature discrepancies, have a tendency to overlook important high-frequency information such as texture and edge details [26]. This is especially problematic for tasks like tumor detection, cancer-type identification through radiomics analysis, as well as treatment response assessment, where abnormalities often manifest in texture. Moreover, self-attention mechanisms have a quadratic computational complexity and may produce redundant features[23].

### Our Contributions:

- ① We propose Laplacian-Former, a novel approach that includes new efficient attention (EF-ATT) consisting of two sub-attention mechanisms: *efficient attention* and *frequency attention*. The efficient attention mechanism reduces the complexity of self-attention to linear while producing the same output. The frequency attention mechanism is modeled using a Laplacian pyramid to emphasize each frequency information's contribution selectively. Then, a parametric frequency attention fusion strategy to balance the importance of shape and texture features by recalibrating the frequency features. These two attention mechanisms work in parallel.
- ② We also introduce a novel efficient enhancement multi-scale bridge that effectively transfers spatial information from the encoder to the decoder while preserving the fundamental features.
- ③ Our method not only alleviates the problem of the traditional self-attention mechanism mentioned above, but also it surpasses all its counterparts in terms of different evaluation metrics for the tasks of medical image segmentation.

## 2 Methods

In our proposed network, illustrated in Figure 1, taking an input image  $X \in \mathbb{R}^{H \times W \times C}$  with spatial dimensions  $H$  and  $W$ , and  $C$  channels, it is first passed through a patch embedding module to obtain overlapping patch tokens of size  $4 \times 4$  from the input image. The proposed model comprises four encoder blocks, each containing two efficient enhancement Transformer layers and a patch merging layer that downsamples the features by merging  $2 \times 2$  patch tokens and increasing the channel dimension. The decoder is composed of three efficient enhancement Transformer blocks and four patch-expanding blocks, followed by a segmentation head to retrieve the final segmentation map. Laplacian-Former then employs a novel efficient enhancement multi-scale bridge to capture local and global correlations of different scale features and effectively transfer the underlying features from the encoder to the decoder.

## 2.1 Efficient Enhancement Transformer Block

In medical imaging, it is important to distinguish different structures and tissues, especially when tissue boundaries are ill-defined. This is often the case for accurate segmentation of small abnormalities, where high-frequency information plays a critical role in defining boundaries by capturing both textures and edges. Inspired by this, we propose an Efficient Enhancement Transformer Block that incorporates an Efficient Frequency Attention (EF-ATT) mechanism to capture contextual information of an image while recalibrating the representation space within an attention mechanism and recovering high-frequency details.

Our efficient enhancement Transformer block first takes a LayerNorm (LN) from the input  $x$ . Then it applies the EF-ATT mechanism to capture contextual information and selectively include various types of frequency information while using the Laplacian pyramid to balance the importance of shape and texture features. Next,  $x$  and diversity-enhanced shortcuts are added to the output of the attention mechanism to increase the diversity of features. It is proved in [24] that as Transformers become deeper, their features become less varied, which restrains their representation capacity and prevents them from attaining optimal performance. To address this issue, we have implemented an *augmented shortcut* method from [19], a Diversity-Enhanced Shortcut (DES), employing a Kronecker decomposition-based projection. This approach involves inserting additional paths with trainable parameters alongside the original shortcut  $x$ , which enhances feature diversity and improves performance while requiring minimal hardware resources. Finally, we apply LayerNorm and MiX-FFN [10] to the resulting feature representation to enhance its power. This final step completes our efficient enhancement Transformer block, as illustrated in Figure 2.

