

HiFormer: Hierarchical Multi-scale Representations Using Transformers for Medical Image Segmentation

Moein Heidari^{*,1} Amirhossein Kazerouni^{*,1} Milad Soltany^{*,1} Reza Azad²

Ehsan Khodapanah Aghdam³ Julien Cohen-Adad⁴ Dorit Merhof^{†,5,6}

¹ School of Electrical Engineering, Iran University of Science and Technology, Tehran, Iran

² Institute of Imaging and Computer Vision, RWTH Aachen University, Aachen, Germany

³ Department of Electrical Engineering, Shahid Beheshti University, Tehran, Iran

⁴ MILA, Quebec AI Institute, Montreal, Canada

⁵ Faculty of Informatics and Data Science, University of Regensburg, Regensburg, Germany

⁶ Fraunhofer Institute for Digital Medicine MEVIS, Bremen, Germany

moein.heidari@elec.iust.ac.ir, {amirhossein477, soltany.m.99, ehsan.khpaghdam}@gmail.com

azad@lfb.rwth-aachen.de, jcohen@polymtl.ca, dorit.merhof@ur.de

Abstract

Convolutional neural networks (CNNs) have been the consensus for medical image segmentation tasks. However, they suffer from the limitation in modeling long-range dependencies and spatial correlations due to the nature of convolution operation. Although transformers were first developed to address this issue, they fail to capture low-level features. In contrast, it is demonstrated that both local and global features are crucial for dense prediction, such as segmenting in challenging contexts. In this paper, we propose HiFormer, a novel method that efficiently bridges a CNN and a transformer for medical image segmentation. Specifically, we design two multi-scale feature representations using the seminal Swin Transformer module and a CNN-based encoder. To secure a fine fusion of global and local features obtained from the two aforementioned representations, we propose a Double-Level Fusion (DLF) module in the skip connection of the encoder-decoder structure. Extensive experiments on various medical image segmentation datasets demonstrate the effectiveness of HiFormer over other CNN-based, transformer-based, and hybrid methods in terms of computational complexity, quantitative and qualitative results. Our code is publicly available at [GitHub](#).

1. Introduction

Medical image segmentation is one of the main challenges in computer vision, which provides valuable information about the areas of anatomy needed for a detailed

analysis. This information can greatly assist doctors in depicting injuries, monitoring disease progression, and assessing the need for appropriate treatment. As a result of the growing use of medical image analysis, highly precise and robust segmentation has become increasingly vital.

With their impressive ability to extract image features, Convolutional Neural Networks (CNNs) have been used widely for different image segmentation tasks. With the rise of encoder-decoder-based networks, like Fully Convolutional Networks (FCNs) [36], U-shaped structures, e.g. U-Net [41], and their variants, CNNs have experienced remarkable success in medical image segmentation tasks. In both structures, skip connections are employed to embody high-level and fine-grained features provided by the encoder and decoder paths, respectively. Despite the success of CNN models in various computer vision tasks, their performance is restricted due to their limited receptive field and the inherent inductive bias [19, 4]. The aforementioned reasons prevent CNNs from building global contexts and long-range dependencies in images and, therefore, capping their performance in image segmentation.

Recently, motivated by the outstanding success of transformers in Natural Language Processing (NLP) [47], vision transformers have been developed to mitigate the drawbacks of CNNs in image recognition tasks [19]. Transformers primarily leverage a multi-head self-attention (MSA) mechanism that can effectively construct long-range dependencies between the sequence of tokens and capture global contexts. The vanilla vision transformer [19] exhibits comparable performance with CNN-based methods but requires large amounts of data to generalize and suffers from quadratic complexity. Several approaches have

^{*}Equal contribution

[†]Corresponding author

been proposed to address these limitations. DeiT [44] proposes an efficient knowledge distillation training scheme to overcome the difficulty of vision transformers demanding a great deal of data to learn. Swin Transformer [35] and pyramid vision transformer [48] attempt to reduce vision transformers' computational complexity by exploiting window-based and spatial reduction attention, respectively.

