

# Beyond Self-Attention: Deformable Large Kernel Attention for Medical Image Segmentation

Reza Azad<sup>1</sup> Leon Niggemeier<sup>1</sup> Michael Hüttemann<sup>1</sup> Amirhossein Kazerouni<sup>2</sup>  
 Ehsan Khodapanah Aghdam<sup>3</sup> Yury Velichko<sup>4</sup> Ulas Bagci<sup>4</sup> Dorit Merhof<sup>5</sup>  
<sup>1</sup>RWTH Aachen University <sup>2</sup>Iran University of Science and Technology  
<sup>3</sup>Shahid Beheshti University <sup>4</sup>Northwestern University <sup>5</sup>University of Regensburg

{reza.azad, leon.niggemeier, michael.huettemann}@rwth-aachen.de, {dorit.merhof}@ur.de  
 {amirhossein477, ehsan.khpaghdam}@gmail.com, {ulas.bagci, y-velichko}@northwestern.edu

## Abstract

*Medical image segmentation has seen significant improvements with transformer models, which excel in grasping far-reaching contexts and global contextual information. However, the increasing computational demands of these models, proportional to the squared token count, limit their depth and resolution capabilities. Most current methods process  $D$  volumetric image data slice-by-slice (called pseudo 3D), missing crucial inter-slice information and thus reducing the model's overall performance. To address these challenges, we introduce the concept of **Deformable Large Kernel Attention (D-LKA Attention)**, a streamlined attention mechanism employing large convolution kernels to fully appreciate volumetric context. This mechanism operates within a receptive field akin to self-attention while sidestepping the computational overhead. Additionally, our proposed attention mechanism benefits from deformable convolutions to flexibly warp the sampling grid, enabling the model to adapt appropriately to diverse data patterns. We designed both 2D and 3D adaptations of the D-LKA Attention, with the latter excelling in cross-depth data understanding. Together, these components shape our novel hierarchical Vision Transformer architecture, the D-LKA Net. Evaluations of our model against leading methods on popular medical segmentation datasets (Synapse, NIH Pancreas, and Skin lesion) demonstrate its superior performance. Our code is publicly available at [GitHub](#).*

## 1. Introduction

Medical image segmentation plays a crucial role in computer-assisted diagnostics, aiding medical professionals in analyzing complex medical images. This process not only reduces the laboriousness of manual tasks and the dependence on medical expertise but also enables faster and

more accurate diagnoses. The automation of segmentation offers the potential for faster and more accurate diagnostic outcomes, facilitating appropriate treatment strategies and enabling the execution of image-guided surgical procedures. Thus, the imperative to create rapid and precise segmentation algorithms serves as a driving force behind the motivation for this research.

Since the mid-2010s, Convolutional Neural Networks (CNNs) have become the preferred technique for many computer vision applications. Their ability to automatically extract complex feature representations from raw data without the need for manual feature engineering has generated significant interest within the medical image analysis community. Many successful CNN architectures such as U-Net [43], Fully Convolutional Networks [39], DeepLab [16], or SegCaps (segmentation capsules) [33] have been developed. These architectures have achieved great success in the task of semantic segmentation and have outperformed state-of-the-art (SOTA) methods previously [3].

The problem of identifying objects at different scales is a key concern in computer vision research [36]. In CNNs, the size of a detectable object is closely linked to the receptive field dimensions of the corresponding network layer. If an object extends beyond the boundaries of this receptive field, this may lead to under-segmentation outcomes. Conversely, using excessively large receptive fields compared to an object's actual size can limit recognition, as background information may exert undue influence on predictions [26].

A promising approach to address this issue involves employing multiple kernels with distinct sizes in parallel, similar to the mechanism of an *Inception Block* [48]. However, increasing the kernel sizes to accommodate larger objects is limited in practice due to the exponential increase in parameters and computational requirements [26]. Consequently, various strategies, including pyramid pooling techniques [27] and dilated convolutions [58] at varying scales,

have emerged to capture multi-scale contextual information.

Another intuitive concept entails directly incorporating multi-scale image pyramids or their associated feature representations into the network architecture. Yet, this approach poses challenges, particularly concerning the feasibility of managing training and inference times [36]. The use of Encoder-Decoder networks, such as U-Net, has proven advantageous in this context. Such networks encode appearance and location in shallower layers, while deeper layers capture higher semantic information and context from the broader receptive fields of neurons [35]. Some methods combine features from different layers or predict features from layers of different sizes to use information from multiple scales [10]. Also, forecasting features from layers of varied scales have emerged, effectively enabling the integration of insights across multiple scales [36]. However, most Encoder-Decoder structures face a challenge: they frequently fail to maintain consistent features across different scales and mainly use the last decoder layer to generate segmentation results [3, 26].

Semantic segmentation is a task that involves predicting the semantic category for each pixel in an image based on a predefined label set. This task requires extracting high-level features while preserving the initial spatial resolution [37, 41]. CNNs are well suited to capture local details and low-level information, albeit at the expense of overlooking global context. This gap in handling global information has been a focus of the vision transformer (ViT) architecture, which has achieved remarkable success in several computer vision tasks, including semantic segmentation.

The cornerstone of the ViT is the *attention mechanism*, which facilitates the aggregation of information across the entire input sequence. This capability empowers the network to incorporate long-range contextual cues beyond the limited receptive field sizes of CNNs [23, 47]. However, this strategy usually limits the ability of ViTs to effectively model local information [9]. This limitation can impede their ability to detect local textures, which is crucial for various diagnostic and prognostic tasks. This lack of local representation can be attributed to the particular way ViT models process images. ViT models partition an image into a sequence of patches and model their dependencies using self-attention mechanisms. This approach may not be as effective as the convolution operations employed by CNN models for extracting local features within receptive fields. This difference in image processing methods between ViT and CNN models may explain the superior performance of CNN models in local feature extraction [6, 21]. In recent years, innovative approaches have been developed to address the insufficient representation of local textures within Transformer models. One such approach involves integrating CNN and ViT features through complementary methods to combine their strengths and mitigate any local represen-

tation shortcomings [15]. TransUNet [15] is an early example of this approach, incorporating Transformer layers within the CNN bottleneck to model both local and global dependencies. HiFormer [28] proposes a solution that combines a Swin Transformer module and a CNN-based encoder to generate two multi-scale feature representations that are integrated via a Double-Level Fusion module. UNETR [25] employs a Transformer-based encoder and a CNN decoder for 3D medical image segmentation. CoTr [56] and TransBTS [52] bridge the CNN encoder and decoder with the Transformer to enhance segmentation performance at low-resolution stages.