## 2.2 Efficient Frequency Attention (EF-ATT)

The traditional self-attention block computes the attention score  $S$  using query ( $\mathbf{Q}$ ) and key ( $\mathbf{K}$ ) values, normalizes the result using Softmax, and then multiplies the normalized attention map with value ( $\mathbf{V}$ ):

$$S(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V}, \quad (1)$$

where  $d_k$  is the embedding dimension. One of the main limitations of the dot-product mechanism is that it generates redundant information, resulting in unnecessary computational complexity. Shen et al. [23] proposed to represent the context more effectively by reducing the computational burden from  $\mathcal{O}(n^2)$  to linear form  $\mathcal{O}(d^2n)$ :

$$E(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \rho_q(\mathbf{Q})\left(\rho_k(\mathbf{K})^T\mathbf{V}\right). \quad (2)$$

Their approach involves applying the Softmax function ( $\rho$ ) to the key and query vectors to obtain normalized scores and formulating the global context by multiplying the key and value matrix. They demonstrate that efficient attention  $E$  can provide an equivalent representation of self-attention while being computationally efficient. By adopting this approach, we can alleviate the issues of feature redundancy and computational complexity associated with self-attention.

Wang et al. [26] explored another major limitation of the self-attention mechanism, where they demonstrated through theoretical analysis that self-attention operates as a low-pass filter that erases high-frequency information, leading to a loss of feature expressiveness in the model's deep layers. Authors found that the Softmax operation causes self-attention to keep low-frequency information and loses its fine details. Motivated by this, we propose a new frequency recalibration technique to address the limitations of self-attention, which only focuses on low-frequency information (which contains shape information) while ignoring the higher frequencies that carry texture and edge information. First, we construct a Laplacian pyramid to determine the different frequency levels of the feature maps. The process begins by extracting  $(L + 1)$  Gaussian representations from the encoded feature using different variance values of the Gaussian function:

$$\mathbf{G}_l(\mathbf{X}) = \mathbf{X} * \frac{1}{\sigma_l \sqrt{2\pi}} e^{-\frac{i^2 + j^2}{2\sigma_l^2}}, \quad (3)$$

where  $\mathbf{X}$  refers to the input feature map,  $(i, j)$  corresponds to the spatial location within the encoded feature map, the variable  $\sigma_l$  denotes the variance of the Gaussian function for the  $l$ -th scale, and the symbol  $*$  represents the convolution operator. The pyramid is then built by subtracting the  $l$ -th Gaussian function ( $\mathbf{G}_l$ ) output from the  $(l + 1)$ -th output ( $\mathbf{G}_l - \mathbf{G}_{l+1}$ ) to encode frequency information at different scales. The Laplacian pyramid is composed of multiple levels, each level containing distinct types of information. To ensure a balanced distribution of low and high-frequency information in the model, it is necessary to efficiently aggregate the features from all levels of the frequency domain. Hence, we present frequency attention that involves multiplying the key and value of each level ( $\mathbf{X}_l$ ) to calculate the attention score and then fuses the resulting attention scores of all levels using a fusion module, which performs summation. The resulting attention score is multiplied by Query ( $\mathbf{Q}$ ) to obtain the final frequency attention result, which subsequently concatenates with the efficient attention result and applies the depth-wise convolution with the kernel size of  $2 \times 1 \times 1$  in order to aggregate both information and recalibrate the feature map, thus allowing for the retrieval of high-frequency information.

### 2.3 Efficient Enhancement Multi-scale Bridge

It is widely known that effectively integrating multi-scale information can lead to improved performance [13]. Thus, we introduce the Efficient Enhancement Multi-scale Bridge as an alternative to simply concatenating the features from the encoder and decoder layers. The proposed bridge, depicted in Figure 1, delivers spatial information to each decoder layer,