Moreover, multi-scale feature representations have lately demonstrated powerful performance in vision transformers. CrossViT [12] proposes a novel dual-branch transformer architecture that extracts multi-scale contextual information and provides more fine-grained feature representations for image classification. Similarly, DS-TransUNet [33] presents a dual-branch Swin Transformer to capture different semantic scale information in the encoder for the task of medical image segmentation. HRViT [23] connects multi-branch high-resolution architectures with vision transformers for semantic segmentation. As a result, such structures can effectively aid in enhancing the modeling of long-range relationships between tokens and obtaining more detailed information.

Despite the vision transformers' ability to model the global contextual representation, the self-attention mechanism induces missing low-level features. Hybrid CNN-transformer approaches have been proposed to alleviate the problem above by leveraging the locality of CNNs and the long-range dependency character of transformers to encode both global and local features, particularly TransUNet [13] and LeViT-Unet [52] in medical image segmentation. However, these approaches have some impediments that prevent them from attaining higher performance: 1) they cannot effectively combine low-level and high-level features while maintaining feature consistency, and 2) they do not use the multi-scale information produced by the hierarchical encoder properly.

In this paper, we propose a novel encoder-decoder CNN-transformer-based framework that efficiently leverages the global long-range relationships of transformers and local feature representations of CNNs for an accurate medical image segmentation task. The encoder comprises three modules: two hierarchical CNN and Swin Transformer modules and the DLF module. Swin Transformer and CNN modules each contain three levels. First, an input image is fed into a CNN module to learn its local semantic representation. To compensate for the lack of global representation, the Swin Transformer module is applied on top of CNN's shallow features to capture long-range dependencies. Next, a pyramid of Swin Transformer modules with varying window sizes is utilized to learn multi-scale interaction. To encourage feature reusability and provide localization information, a skip connection module is designed to transfer CNN's local features into the Transformer blocks. The resulting representation of the smallest and largest pyramid levels is then

entered into the DLF module. The novel proposed DLF module is a multi-scale vision transformer that fuses two obtained feature maps using a cross-attention mechanism. Finally, both recalibrated feature maps are passed into the decoder block to produce the final segmentation mask. Our proposed HiFormer not only alleviates the problem mentioned above but also surpasses all its counterparts in terms of different evaluation metrics. Our main contributions:

- A novel hybrid method that merges the long-range contextual interactions of the transformer and the local semantic information of CNN.
- A DLF module to establish effective feature fusion between coarse and fine-grained feature representations.
- Experimental results demonstrate the effectiveness and superiority of the proposed HiFormer compared to the competing methods on medical image segmentation datasets.

2. Related Works

2.1. CNN-based Segmentation Networks

Convolutional Neural Networks are considered the de-facto standard for different computer vision tasks. One area where CNNs have achieved excellent results is image segmentation, where class labels are assigned to each pixel. Long et al. [36] showed that fully convolutional networks (FCNs) can be used to segment images without fully connected layers. Given that the output from vanilla FCNs, where the convolutional layers are stacked sequentially, is usually coarse, other models were proposed that fuse the output of different layers [6, 39, 41]. Several approaches have been introduced to improve the limited receptive field of FCN, including dilated convolution [14, 54] and context modeling [56, 15]. CNN models have shown outstanding performance in medical imaging tasks. After the introduction U-net [41], other researchers focused on utilizing U-shaped encoder-decoder structures. In [46], an over-complete network is augmented with U-net, and in U-net++ [57], the encoder-decoder architecture is re-designed by adding dense skip connection between the modules. This structure has been further improved and utilized in different medical domains [8, 30, 21, 5].

2.2. Vision Transformers

Following the remarkable success of transformers in NLP [47], Dosovitskiy et al. [19] propose the Vision Transformer (ViT), which achieved state-of-the-art performance on image classification tasks by employing self-attention mechanisms to learn global information. Several derivatives of vision transformers have been introduced to make them more efficient and less dependent on a large-sized dataset to achieve generalization [44, 55, 51].