An alternate strategy to enhance local feature representation is to redesign the self-attention mechanism within pure Transformer models. In this vein, Swin-Unet [13] integrates a Swin Transformer [38] block of linear computational complexity within a U-shaped structure as a multi-scale backbone. MISSFormer [31] employs the Efficient Transformer [55] to address parameter issues in vision transformers by incorporating a non-invertible downsampling operation on input blocks. D-Former [54] introduces a pure transformer-based pipeline featuring a double attention module to capture fine-grained local attention and interaction with diverse units in a dilated manner. Nonetheless, certain specific limitations remain. These include computational inefficiency, as evidenced in the TransUNet model, the heavy dependence on a CNN backbone, as observed in HiFormer, and the lack of consideration of multi-scale information. Additionally, current segmentation architectures often take a slice-by-slice approach to process 3D input volumes, inadvertently disregarding potential correlations between neighboring slices. This oversight limits the comprehensive use of volumetric information, consequently compromising both localization accuracy and context integration. Furthermore, it is crucial to recognize that lesions in the medical domain often exhibit deformations in their shape. Therefore, any learning algorithm intended for medical image analysis must be endowed with the ability to capture and comprehend these deformations. Simultaneously, the algorithm should maintain computational efficiency to facilitate the processing of 3D volumetric data.

**Our contributions:** To address the challenges outlined above, we propose a solution in the form of the Deformable-LKA module (●), which serves as a fundamental building block within our network design. This module is explicitly designed to effectively handle contextual information while simultaneously preserving local descriptors. This balance between the two aspects within our architecture enhances its ability to achieve precise semantic segmentation. Notably, our model introduces a dynamic adaptation of receptive fields based on the data, diverging from the conventional fixed filter masks found in traditional convolutional operations. This adaptive approach allows us to overcome

the inherent limitations associated with static methods. This innovative approach extends to the development of both 2D and 3D versions of the D-LKA Net architecture (⊕). In the case of the 3D model, the D-LKA mechanism is tailored to suit a 3D context, thus enabling the seamless exchange of information across different volumetric slices. (⊕). Finally, our contribution is further emphasized by its computational efficiency. We achieve this through a design that leans solely on the D-LKA concept, resulting in a remarkable performance on various segmentation benchmarks that establish our method as a new SOTA approach.

## 2. Method

In this section, we begin by providing an overview of the methodology. First, we revisit the concept of Large Kernel Attention (LKA) as introduced by Guo et al. [22]. Then, we introduce our innovative exploration of the deformable LKA module. Built up on this, we introduce both 2D and 3D network architectures for segmentation tasks.

### 2.1. Large Kernel Attention

Large convolution kernels provide a similar receptive field as the self-attention mechanism. A large convolution kernel can be constructed with much less parameters and computation by using a depth-wise convolution, a depth-wise dilated convolution, and a  $1 \times 1$  convolution. The equations for the kernel sizes of the depthwise convolution and the depthwise dilated convolution to construct a  $K \times K$  kernel for an input of dimension  $H \times W$  and channels  $C$  are:

$$DW = (2d - 1) \times (2d - 1), \quad (1)$$

$$DW-D = \left\lceil \frac{K}{d} \right\rceil \times \left\lceil \frac{K}{d} \right\rceil, \quad (2)$$

with a kernel size of  $K$  and a dilation rate of  $d$ . The number of parameters  $P(K, d)$  and floating-point operations (FLOPs)  $F(K, d)$  is calculated by:

$$P(K, d) = C \left( \left\lceil \frac{K}{d} \right\rceil^2 + (2d - 1)^2 + 3 + C \right), \quad (3)$$

$$F(K, d) = P(K, d) \times H \times W. \quad (4)$$

The number of FLOPs grows linearly with the size of the input image. The number of parameters increases quadratically with the number of channels and kernel size. However, as both are usually so small, they are not restricting factors.

To minimize the number of parameters for a fixed kernel size  $K$ , the derivative of equation 3 with respect to dilation rate  $d$  can be set to zero:

$$\frac{d}{d\hat{d}} P(K, \hat{d}) \stackrel{!}{=} 0 = C \cdot \left( 8\hat{d} + \frac{2K^2}{\hat{d}^3} - 4 \right). \quad (5)$$

For example, when the kernel size is  $K = 21$ , this results in  $d \approx 3.37$ . Extending the formulas to the three-dimensional case is straightforward. For an input of size  $H \times W \times D$  and channels  $C$ , then the equations for the number of parameters  $P_{3d}(K, d)$  and FLOPs  $F_{3d}(K, d)$  are:

$$P_{3d}(K, d) = C \left( \left\lceil \frac{K}{d} \right\rceil^3 + (2d - 1)^3 + 3 + C \right), \quad (6)$$

$$F_{3d}(K, d) = P_{3d}(K, d) \times H \times W \times D, \quad (7)$$

with kernel size  $K$  and dilation  $d$ .

### 2.2. Deformable Large Kernel Attention

The concept of utilizing Large Kernels for medical image segmentation is extended by incorporating Deformable Convolutions [20]. Deformable Convolutions enable adjusting the sampling grid with whole-numbered offsets for free deformation. An additional convolutional layer learns The deformation from the feature maps, which creates an offset field. Learning the deformation based on the features itself results in an adaptive convolution kernel. This flexible kernel shape can improve the representation of lesions or organ deformations, resulting in an enhanced definition of object boundaries. The convolutional layer responsible for calculating offsets follows the kernel size and dilation of its corresponding convolutional layer. Bilinear interpolation is employed to compute pixel values for offsets that are not found on the image grid. As shown in Figure 2, the D-LKA module can be formulated as:

$$\begin{aligned} \text{Attention} &= \text{Conv}_{1 \times 1}(\text{Conv}_{DDW-D}(\text{Conv}_{DDW}(F'))), \\ \text{Output} &= \text{Conv}_{1 \times 1}(\text{Attention} \otimes F') + F, \end{aligned} \quad (8)$$

where the input feature is denoted by  $F \in \mathbb{R}^{C \times H \times W}$  and  $F' = \text{GELU}(\text{Conv}(F))$ . The attention component  $\in \mathbb{R}^{C \times H \times W}$  is represented as an attention map, with each value indicating the relative importance of corresponding features. The operator  $\otimes$  denotes an element-wise product operation. Notably, LKA departs from conventional attention methods by not requiring additional normalization functions such as sigmoid or Softmax. According to [51], these normalization functions tend to neglect high-frequency information, thereby decreasing the performance of self-attention-based methods.

In the 2D version of the approach, the convolution layers are substituted with deformable convolutions because deformable convolutions improve the ability to capture objects characterized by irregular shapes and sizes. Such objects are commonly found in medical image data, making this augmentation especially significant.