enabling the recovery of intricate details while generating output segmentation masks. In this approach, we aim to calculate the efficient attention mechanism for each level and fuse the multi-scale information in their context; thus, it is important that all levels' embedding dimension is of the same size. Therefore, in order to calculate the global context ( $G_i$ ), we parametrize the query and value of each level using a convolution  $1 \times 1$  where it gets the size of  $mC$  and outputs  $C$ , where  $m$  equals 1, 2, 5, and 8 for the first to fourth levels, respectively. We multiply the new key and value to each other to attain the global context. We then use a summation module to aggregate the global context of all levels and reshape the query for matrix multiplication with the augmented global context. Taking the second level with the dimension of  $\frac{H}{8} \times \frac{W}{8} \times 2C$ , the key and value are mapped to  $(\frac{H}{8} \frac{W}{8}) \times C$ , and the query to  $(2\frac{H}{8} \frac{W}{8}) \times C$ . The augmented global context with the shape of  $C \times C$  is then multiplied by the query, resulting in an enriched feature map with the shape of  $(2\frac{H}{8} \frac{W}{8}) \times C$ . We reshape the obtained feature map into  $\frac{H}{8} \times \frac{W}{8} \times 2C$  and feed it through an LN and MiX-FFN module with a skip connection to empower the feature representations. The resulting output is combined with the expanded feature map, and then projected using a linear layer onto the same size as the encoder block corresponding to that level.

### 3 Results

Our proposed technique was developed using the PyTorch library and executed on a single RTX 3090 GPU. A batch size of 24 and a stochastic gradient descent algorithm with a base learning rate of 0.05, a momentum of 0.9, and a weight decay of 0.0001 was utilized during the training process, which was carried out for 400 epochs. For the loss function, we used both cross-entropy and Dice losses ( $Loss = \gamma \cdot L_{dice} + (1 - \gamma) \cdot L_{ce}$ ),  $\gamma$  set to 0.6 empirically.

#### Datasets:

We tested our model using the *Synapse* dataset [17], which comprises 30 cases of contrast-enhanced abdominal clinical CT scans (a total of 3,779 axial slices). Each CT scan consists of 85 ~ 198 slices of the in-plane size of  $512 \times 512$  and has annotations for eight different organs. We followed the same preferences for data preparation analogous to [5]. We also followed [2,8] experiments to evaluate our method on the ISIC 2018 skin lesion dataset [6] with 2,694 images.

#### Synapse Multi-Organ Segmentation:

Table 1 presents a comparison of our proposal with previous SOTA methods using the DSC and HD metrics across eight abdominal organs. Laplacian-Former clearly outperforms SOTA CNN-based methods. We extensively evaluated EfficientFormer (EffFormer) plus another drift of Laplacian-Former without utilizing the bridge connections to endorse the superiority of Laplacian-Former. Laplacian-Former exhibits superior learning ability on the Dice score metric compared to other transformer-based models, achieving an increase of +1.59% and +2.77% in Dice scores compared to HiFormer and Swin-Unet, respectively. Figure 3 illustrates a qualitative result of our method for different organ segmentation, specifically we can observe that the LalacianFormer produces a precise

boundary segmentation on Gallbladder, Liver, and Stomach organs. It is noteworthy to mention that our pipeline, as a pure transformer-based architecture trained from scratch without pretraining weights, outperforms all previously presented network architectures.

### **Skin Lesion Segmentation:**

Table 2a shows the comparison results of our proposed method, Laplacian-Former, against leading methods on the skin lesion segmentation benchmark. Our approach outperforms other competitors across most evaluation metrics, indicating its excellent generalization ability across different datasets. In particular, our approach performs better than hybrid methods such as TMU-Net [19] and pure transformer-based methods such as Swin-Unet [4]. Our method achieves superior performance by utilizing the frequency attention in a pyramid scale to model local textures. Specifically, our frequency attention emphasizes the fine details and texture characteristics that are indicative of skin lesion structures and amplifies regions with significant intensity variations, thus accentuating the texture patterns present in the image and resulting in better performance. In addition, we provided the spectral response of LaplacianFormer vs. Standard Transformer in identical layers in Table 2b. It is evident Standard design frequency response in deep layers of structure attenuates more than the LaplacianFormer, which is a visual endorsement of the capability of LaplacianFormer for its ability to preserve high-frequency details. The supplementary provides more visualization results.

## **4 Conclusion**

In this paper, we introduce Laplacian-Former, a novel standalone transformer-based U-shaped architecture for medical image analysis. Specifically, we address the transformer's inability to capture local context as high-frequency details, e.g., edges and boundaries, by developing a new design within a scaled dot attention block. Our pipeline benefits the multi-resolution Laplacian module to compensate for the lack of frequency attention in transformers. Moreover, while our design takes advantage of the efficiency of transformer architectures, it keeps the parameter numbers low.