In addition, many approaches have been presented, focusing on multi-scale representations to improve the accuracy and efficiency via extracting information from different scales. Inspired by the pyramid structure in CNNs

[38, 53, 9, 34], PVT [48] was the first introduced pyramid vision transformer. Later, Swin Transformer [35] proposes a hierarchical vision transformer using an efficient shifted windowing approach for computing self-attention locally. CrossViT [12] suggests using a dual-branch vision transformer followed by a cross-attention module for richer feature representations while performing in linear time. Vision transformers have also shown impressive results in other vision tasks, including [58, 20], which offer end-to-end transformer-based models for object detection, and [43, 24] for semantic and instance segmentation.

2.3. Transformers for Medical Image Segmentation

Despite the encouraging results of CNN models, such approaches generally demonstrate restrictions for modeling long-range dependencies due to their limited receptive field, thereby yielding weak performance. Recently, transformer-based models have gained significant popularity over CNN models in medical image segmentation. Swin-UNet [10] and DS-TransUNet [33] propose pure transformer models with a U-shaped architecture based on Swin Transformer for 2D segmentation. In addition to fully transformer models, TransUNet [13] takes advantage of both CNNs and transformers to capture both low-level and high-level features. UNETR [27] uses a transformer-based encoder to embed input 3D patches and a CNN-based decoder to achieve the final 3D segmentation results. Most prior works utilize either CNN, lacking in global features, or transformers, limited in local feature representation for feature extraction; this renders ineffective feature maps that do not contain rich information. In hybrid works, simple feature-fusing mechanisms are employed that cannot guarantee feature consistency between different scales. Motivated by multi-scale representations, we propose HiFormer, a CNN-transformer-based architecture that effectively incorporates both global and local information and utilizes a novel transformer-based fusing scheme to maintain feature richness and consistency for the task of 2D medical image segmentation.

3. Method

An overview of the proposed HiFormer is presented in this section. As illustrated in Fig. 1a, our proposed architecture provides an end-to-end training strategy that integrates global contextual representations from Swin Transformer and local representative features from the CNN module in the encoder. A richer feature representation is then obtained using the Double-level Fusion module (DLF). Afterward, the decoder outputs the final segmentation map.

3.1. Encoder

As shown in Fig. 1a, the proposed encoder is composed of two hierarchical models, CNN and Swin Transformer, with the DLF module that enriches the retrieved features

and prepares them to be fed into the decoder. Since using CNNs or transformers separately causes either local or global features to be neglected, which affects the model's performance, we first utilize the CNN locality trait to obtain local features. Here, the CNN and Swin Transformer each include three distinct levels. We transfer local features of each level to the corresponding Swin Transformer's level via a skip connection to attain universal representations. Then each transferred CNN level is added with its parallel transformer level and passes through a Patch Merging module to produce a hierarchical representation (see Fig. 1a). We exploit the hierarchical design to take advantage of multi-scale representations. The largest and smallest levels go into the DLF module to exchange information from different scales and generate more powerful features. In the following, we will discuss our CNN, Swin Transformer, and DLF modules deeply and in detail.

3.1.1 CNN Module

The proposed encoder begins by employing a CNN as the feature extractor to build a pyramid of intermediate CNN feature maps of different resolutions. Taking an input image $X \in R^{H \times W \times C}$ with spatial dimensions H and W , and C channels, it is first fed into the CNN module. CNN module consists of three levels, from which a skip connection is connected to the associated transformer's level using a Conv 1×1 to compensate for low-level missing information of transformers and recover localized spatial information.