However, extending the concept of deformable LKA to the 3D domain presents certain challenges. The primary constraint arises from the additional convolution layer

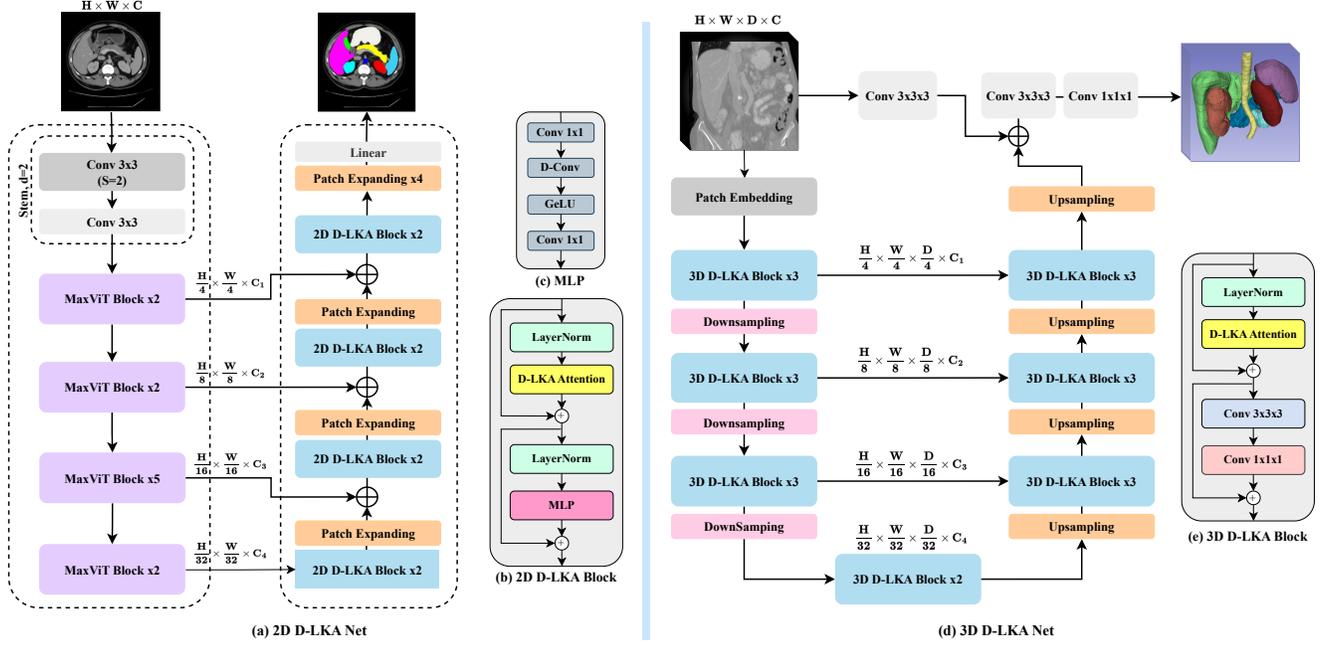


Figure 1. Proposed network architecture of the 2D D-LKA model on the left and the 3D D-LKA model on the right.

needed for offset generation. In contrast to the 2D case, this layer cannot be executed in a depth-wise manner due to the nature of input and output channels. In the 3D context, input channels correspond to features, and the output channels enlarge to  $3 \cdot k \times k \times k$ , with a kernel size of  $k$ . The intricate nature of large kernels leads to an expansion of channel count along the third dimension, causing a substantial rise in parameters and FLOPs. Consequently, an alternative approach is implemented for the 3D scenario. A sole deformable convolution layer is introduced into the existing LKA framework, following the depth-wise convolutions. This strategic design adaptation aims to mitigate the challenges posed by the extension to three dimensions.

### 2.3. 2D D-LKA Net

The architecture of the 2D network is illustrated in Figure 1. The first variant uses a MaxViT [49] as the encoder component for efficient feature extraction, while the second variant incorporates deformable LKA layers for more refined, superior segmentation.

In a more formal description, the encoder generates four hierarchical output representations. A convolutional stem first reduces the dimensions of the input image to  $\frac{H}{4} \times \frac{W}{4} \times C$ . Subsequently, feature extraction is carried out through four stages of MaxViT blocks, each followed by downsampling layers. As the process advances to the decoder, four stages of D-LKA layers are implemented, each stage encompassing two D-LKA blocks. Patch-expanding layers are then applied to achieve resolution upsampling while also decreasing channel dimensions. Finally, a linear layer is re-

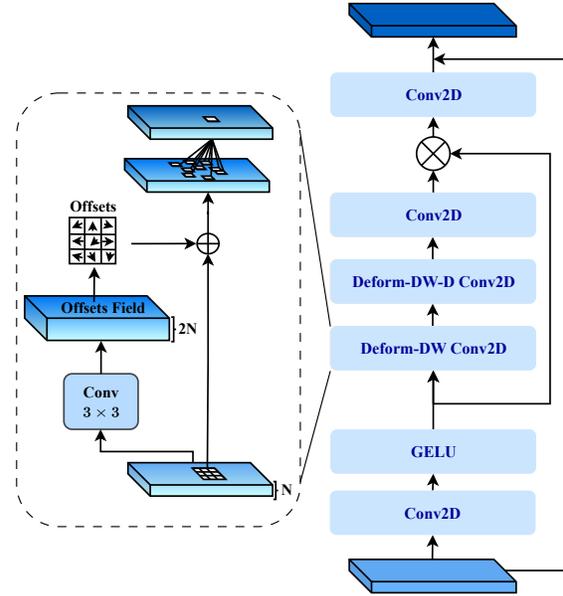


Figure 2. Architecture of the deformable LKA module.

sponsible for generating the ultimate output. The structure of the 2D D-LKA block comprises LayerNorm, deformable LKA, and a Multi-Layer Perceptron (MLP). The integration of residual connections ensures effective feature propagation, even across deeper layers. This arrangement can be mathematically represented as:

$$x_1 = D-LKA-Attn(LN(x_{in})) + x_{in}, \quad (9)$$

$$x_{out} = MLP(LN(x_1)) + x_1, \quad (10)$$

$$MLP = Conv_1(GeLU(Conv_d(Conv_1(x)))), \quad (11)$$

with input features  $x_{in}$ , layer normalization  $LN$ , deformable LKA attention  $D-LKA-Attn$ , depthwise convolution  $Conv_d$ , linear layers  $Conv_1$  and GeLU activation function  $GeLU$ .