## **Supplementary Material**

Refer to Web version on PubMed Central for supplementary material.

## **References**

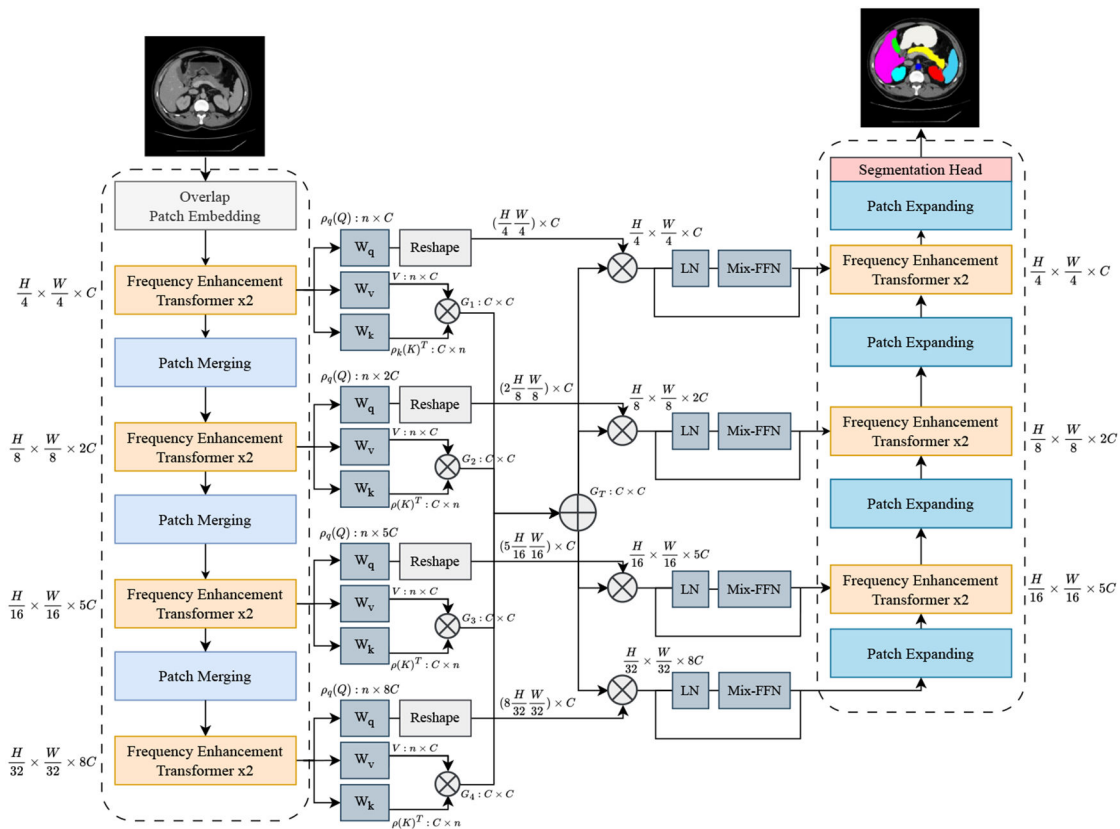
1. Azad R, Fayjie AR, Kauffmann C, Ben Ayed I, Pedersoli M, Dolz J: On the texture bias for few-shot cnn segmentation. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2674–2683 (2021)
2. Azad R, Heidari M, Shariatnia M, Aghdam EK, Karimijafarbigloo S, Adeli E, Merhof D: Transdeeplab: Convolution-free transformer-based deeplab v3+ for medical image segmentation. In: Predictive Intelligence in Medicine. pp. 91–102. Springer Nature Switzerland (2022)
3. Bai J, Yuan L, Xia ST, Yan S, Li Z, Liu W: Improving vision transformers by revisiting high-frequency components. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV. pp. 1–18. Springer (2022)