3.1.2 Swin Transformer Module

The vanilla transformer encoder block [19] consists of two main modules: a multi-head self-attention (MSA) and a multi-layer perceptron (MLP). The vanilla transformer is composed of N identical transformer encoder blocks. In each block, before the MSA and the MLP blocks, Layer-Norm (LN) is applied. Additionally, a copy of the activations is also added to the output of the MSA or MLP block through skip-connections. One major problem with the vanilla ViT, which uses the standard MSA, is its quadratic complexity, rendering it inefficient for high-resolution computer vision tasks like image segmentation. To overcome this limitation, Swin Transformer [10] introduced the W-MSA and SW-MSA.

The Swin Transformer module includes two successive modified transformer blocks; the MSA block is replaced with the window-based multi-head self-attention (W-MSA) and the shifted window-based multi-head self-attention (SW-MSA). In the W-MSA module, self-attention will be applied to local windows of size $M \times M$. The W-MSA module has linear complexity; however, given that there is no connection across windows, it has limited modeling power. To alleviate this, SW-MSA is introduced that utilizes a windowing configuration that is shifted compared to the input of the W-MSA module; this is to ensure that we have cross-window connections. This process is depicted in

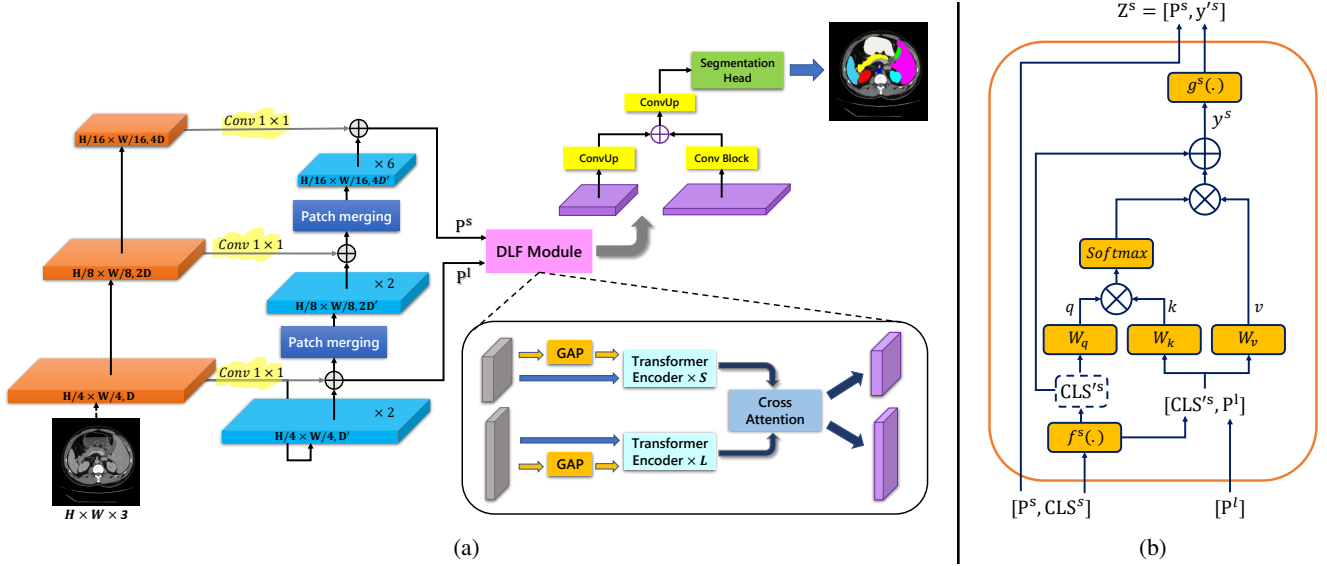


Figure 1: **(a) The overview of the proposed HiFormer.** HiFormer consists of a hierarchical CNN-transformer feature extractor module; outputs of the first and last levels are fed through the DLF feature fusion module. Afterward, the decoder uses the DLF’s output to generate accurate segmentation maps. In the diagram, **blue** and **orange** blocks denote Swin Transformer and CNN levels, respectively. **(b) The overview of Cross Attention.** The class token of the small level, CLS^s , is first projected for dimension alignment and then appended to P^l . The resulting embedding performs as a key and value. Moreover, CLS^s is utilized for the query. Finally, after computing attention and back projection, Z^s is obtained. This process can also be extended to the large level.