## 2.4. 3D D-LKA Net

The 3D network architecture, depicted in Figure 1, is structured hierarchically using an encoder-decoder design. Initially, a patch embedding layer reduces the input image dimensions from  $(H \times W \times D)$  to  $(\frac{H}{4} \times \frac{W}{4} \times \frac{D}{2})$ . Within the encoder, a sequence of three D-LKA stages is employed, each containing three D-LKA blocks. After each stage, a downsampling step reduces spatial resolution by half while doubling the channel dimension. The central bottleneck includes another set of two D-LKA blocks. The decoder structure is symmetric to that of the encoder. In order to double the feature resolution while reducing channel count, transpose convolutions are utilized. Each decoder stage employs three D-LKA blocks to promote long-range feature dependencies. The final segmentation output is produced by a  $3 \times 3 \times 3$  convolutional layer, succeeded by a  $1 \times 1 \times 1$  convolutional layer to match class-specific channel requirements. To establish a direct connection between the input image and the segmentation output, a skip connection is formed using convolutions. Additional skip connections perform a fusion of features from other stages based on simple addition. The ultimate segmentation map is produced through a combination of  $3 \times 3 \times 3$  and  $1 \times 1 \times 1$  convolutional layers.

The 3D D-LKA block includes layer normalization followed by D-LKA Attention, with residual connections applied. The subsequent section employs a  $3 \times 3 \times 3$  convolutional layer, followed by a  $1 \times 1 \times 1$  convolutional layer, both accompanied by residual connections. This entire process can be summarized as follows:

$$x_1 = DAttn(LN(x_{in})) + x_{in}, \quad (12)$$

$$x_{out} = Conv_1(Conv_3(x_1)) + x_1, \quad (13)$$

with input features  $x_{in}$ , layer normalization  $LN$ , deformable LKA  $DAttn$ , convolutional layer  $Conv_1$ , and output features  $x_{out}$ .  $Conv_3$  refers to a feed forward network, with two convolutional layers and activation functions.

## 3. Experiments

### 3.1. Experimental Setup

We have implemented both 2D and 3D models using the PyTorch framework and performed training on a single RTX 3090 GPU. For the 2D method, a batch size of 20 was used, along with Stochastic Gradient Descent (SGD)

employing a base learning rate of 0.05, a momentum of 0.9, and a weight decay of 0.0001. The training process consisted of 400 epochs, employing a combination of cross-entropy and Dice loss, represented as follows:

$$\mathcal{L}_{total} = 0.6 \cdot \mathcal{L}_{dice} + 0.4 \cdot \mathcal{L}_{ce}. \quad (14)$$

Consistent with [15], identical data augmentation techniques were applied. For the 3D model, a batch size of 2 was chosen, and stochastic gradient descent was employed with a base learning rate of 0.01 and a weight decay of  $3e^{-5}$ . The input images were in the form of patches sized  $128 \times 128 \times 64$ . The training process consisted of 1000 epochs, with 250 patches utilized per epoch. Data augmentation techniques consistent with nnFormer [59] and UNETR++ [46] were employed.

### 3.2. Datasets

**Synapse Multi-Organ Segmentation:** First, we evaluate the performance of our method using the well-established synapse multi-organ segmentation dataset [14]. This dataset consists of 30 cases with a total of 3779 axial abdominal clinical CT images. Each CT volume comprises a range of 85 to 198 slices, each with dimensions of  $512 \times 512$  pixels. The voxel spatial resolution varies in the range of  $([0.54 \sim 0.54] \times [0.98 \sim 0.98] \times [2.5 \sim 5.0])$  mm<sup>3</sup>. Our evaluation follows the setting presented in [15, 46] for 2D and 3D versions.

**Skin Lesion Segmentation:** Our comprehensive experiments also extend to skin lesion segmentation datasets. Specifically, we leverage the ISIC 2017 dataset [19], which contains 2000 dermoscopic images for training, 150 for validation, and 600 for testing. Furthermore, we adopt the division scheme described in previous literature [1, 2, 5] for the ISIC 2018 dataset [18]. Additionally, the PH<sup>2</sup> dataset [40] is used as a dermoscopic image repository designed for both segmentation and classification tasks. This dataset consists of 200 dermoscopic images containing 160 nevi and 40 melanomas.

**NIH Pancreas Segmentation:** The publicly available NIH Pancreas dataset consists of 82 contrast-enhanced 3D abdominal CT volumes, each accompanied by manual annotations [44]. In our configuration, we use 62 samples for training and reserve the remaining samples for testing.

### 3.3. Quantitative and Qualitative Results

#### 3.3.1 2D results

**Synapse Dataset:** In Table 1, we present a comprehensive comparison of the leading performances achieved by other SOTA techniques in contrast to our proposed approach.

The results in terms of the Dice Similarity Coefficient (DSC) reveal that D-LKA Net exhibits superiority over the

Table 1. Comparison results of the proposed method evaluated on the Synapse dataset. **Blue** indicates the best result, and **red** displays the second-best. Parameters are reported in millions (M) and FLOPs in billions (G). DSC is presented for abdominal organs spleen (Spl), right kidney (RKid), left kidney (LKid), gallbladder (Gal), liver (Liv), stomach (Sto), aorta (Aor), and pancreas (Pan).

Methods	Params (M)	FLOPs (G)	Spl	RKid	LKid	Gal	Liv	Sto	Aor	Pan	Average	
											DSC $\uparrow$	HD95 $\downarrow$
TransUNet [15]	96.07	88.91	85.08	77.02	81.87	63.16	94.08	75.62	87.23	55.86	77.49	31.69
Swin-UNet [11]	27.17	6.16	90.66	79.61	83.28	66.53	94.29	76.60	85.47	56.58	79.13	21.55
LeViT-UNet-384 [57]	52.17	25.55	88.86	80.25	84.61	62.23	93.11	72.76	87.33	59.07	78.53	16.84
MISSFormer [32]	42.46	9.89	<b>91.92</b>	82.00	85.21	68.65	94.41	80.81	86.99	65.67	81.96	18.20
ScaleFormer [30]	111.6	48.93	89.40	83.31	86.36	<b>74.97</b>	<b>95.12</b>	80.14	<b>88.73</b>	64.85	<b>82.86</b>	16.81
HiFormer-B [29]	25.51	8.045	90.99	79.77	85.23	65.23	94.61	81.08	86.21	59.52	80.39	<b>14.70</b>
DAEFormer [4]	48.07	27.89	<b>91.82</b>	<b>82.39</b>	<b>87.66</b>	71.65	<b>95.08</b>	80.77	87.84	63.93	82.63	<b>16.39</b>
TransDeepLab [7]	21.14	16.31	89.00	79.88	84.08	69.16	93.53	78.40	86.04	61.19	80.16	21.25
PVT-CASCADE [42]	35.28	6.40	90.1	80.37	82.23	70.59	94.08	83.69	83.01	64.43	81.06	20.23
LKA Baseline	85.82	13.62	91.45	81.93	84.93	71.05	94.87	<b>83.71</b>	87.48	<b>66.76</b>	82.77	17.42
<b>2D D-LKA Net</b>	101.64	19.92	91.22	<b>84.92</b>	<b>88.38</b>	<b>73.79</b>	94.88	<b>84.94</b>	<b>88.34</b>	<b>67.71</b>	<b>84.27</b>	20.04