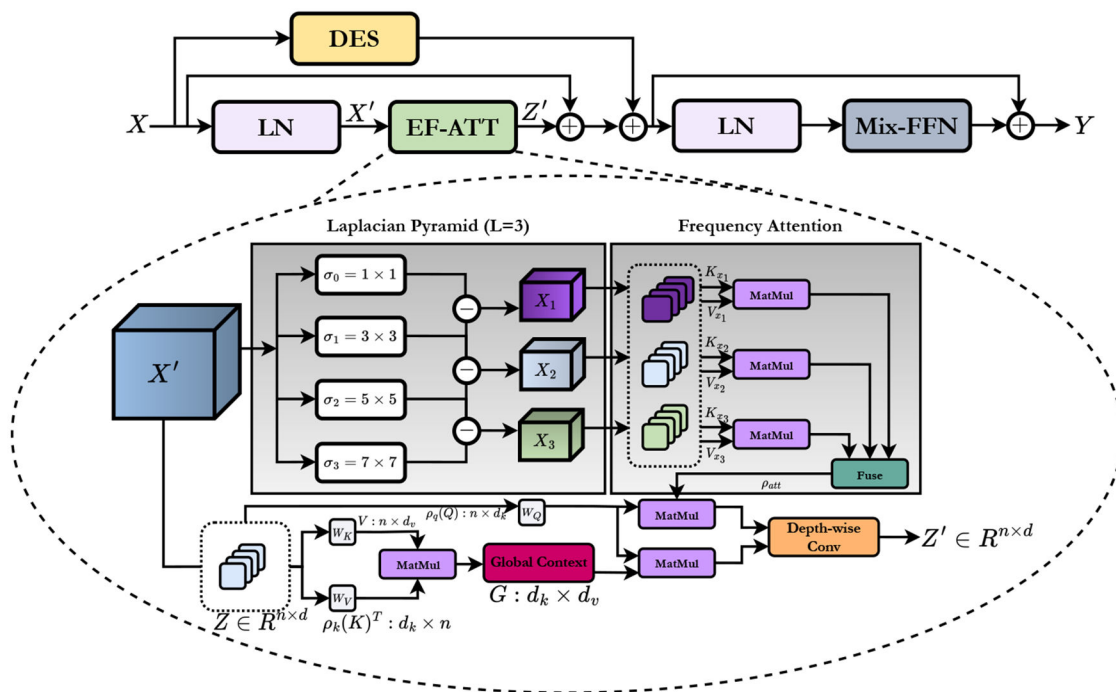
4. Cao H, Wang Y, Chen J, Jiang D, Zhang X, Tian Q, Wang M: Swin-UNET: Unet-like pure transformer for medical image segmentation. In: Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III. pp. 205–218. Springer (2023)
5. Chen J, Lu Y, Yu Q, Luo X, Adeli E, Wang Y, Lu L, Yuille AL, Zhou Y: TransUNet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306 (2021)
6. Codella N, Rotemberg V, Tschandl P, Celebi ME, Dusza S, Gutman D, Helba B, Kalloo A, Liopyris K, Marchetti M, et al. : Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). arXiv preprint arXiv:1902.03368 (2019)
7. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, Uszkoreit J, Houlsby N: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=YicbFdNTTy>
8. Eskandari S, Lumpp J: Inter-scale dependency modeling for skin lesion segmentation with transformer-based networks. In: Medical Imaging with Deep Learning, short paper track (2023)
9. Geirhos R, Rubisch P, Michaelis C, Bethge M, Wichmann FA, Brendel W: Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: International Conference on Learning Representations
10. Gu J, Kwon H, Wang D, Ye W, Li M, Chen YH, Lai L, Chandra V, Pan DZ: Multi-scale high-resolution vision transformer for semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12094–12103 (2022)
11. Hatamizadeh A, Tang Y, Nath V, Yang D, Myronenko A, Landman B, Roth HR, Xu D: Unetr: Transformers for 3d medical image segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 574–584 (2022)
12. Heidari M, Kazerouni A, Soltany M, Azad R, Aghdam EK, Cohen-Adad J, Merhof D: Hiformer: Hierarchical multi-scale representations using transformers for medical image segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 6202–6212 (2023)
13. Huang X, Deng Z, Li D, Yuan X, Fu Y: Missformer: An effective transformer for 2d medical image segmentation. IEEE Transactions on Medical Imaging pp. 1–1 (2022). 10.1109/TMI.2022.3230943
14. Karimijafarbigloo S, Azad R, Kazerouni A, Ebadollahi S, Merhof D: Mmcformer: Missing modality compensation transformer for brain tumor segmentation. In: Medical Imaging with Deep Learning (2023)
15. Karimijafarbigloo S, Azad R, Kazerouni A, Merhof D: Ms-former: Multi-scale self-guided transformer for medical image segmentation. In: Medical Imaging with Deep Learning (2023)
16. Karimijafarbigloo S, Azad R, Merhof D: Self-supervised few-shot learning for semantic segmentation: An annotation-free approach. In: MICCAI 2023 workshop (2023)
17. Landman B, Xu Z, Igelsias J, Styner M, Langerak T, Klein A: Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge. In: Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge. vol. 5, p. 12 (2015)
18. 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: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
19. Reza A, Moein H, Yuli W, Dorit M: Contextual attention network: Transformer meets u-net. arXiv preprint arXiv:2203.01932 (2022)
20. 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)
21. 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)
22. Schlemper J, Oktay O, Schaap M, Heinrich M, Kainz B, Glocker B, Rueckert D: Attention gated networks: Learning to leverage salient regions in medical images. Medical image analysis 53, 197–207 (2019) [PubMed: 30802813]