Eq 1.

$$\begin{aligned}
 \hat{z}^l &= \text{W-MSA}(\text{LN}(z^{l-1})) + z^{l-1}, \\
 z^l &= \text{MLP}(\text{LN}(\hat{z}^l)) + \hat{z}^l, \\
 \hat{z}^{l+1} &= \text{SW-MSA}(\text{LN}(z^l)) + z^l, \\
 z^{l+1} &= \text{MLP}(\text{LN}(\hat{z}^{l+1})) + \hat{z}^{l+1},
 \end{aligned} \quad (1)$$

The output of the first level in the CNN pyramid will be fed into a 1×1 convolution to generate $(H/4 \times W/4)$ patches (tokens) of length D' . These patches go through the first Swin Transformer block, generating the first attention-based feature maps. A skip-connection adds the previous activations to the obtained feature maps, resulting in the largest branch feature map P^l . Next, patch-merging is applied, which concatenates 2×2 groups of adjacent patches, applies a linear layer, and increases the embedding dimensions from D' to $2D'$ while reducing resolution. Similarly, higher-level feature maps of both the CNN and attention-based feature maps are fused and fed into Swin Transformer blocks to generate higher-level outputs. The latter is denoted as P^s , the smallest level feature map.

3.1.3 Double-Level Fusion Module (DLF)

The main challenge is efficiently fusing CNN and Swin Transformer level features while preserving feature consistency. A straightforward approach is to directly feed the summation of CNN levels with their matching Swin Transformer levels through a decoder and attain the segmentation map. Such approach, however, fails to ensure feature

consistency between them, leading to subpar performance. Hence, we propose a novel Double-Level Fusion (DLF) module, which takes the resultant smallest (P^s) and largest (P^l) levels as inputs and employs a cross-attention mechanism to fuse information across scales.

In general, shallow levels have better localization information, and as we approach deeper levels, semantic information becomes more prevalent and is better suited for the decoder part. Faced with the dilemma of extensive computational cost and imperceptible effect of middle-level feature map in model accuracy, we did not consider using the middle level in feature fusion to save computational costs. As a result, we encourage representation by multiscale the shallowest (P^s) and last (P^l) levels while preserving localization information.

In the proposed DLF module, the class token plays a significant role since it summarizes all the information of input features. We assign each level a class token derived from global average pooling (GAP) over the level’s norm. We obtain class tokens as demonstrated below:

$$\begin{aligned}
 CLS^s &= \text{GAP}(\text{Norm}(P^s)) \\
 CLS^l &= \text{GAP}(\text{Norm}(P^l))
 \end{aligned} \quad (2)$$

where $CLS^s \in R^{4D' \times 1}$ and $CLS^l \in R^{D' \times 1}$. Class tokens are then concatenated with associated level embeddings before being passed into the transformer encoders. The small level is followed by S and the large level by L transformer

encoders for computing global self-attention. Notably, we also add a learnable position embedding for each token of both levels before giving them to the transformer encoders for learning position information.

After passing embeddings through the transformer encoders, features of each level are fused using the cross-attention module. Specifically, before fusion, two-level class tokens are swapped, which means the class token of one level concatenates with the tokens of the another level. Then each new embedding is separately fed through the module for fusion and finally back projected to its own level. This interaction with other level tokens enables class tokens to share rich information with their cross-level.