previously established SOTA methods. Specifically, it outperforms ScaleFormer [30] by 1.41%, DAEFormer [4] by 1.64%, and by even larger margins when compared to other approaches. Notably, significant improvements are seen in the segmentation of specific anatomical regions such as the right kidney, left kidney, stomach, and pancreas. Notably, the pancreas achieves a significantly improved segmentation result, showing an impressive 2.04% improvement over the second-best performing method. Given that the segmentation of smaller abdominal organs, such as the gallbladder or pancreas, has historically been a challenge for existing SOTA approaches, this notable performance improvement represents a significant step forward in achieving more accurate segmentation results. A qualitative comparison of different methods is shown in Figure 3. Compared to DAEFormer [4], our approach results in less misclassifications for the *stomach*. While Unet [43] and Swin-UNet [12] sometimes classify distant tissue as parts of the *liver*, *gallbladder* or *stomach*, our approach reduces the misclassifications and better represents the shape of the organs.

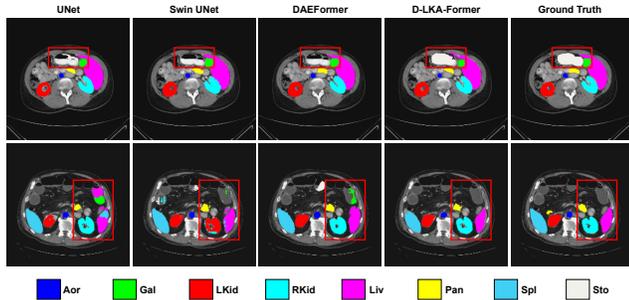


Figure 3. Qualitative comparison of the results from UNet [43], Swin UNet [12], DAEFormer [4] and D-LKA Net.

**Skin Lesion Segmentation Results:** The comparative outcomes for skin lesion segmentation benchmarks, includ-

ing ISIC 2017, ISIC 2018, and PH<sup>2</sup>, in contrast to leading methods, are detailed in Table 2. Notably, our D-LKA Net consistently outperforms its competitors across various evaluation metrics. This consistent superiority observed across different datasets underscores D-LKA Net’s robust generalization capabilities.

A qualitative comparison of the results is presented in Figure 4. In comparison to the baseline method, D-LKA Net better follows the complex outline of the lesions. In contrast to Swin-UNet and HiFormer-B, which tend to over- or under-segment certain areas, our approach achieves a more accurate segmentation.

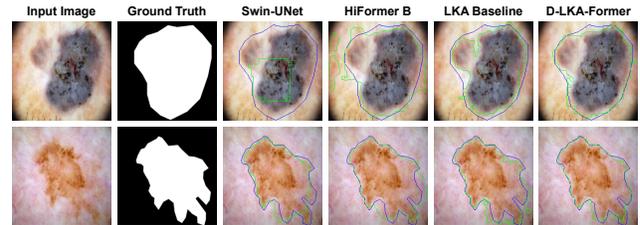


Figure 4. 2D visualizations on the ISIC 2018 [18] dataset.

### 3.3.2 3D results

**Synapse Dataset:** We compare our 3D approach with previous SOTA methods on the Synapse dataset. The results are presented in Table 3. We achieve a 0.27% improvement in *DSC* over the previous SOTA methods UNETR++ [46]. Compared to nnFormer [59], an improvement of 0.92% is achieved. For the HD95 metric, D-LKA Net reaches the second-best result. In comparison to UNETR++, a small performance increase is observed for *spleen*, *left kidney* and *aorta*. A significant increase is reported of *right kidney* and the small organs *gallbladder* and *pancreas*. An increase in the segmentation performance of these small organs is especially important.

Table 2. Performance comparison of the proposed method against the SOTA approaches on skin lesion segmentation benchmarks.

Methods	ISIC 2017				ISIC 2018				PH <sup>2</sup>			
	DSC	SE	SP	ACC	DSC	SE	SP	ACC	DSC	SE	SP	ACC
U-Net [43]	81.59	81.72	96.80	91.64	85.45	88.00	96.97	94.04	89.36	91.25	95.88	92.33
Att-UNet [45]	80.82	79.98	97.76	91.45	85.66	86.74	98.63	93.76	90.03	92.05	96.40	92.76
DAGAN [34]	84.25	83.63	97.16	93.04	88.07	90.72	95.88	93.24	92.01	83.20	96.40	94.25
TransUNet [15]	81.23	82.63	95.77	92.07	84.99	85.78	96.53	94.52	88.40	90.63	94.27	92.00
MCGU-Net [2]	89.27	85.02	98.55	95.70	89.50	84.80	98.60	95.50	92.63	83.22	97.14	95.37
MedT [50]	80.37	80.64	95.46	90.90	83.89	82.52	96.37	93.58	91.22	84.72	96.57	94.16
FAT-Net [53]	85.00	83.92	97.25	93.26	89.03	91.00	96.99	95.78	94.40	94.41	97.41	97.03
TMU-Net [8]	91.64	91.28	97.89	96.60	90.59	90.38	97.46	96.03	94.14	93.95	97.56	96.47
Swin-Unet [11]	91.83	91.42	97.98	97.01	89.46	90.56	97.98	96.45	94.49	94.10	95.64	96.78
DeepLabv3+ (CNN) [17]	91.62	87.33	99.21	96.91	88.20	85.60	97.70	95.10	92.02	88.18	98.32	95.03
HiFormer-B [28]	92.53	91.55	98.40	97.02	91.02	91.19	97.55	96.21	94.60	94.20	97.72	96.61
LKA (Baseline)	92.29	91.92	98.24	97.04	91.18	89.84	98.02	96.32	94.09	93.12	97.59	96.12
<b>2D D-LKA Net</b>	<b>92.54</b>	<b>93.27</b>	97.93	<b>97.05</b>	<b>91.77</b>	<b>91.64</b>	97.73	<b>96.47</b>	<b>94.90</b>	<b>94.30</b>	<b>97.75</b>	96.59

Table 3. Results on the Synapse dataset using 3D models. Blue indicates the best result, and red displays the second-best.