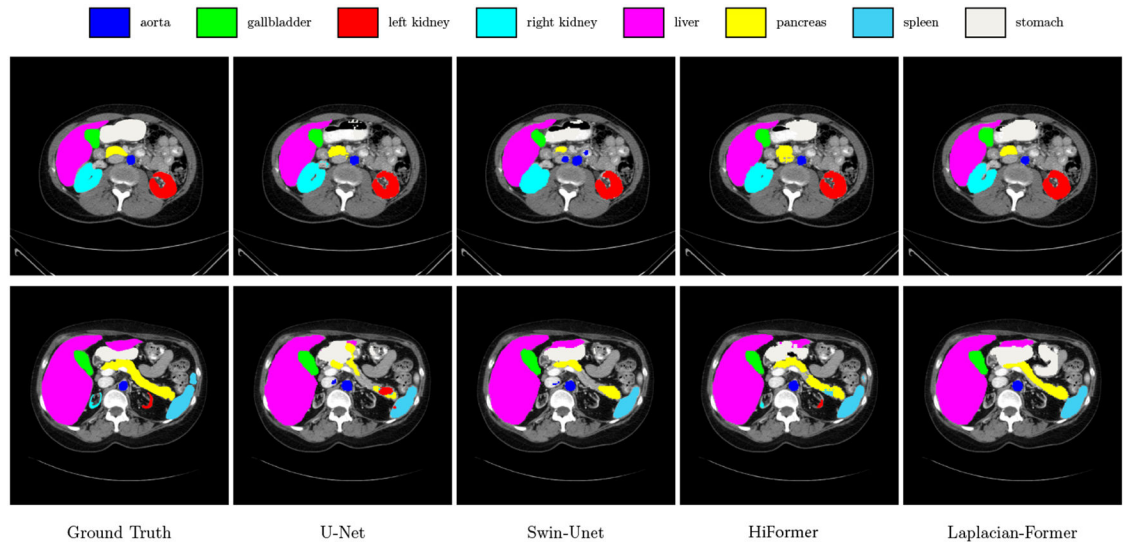
23. Shen Z, Zhang M, Zhao H, Yi S, Li H: Efficient attention: Attention with linear complexities. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 3531–3539 (2021)
24. Tang Y, Han K, Xu C, Xiao A, Deng Y, Xu C, Wang Y: Augmented shortcuts for vision transformers. *Advances in Neural Information Processing Systems* 34, 15316–15327 (2021)
25. Wang P, Zheng W, Chen T, Wang Z: Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. In: International Conference on Learning Representations
26. Wang P, Zheng W, Chen T, Wang Z: Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=O476oWmiNNp>
27. Wang W, Chen C, Ding M, Yu H, Zha S, Li J: Transbts: Multimodal brain tumor segmentation using transformer. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24. pp. 109–119. Springer (2021)
28. Wu H, Chen S, Chen G, Wang W, Lei B, Wen Z: Fat-net: Feature adaptive transformers for automated skin lesion segmentation. *Medical Image Analysis* 76, 102327 (2022) [PubMed: 34923250]
29. Wu Y, Liao K, Chen J, Wang J, Chen DZ, Gao H, Wu J: D-former: A u-shaped dilated transformer for 3d medical image segmentation. *Neural Computing and Applications* pp. 1–14 (2022)
30. Xie E, Wang W, Yu Z, Anandkumar A, Alvarez JM, Luo P: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems* 34, 12077–12090 (2021)
31. Xie Y, Zhang J, Shen C, Xia Y: Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24. pp. 171–180. Springer (2021)
32. Xu G, Wu X, Zhang X, He X: Levit-unet: Make faster encoders with transformer for medical image segmentation. arXiv preprint arXiv:2107.08623 (2021)