In particular, this displacement for the small level is shown in Fig. 1b. $f^s(\cdot)$ first projects CLS^s to the dimensionality of P^l , and the output is denoted as CLS'^s . CLS'^s concatenated with P^l serves as a key and value and independently performs as a query for computing attention. Since we only query the class token, the cross-attention mechanism operates in linear time. The final output Z^s can be mathematically written as follows:

$$\begin{aligned} y^s &= f^s(CLS^s) + MCA(LN([f^s(CLS^s) \parallel P^l])) \\ Z^s &= [P^s \parallel g^s(y^s)] \end{aligned} \quad (3)$$

3.2. Decoder

Motivated by Semantic FPN [31], we design a decoder that combines features from the P^s and P^l levels into a unified mask feature. First, the low and high-resolution feature maps, P^s and P^l , are received from the DLF module. P^s (H/16, W/16) is followed by a ConvUp block which applies two stages of 3×3 Conv, $2 \times$ bilinear upsampling, Group Norm [50], and ReLU to attain (H/4, W/4) resolution. P^l (H/4, W/4) is also followed by a Conv Block, which employs a 3×3 Conv, Group Norm, and ReLU and remains at (H/4, W/4) resolution. The summation of both processed P^s and P^l is headed through another ConvUp block to achieve the final unified $H \times W$ feature map. After passing the acquired feature map through a 3×3 Conv in the segmentation head, the final segmentation map is generated.

4. Experiments

4.1. Dataset

Synapse Multi-Organ Segmentation: First, we evaluate HiFormer’s performance on the benchmarked synapse multi-organ segmentation dataset [11]. This dataset includes 30 cases with 3779 axial abdominal clinical CT images where each CT volume involves $85 \sim 198$ slices of 512×512 pixels, with a voxel spatial resolution of $([0.54 \sim 0.54] \times [0.98 \sim 0.98] \times [2.5 \sim 5.0]) \text{ mm}^3$.

Skin Lesion Segmentation: We conduct extensive experiments on the skin lesion segmentation datasets. Specifically, we utilize the ISIC 2017 dataset [18] comprising 2000 dermoscopic images for training, 150 for validation, and

600 for testing. Moreover, we adopt the ISIC 2018 [17] and follow the literature work [1, 2] to divide the dataset into train, validation, and test sets accordingly. Besides, the PH² dataset [37] is used, a dermoscopic image database introduced for both segmentation and classification tasks.

Multiple Myeloma Segmentation: We also evaluate our methodology on multiple myeloma cell segmentation grand challenges provided by SegPC 2021 [25, 26]. The challenge dataset includes a training set with 290 samples and validation and test sets with 200 and 277 samples, respectively.

4.2. Implementation Details

We implemented our framework in PyTorch and trained on a single Nvidia RTX 3090 GPU with 24 GB of memory. The input image size is 224×224 , and we set the batch size and learning rate to 10 and 0.01 during training, respectively. In addition, we use the weights pre-trained on ImageNet for the CNN and Swin Transformer modules to initialize their parameters. Our model is optimized using the SGD optimizer with a momentum of 0.9 and weight decay of 0.0001. Moreover, data augmentations such as flipping and rotating are employed during training to boost diversity. Table 1 depicts the suggested model’s final configurations.

Table 1: **The proposed model configurations.** WS represents window size, D' expresses the embedding dimension, and r denotes the MLP expanding ratio used in the transformer block. The number of heads in the DLF module is the same for both levels.

Model	CNN	Swin Transformer				DLF					
		D'	# Layer	# Head	WS	Dimension	P^s	P^l	S	L	# Head
HiFormer-S	ResNet34	96	[2,2,6]	[3,6,12]	7	384	96	1	1	1	3
HiFormer-B	ResNet50	96	[2,2,6]	[3,6,12]	7	384	96	2	1	2	6
HiFormer-L	ResNet34	96	[2,2,6]	[3,6,12]	7	384	96	4	1	4	6

4.3. Evaluation Results

We adopt a task-specific paradigm in terms of evaluation metrics in each experiment. Specifically, these metrics include the Dice score, 95% Hausdorff Distance (HD), Sensitivity and Specificity, Accuracy, and mIOU. To ensure an unprejudiced comparison, we contrast HiFormer against both CNN and transformer-based methods, along with the models formulated on an amalgamation of both.