Methods	Params (M)	FLOPs (G)	Spl	RKid	LKid	Gal	Liv	Sto	Aor	Pan	Average	
											DSC ↑	HD95 ↓
UNETR [25]	92.49	75.76	85.00	84.52	85.60	56.30	94.57	70.46	89.80	60.47	78.35	18.59
Swin-UNETR [24]	62.83	384.2	95.37	86.26	86.99	66.54	95.72	77.01	91.12	68.80	83.48	10.55
nnFormer [59]	150.5	213.4	90.51	86.25	86.57	70.17	96.84	86.83	92.04	83.35	86.57	10.63
UNETR++ [46]	42.96	47.98	95.77	87.18	87.54	71.25	96.42	86.01	92.52	81.10	87.22	7.53
LKA Baseline	28.94	48.79	90.49	87.54	87.57	63.81	96.96	84.89	93.22	82.46	85.87	14.35
<b>D-LKA Net</b>	<b>42.35</b>	<b>66.96</b>	<b>95.88</b>	<b>88.50</b>	<b>87.64</b>	<b>72.14</b>	96.25	85.03	<b>92.87</b>	81.64	<b>87.49</b>	<b>9.57</b>

In terms of parameters, we have the lowest number of parameters with only 42.35 M while still achieving excellent segmentation performance. The number of FLOPs is 66.96 G, the second lowest. Only UNETR++ has less FLOPs. Compared to SOTA approaches like Swin-UNETR and nnFormer, we need only about 17% and 31% of the computations while achieving a better performance.

Figure 5 shows qualitative results on the Synapse dataset. In comparison to the nnFormer, we better capture the *aorta* as a whole and don't confuse other tissue as the organ. In comparison to UNETR++, we better segment the *pancreas*, whereas UNETR++ tends to under-segment. Also, our approach segments the *liver* and *stomach* more accurately than UNETR++, which tends to over-segment these organs.

**Pancreas Dataset:** The results of the NIH Pancreas dataset are presented in Table 4. Our approach achieves the best performance in all four metrics. In comparison to the closest competitor UNETR++, an increase of 0.63% in DSC, 0.82% in Jaccard, and a decrease of 1.04 in HD95 and 0.26 in ASD can be noted. D-LKA Net also has the lowest number of Parameters, with 62.07 M.

Figure 6 shows a qualitative comparison of different approaches. UNETR fails to segment the Pancreas as a single object. UNETR++ has smaller artifacts in the segmentation

Table 4. Results on the Pancreas dataset. The dice score, Jaccard index, HD95, and ASD are reported.

Methods	Params (M)	FLOPs (G)	DSC ↑	Jaccard ↑	HD95 ↓	ASD ↓
UNETR [25]	92.45	63.53	77.42	63.95	15.07	5.09
UNETR++ [46]	96.77	102.44	80.59	68.08	8.63	2.25
<b>D-LKA Net</b>	<b>62.07</b>	166.63	<b>81.22</b>	<b>68.90</b>	<b>7.59</b>	<b>1.99</b>

result. Our approach follows the highly irregular shape of the organ better than the other approaches.

### 3.4. Ablations Studies

**Robustness.** To enhance the robustness of our evaluation and to analyze statistical significance, we undertake 5 training runs for each method on the Synapse 2D version. This practice not only ensures a more comprehensive assessment but also enables us to visualize the variations in performance (Please refer to the supplementary file for the visualizations). In our evaluation, we observed a stable increase in performance for the aorta, gallbladder, left and right kidney, liver, pancreas, and stomach, where the median performance is higher than other SOTA methods. Solely the spleen segmentation performance is slightly worse. Furthermore, significant performance improvements are achieved for the gallbladder, pancreas, and stomach.

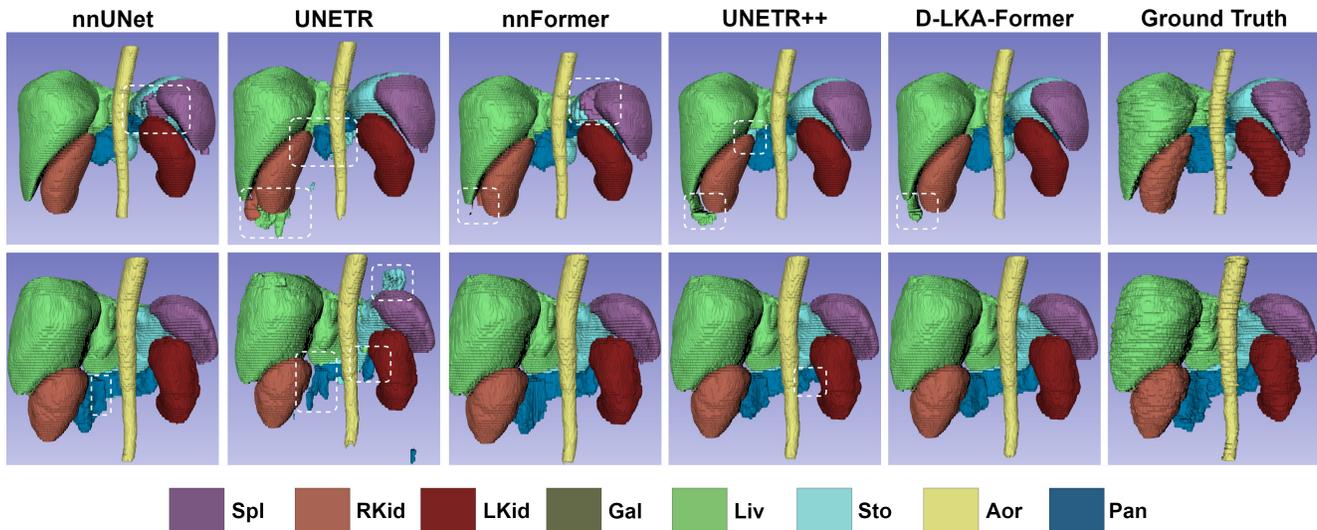


Figure 5. Qualitative results on the Synapse dataset. The boxes mark regions of failure.

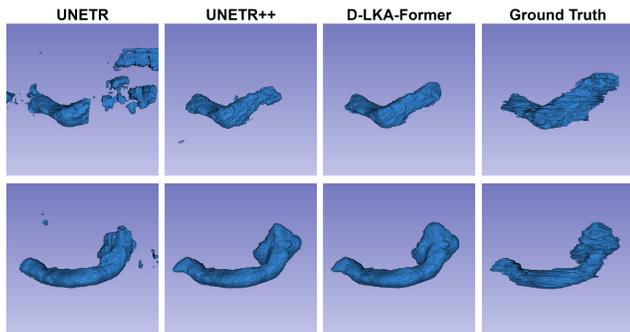


Figure 6. Qualitative results on the Pancreas Dataset.