**Fig. 1:**  
Architecture of our proposed Laplacian-Former.



**Fig. 2:**  
The structure of our frequency enhancement Transformer block.



**Fig. 3:** Segmentation results of the proposed method on the *Synapse* dataset. Our Laplacian-Former shows finer boundaries (high-frequency details) for the region of the stomach and less false positive prediction for the pancreas.

**Table 1:**

Comparison results of the proposed method on the *Synapse* dataset. **Blue** indicates the best result, and **red** indicates the second-best.

Methods	# Params (M)	DSC ↑	HD ↓	Aorta	Gallbladder	Kidney(L)	Kidney(R)	Liver	Pancreas	Spleen	Stomach
R50 U-Net [5]	30.42	74.68	36.87	87.74	63.66	80.60	78.19	93.74	56.90	85.87	74.16
U-Net [20]	14.8	76.85	39.70	89.07	69.72	77.77	68.60	93.43	53.98	86.67	75.58
Att-UNet [22]	34.9	77.77	36.02	89.55	68.88	77.98	71.11	93.57	58.04	87.30	75.75
TransUNet [5]	105.28	77.48	31.69	87.23	63.13	81.87	77.02	94.08	85.86	85.08	75.62
Swin-Unet [4]	27.17	79.13	21.55	85.47	66.53	83.28	79.61	94.29	56.58	90.66	76.60
LeVit-Unet [32]	52.17	78.53	16.84	78.53	62.23	84.61	80.25	93.11	59.07	88.86	72.76
TransDeepLab [2]	21.14	80.16	21.25	86.04	69.16	84.08	79.88	93.53	61.19	89.00	78.40
HiFormer [12]	25.51	80.39	14.70	86.21	65.69	85.23	79.77	94.61	59.52	90.99	81.08
EffFormer	22.31	80.79	17.00	85.81	66.89	84.10	81.81	94.80	62.25	91.05	79.58
LaplacianFormer (without bridge)	23.87	81.59	17.31	87.41	69.57	85.22	80.46	94.68	63.71	91.47	78.23
<b>LaplacianFormer</b>	27.54	81.90	18.66	86.55	71.19	84.23	80.52	94.90	64.75	91.91	81.14

**Table 2:**

(a) Performance comparison of Laplacian-Former against the SOTA approaches on *ISIC 2018* skin lesion dataset. Blue and red indicates the best and the second-best results. (b) Frequency response analysis on the LaplacianFormer (up) vs. Standard Transformer (down).