4.3.1 Results of Synapse Multi-Organ Segmentation

The comparison of the proposal with previous state-of-the-art (SOTA) methods in terms of the average Dice-Similarity Coefficient (DSC) and average Hausdorff Distance (HD) on eight abdominal organs is shown in Table 2. HiFormer outperforms CNN-based SOTA methods by a large margin. Compared to other transformer-based models, our HiFormer-B shows superior learning ability on both evaluation metrics, observing an increase of 2.91% and 1.26% in Dice score and a decrease of 16.99 and 6.85 in average HD compared to TransUnet and Swin-Unet, respectively. Concretely, HiFormer steadily beats the literature work in the

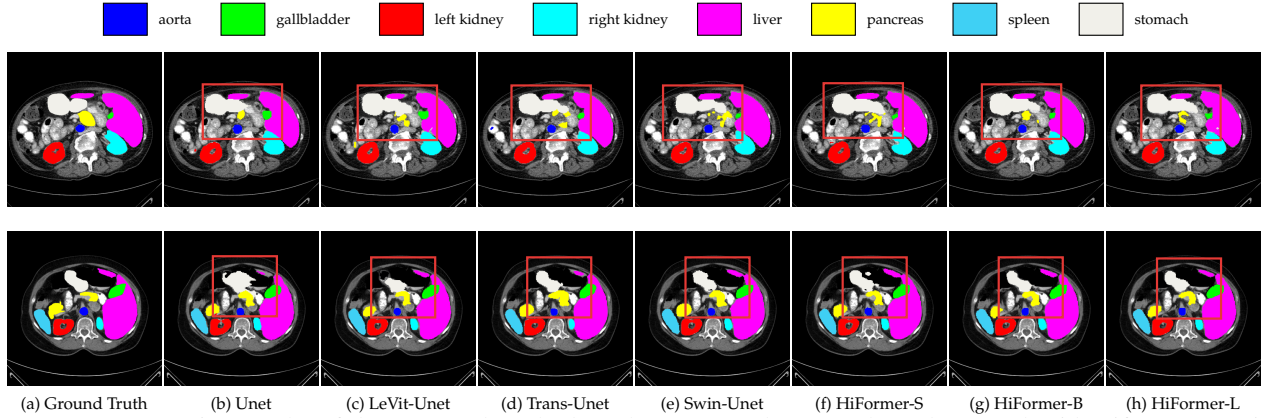


Figure 2: Segmentation results of the proposed method on the *Synapse* dataset. The red rectangles identify organ regions where the superiority of our proposed method can be clearly seen.

segmentation of most organs, particularly for the stomach, kidney, and liver segmentation. One can observe that HiFormer has distinct advantages over other methods in terms of average HD. Besides, the efficiency in terms of the number of parameters is indicated in Table 2, which will be discussed in the following sections. A characteristic qualitative example of the results is given in Fig. 2. We have observed that the proposed method can accurately segment fine and complex structures and output more accurate segmentation results, which are more robust to complicated backgrounds.

4.3.2 Results of Skin Lesion Segmentation

The comparison results for benchmarks of ISIC 2017, ISIC 2018, and PH² skin lesion segmentation task against leading methods are presented in Table 3. Our HiFormer performs much better than other competitors w.r.t. most of the evaluation metrics. Specifically, the superiority of HiFormer across different datasets highlights its satisfactory generalization ability. We also show a visual comparison of the skin lesion segmentation results in Fig. 3 which indicates that our proposed method is able to capture finer structures and generates more precise contours. Specifically, as in Fig. 3, our approach performs better than hybrid methods such as TMU-Net [40] in boundary areas. Moreover, showcased in Fig. 3, HiFormer is robust to noisy items compared to pure transformer-based methods such as Swin-Unet [10], where the performance degrades due to lack of locality modeling. The superior performance is achieved by an expedient combination of transformer and CNN for modeling global relationships and local representations.