**Deformable LKA Influence.** We continue our ablation study by determining the effectiveness of D-LKA. For this purpose, we construct a network employing 3D LKA without a deformable layer and another version using 3D LKA with an extra convolutional layer instead of the deformable one. The results of this analysis are presented in Table 5. Introducing the additional 3D convolutional layer results in a notable improvement in the performance of 0.99% in DSC when compared to the 3D LKA baseline. However, this modification also increases the number of parameters within the network. Replacing the 3D convolutional layer with a deformable convolutional layer leads to an additional performance boost, as indicated by a 0.63% increase in DSC. This alteration, similar to the previous one, also introduces more parameters and FLOPs into the network. Since the network’s size remains relatively small, the increase in these metrics is acceptable.

**Skip Connections.** Lastly, we assess the effect of the skip connections on the segmentation process. The results are shown in Table 6. We remove all skip connections and gradually add them to the network, starting with the highest-

Table 5. Results for the experiments of the influence of the deformable convolution layer with the Synapse Dataset.

Methods	Params	FLOPs	DSC $\uparrow$	HD95 $\downarrow$
3D LKA	28.94	48.79	85.87	14.35
3D LKA + 3D conv.	37.73	58.64	86.86	12.69
3D LKA + 3D deform. conv.	42.35	66.96	87.49	9.57

level skip connection. The results yield that skip connections are crucial for obtaining optimal segmentation performance. Additionally, we highlight that the highest-level skip connection is vital for achieving the best segmentation result, improving the DSC performance by 0.42%.

Table 6. Results for the experiments of the influence of the number of skip connections with the Synapse Dataset.

Metrics	0	1	2	3	4
DSC $\uparrow$	84.78	85.99	86.56	87.07	87.49
HD95 $\downarrow$	11.18	12.21	14.89	9.25	9.57

## 4. Conclusion

In this paper, we propose a novel Deformable LKA module. This attention mechanism enables the network to learn a deformation grid for accessing more relevant information than conventional attention strategies. Furthermore, the LKA mechanism can aggregate global information similar to that of self-attention to overcome the local restrictions of CNN mechanisms. Further, we present a 3D version of our proposed network, which includes cross-slice feature extraction with SOTA results.

**Acknowledgement:** The project is supported by the NIH funding: R01-CA246704, R01-CA240639, U01 DK127384-02S1, and U01-CA268808.

## References

- [1] Ehsan Khodapanah Aghdam, Reza Azad, Maral Zarvani, and Dorit Merhof. Attention swin u-net: Cross-contextual attention mechanism for skin lesion segmentation. *arXiv preprint arXiv:2210.16898*, 2022. 5
- [2] Maryam Asadi-Aghbolaghi, Reza Azad, Mahmood Fathy, and Sergio Escalera. Multi-level context gating of embedded collective knowledge for medical image segmentation. *arXiv preprint arXiv:2003.05056*, 2020. 5, 7
- [3] Reza Azad, Ehsan Khodapanah Aghdam, Amelie Rauland, Yiwei Jia, Atlas Haddadi Avval, Afshin Bozorgpour, Sanaz Karimijafarbigloo, Joseph Paul Cohen, Ehsan Adeli, and Dorit Merhof. Medical image segmentation review: The success of u-net. *arXiv preprint arXiv:2211.14830*, 2022. 1, 2
- [4] Reza Azad, René Arimond, Ehsan Khodapanah Aghdam, Amirhosein Kazerouni, and Dorit Merhof. Dae-former: Dual attention-guided efficient transformer for medical image segmentation. *arXiv preprint arXiv:2212.13504*, 2022. 6
- [5] Reza Azad, Maryam Asadi-Aghbolaghi, Mahmood Fathy, and Sergio Escalera. Bi-directional convlstm u-net with densley connected convolutions. In *ICCV 2019, IEEE International Conference on Computer Vision 2019*, 2019. 5
- [6] Reza Azad, Abdur R Fayjie, Claude Kauffmann, Ismail Ben Ayed, Marco Pedersoli, and Jose Dolz. On the texture bias for few-shot cnn segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2674–2683, 2021. 2
- [7] Reza Azad, Moein Heidari, Moein Shariatnia, Ehsan Khodapanah Aghdam, Sanaz Karimijafarbigloo, Ehsan Adeli, and Dorit Merhof. Transdeeplab: Convolution-free transformer-based deeplab v3+ for medical image segmentation. In *Predictive Intelligence in Medicine*, pages 91–102. Springer Nature Switzerland, 2022. 6
- [8] Reza Azad, Moein Heidari, Yuli Wu, and Dorit Merhof. Contextual attention network: Transformer meets u-net. In *Machine Learning in Medical Imaging: 13th International Workshop, MLMI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 18, 2022, Proceedings*, pages 377–386. Springer, 2022. 7
- [9] Reza Azad, Amirhossein Kazerouni, Moein Heidari, Ehsan Khodapanah Aghdam, Amirali Molaei, Yiwei Jia, Abin Jose, Rijo Roy, and Dorit Merhof. Advances in medical image analysis with vision transformers: A comprehensive review. *arXiv preprint arXiv:2301.03505*, 2023. 2
- [10] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation, 2016. 2
- [11] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation, 2021. 6, 7
- [12] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. In *European conference on computer vision*, pages 205–218. Springer, 2022. 6
- [13] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiopeng Zhang, Qi Tian, and Manning Wang. 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*, pages 205–218. Springer, 2023. 2
- [14] MICCAI 2015 Multi-Atlas Abdomen Labeling Challenge. Synapse multi-organ segmentation dataset. <https://www.synapse.org/#!Synapse:syn3193805/wiki/217789>, 2015. Accessed: 2022-04-20. 5
- [15] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021. 2, 5, 6, 7
- [16] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs, 2016. 1
- [17] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 7
- [18] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kaloo, Konstantinos Liopyris, Michael Marchetti, 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. 5, 6
- [19] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kaloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, pages 168–172. IEEE, 2018. 5
- [20] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. 3
- [21] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*. 2
- [22] Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual attention network. *arXiv preprint arXiv:2202.09741*, 2022. 3
- [23] Kai Han, Yunhe Wang, Hanqing Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, Zhaohui Yang, Yiman Zhang, and Dacheng Tao. A survey on vision transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):87–110, jan 2023. 2

- [24] Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger R Roth, and Daguang Xu. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In *International MICCAI Brainlesion Workshop*, pages 272–284. Springer, 2021. [7](#)
- [25] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 574–584, 2022. [2](#), [7](#)
- [26] Junjun He, Zhongying Deng, and Yu Qiao. Dynamic multi-scale filters for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [1](#), [2](#)
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015. [1](#)
- [28] Moein Heidari, Amirhossein Kazerouni, Milad Soltany, Reza Azad, Ehsan Khodapanah Aghdam, Julien Cohen-Adad, and Dorit Merhof. Hiformer: Hierarchical multi-scale representations using transformers for medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6202–6212, 2023. [2](#), [7](#)
- [29] Moein Heidari, Amirhossein Kazerouni, Milad Soltany, Reza Azad, Ehsan Khodapanah Aghdam, Julien Cohen-Adad, and Dorit Merhof. Hiformer: Hierarchical multi-scale representations using transformers for medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6202–6212, 2023. [6](#)
- [30] Huimin Huang, Shiao Xie, Lanfen Lin, Yutaro Iwamoto, Xi-anhua Han, Yen-Wei Chen, and Ruofeng Tong. Scaleformer: revisiting the transformer-based backbones from a scale-wise perspective for medical image segmentation. *arXiv preprint arXiv:2207.14552*, 2022. [6](#)
- [31] Xiaohong Huang, Zhifang Deng, Dandan Li, Xueguang Yuan, and Ying Fu. Missformer: An effective transformer for 2d medical image segmentation. *IEEE Transactions on Medical Imaging*, pages 1–1, 2022. [2](#)
- [32] Xiaohong Huang, Zhifang Deng, Dandan Li, Xueguang Yuan, and Ying Fu. Missformer: An effective transformer for 2d medical image segmentation. *IEEE Transactions on Medical Imaging*, 2022. [6](#)
- [33] Rodney LaLonde, Ziyue Xu, Ismail Irmakci, Sanjay Jain, and Ulas Bagci. Capsules for biomedical image segmentation. *Medical image analysis*, 68:101889, 2021. [1](#)
- [34] Baiying Lei, Zaimin Xia, Feng Jiang, Xudong Jiang, Zongyuan Ge, Yanwu Xu, Jing Qin, Siping Chen, Tianfu Wang, and Shuqiang Wang. Skin lesion segmentation via generative adversarial networks with dual discriminators. *Medical Image Analysis*, 64:101716, 2020. [7](#)
- [35] Di Lin, Yuanfeng Ji, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Multi-scale context intertwining for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [2](#)
- [36] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection, 2017. [1](#), [2](#)
- [37] Yang Liu, Yao Zhang, Yixin Wang, Feng Hou, Jin Yuan, Jiang Tian, Yang Zhang, Zhongchao Shi, Jianping Fan, and Zhiqiang He. A survey of visual transformers. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21, 2023. [2](#)
- [38] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. [2](#)
- [39] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [1](#)
- [40] Teresa Mendonça, Pedro M Ferreira, Jorge S Marques, André RS Marcal, and Jorge Rozeira. Ph 2-a dermoscopic image database for research and benchmarking. In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pages 5437–5440. IEEE, 2013. [5](#)
- [41] Yujian Mo, Yan Wu, Xinneng Yang, Feilin Liu, and Yujun Liao. Review the state-of-the-art technologies of semantic segmentation based on deep learning. *Neurocomputing*, 493, 01 2022. [2](#)
- [42] Md Mostafijur Rahman and Radu Marculescu. Medical image segmentation via cascaded attention decoding. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6222–6231, 2023. [6](#)
- [43] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [1](#), [6](#), [7](#)
- [44] Holger R Roth, Le Lu, Amal Farag, Hoo-Chang Shin, Jiamin Liu, Evrim B Turkbey, and Ronald M Summers. Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part I 18*, pages 556–564. Springer, 2015. [5](#)
- [45] Jo Schlemper, Ozan Oktay, Michiel Schaap, Mattias Heinrich, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention gated networks: Learning to leverage salient regions in medical images. *Medical image analysis*, 53:197–207, 2019. [7](#)
- [46] Abdelrahman Shaker, Muhammad Maaz, Hanoona Rasheed, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Unetr++: delving into efficient and accurate 3d medical image segmentation. *arXiv preprint arXiv:2212.04497*, 2022. [5](#), [6](#), [7](#)
- [47] Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu. Transformers in medical imaging: A survey, 2022. [2](#)

- [48] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions, 2014. [1](#)
- [49] Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxvit: Multi-axis vision transformer. In *European conference on computer vision*, pages 459–479. Springer, 2022. [4](#)
- [50] Jeya Maria Jose Valanarasu, Poojan Oza, Ilker Hacihaliloglu, and Vishal M Patel. Medical transformer: Gated axial-attention for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 36–46. Springer, 2021. [7](#)
- [51] Peihao Wang, Wenqing Zheng, Tianlong Chen, and Zhangyang Wang. Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. In *International Conference on Learning Representations*, 2022. [3](#)
- [52] Wenxuan Wang, Chen Chen, Meng Ding, Hong Yu, Sen Zha, and Jiangyun Li. 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*, pages 109–119. Springer, 2021. [2](#)
- [53] Huisi Wu, Shihuai Chen, Guilian Chen, Wei Wang, Baiying Lei, and Zhenkun Wen. Fat-net: Feature adaptive transformers for automated skin lesion segmentation. *Medical Image Analysis*, 76:102327, 2022. [7](#)
- [54] Yixuan Wu, Kuanlun Liao, Jintai Chen, Jinhong Wang, Danny Z Chen, Honghao Gao, and Jian Wu. D-former: A u-shaped dilated transformer for 3d medical image segmentation. *Neural Computing and Applications*, pages 1–14, 2022. [2](#)
- [55] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. [2](#)
- [56] Yutong Xie, Jianpeng Zhang, Chunhua Shen, and Yong Xia. 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*, pages 171–180. Springer, 2021. [2](#)
- [57] Guoping Xu, Xingrong Wu, Xuan Zhang, and Xinwei He. Levit-unet: Make faster encoders with transformer for medical image segmentation. *arXiv preprint arXiv:2107.08623*, 2021. [6](#)
- [58] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015. [1](#)
- [59] Hong-Yu Zhou, Jiansen Guo, Yinghao Zhang, Lequan Yu, Liansheng Wang, and Yizhou Yu. nnformer: Interleaved transformer for volumetric segmentation. *arXiv preprint arXiv:2109.03201*, 2021. [5](#), [6](#), [7](#)