4.3.3 Results of Multiple Myeloma Segmentation

In Table 4, we include the results based on the mean IoU metric. The HiFormer structure consistently outperformed the challenge leader board in all configurations we tested. In addition, some segmentation outputs of the proposed HiFormer are illustrated in Fig. 4. As shown, our predictions adjust well to the provided GT masks. One of the key advantages of HiFormer is its ability to model multi-scale representation. It restrains the background noise, which is the

case in datasets with highly overlapped backgrounds (such as SegPC). Stated succinctly, HiFormer exceeds CNN-based methods with only local information modeling ability and transformer-based counterparts, which render poor performance in boundary areas.

4.4. Comparison of Model Parameters

In 5, we compare the numbers of parameters of our proposed method with those of medical image segmentation models. Our lightweight HiFormer shows great superiority in terms of model complexity while attaining eminent or on-par performance compared to the literature works.

5. Ablation Study

Comparison of different CNN backbones. We first investigate the contribution of different CNN backbones. Specifically, we employ variants of ResNet [28] and DenseNet [29] as two prior arts of convolutional architectures. As shown in Table 7, utilizing the ResNet backbone results in the best performance. Moreover, we have witnessed that a larger CNN backbone does not necessarily result in a performance boost (see row 3, 4 in Table 7), which gives us the insight to use ResNet50 architecture as the default.

Impact of the DLF module. Next, we evaluate the importance of the DLF module on segmentation performance. The experimental results reported in Table 6 reveal the non-negligible role of the DLF module during the encoding and decoding process. Specifically, the DLF module brings significant improvements (3.24% and 2.18%) to the dice score and HD, respectively. Through the cross-attention mechanism, the DLF module assists the network in incorporating global and local features. The results prove that an expedient combination of CNN and transformer is helpful in segmenting target lesions. In addition, the impact of the DLF module on the SegPc and Skin datasets is provided in the Supplementary Material (SM) (see Table 1-2).

Ablation on different DLF module configurations. Table 8 shows the performance of different DLF module configurations. We test different values for the number of heads

Table 2: Comparison results of the proposed method on the *Synapse* dataset. **Blue** indicates the best result, and **red** displays the second-best.

Methods	DSC \uparrow	HD \downarrow	Aorta	Gallbladder	Kidney(L)	Kidney(R)	Liver	Pancreas	Spleen	Stomach
DARR [22]	69.77	-	74.74	53.77	72.31	73.24	94.08	54.18	89.90	45.96
R50 U-Net [13]	74.68	36.87	87.74	63.66	80.60	78.19	93.74	56.90	85.87	74.16
U-Net [41]	76.85	39.70	89.07	69.72	77.77	68.60	93.43	53.98	86.67	75.58
R50 Att-UNet [13]	75.57	36.97	55.92	63.91	79.20	72.71	93.56	49.37	87.19	74.95
Att-UNet [42]	77.77	36.02	89.55	68.88	77.98	71.11	93.57	58.04	87.30	75.75
R50 ViT [13]	71.29	32.87	73.73	55.13	75.80	72.20	91.51	45.99	81.99	73.95
TransUnet [13]	77.48	31.69	87.23	63.13	81.87	77.02	94.08	55.86	85.08	75.62
Swin-Unet [10]	79.13	21.55	85.47	66.53	83.28	79.61	94.29	56.58	90.66	76.60
LeVit-Unet [52]	78.53	16.84	78.53	62.23	84.61	80.25	93.11	59.07	88.86	72.76
DeepLabv3+ (CNN) [16]	77.63	39.95	88.04	66.51	82.76	74.21	91.23	58.32	87.43	73.53
HiFormer-S	80.29	18.85	85.63	73.29	82.39	64.84	94.22	60.84	91.03	78.07
HiFormer-B	80.39	14.70	86.21	65.69	85.23	79.77	94.61	59.52	90.99	81.08
HiFormer-L	80.69	19.14	87.03	68.61	84.23	78.37	94.07	60.77	90.44	82.03

Table 3: Performance comparison of the proposed method against the SOTA approaches on skin lesion segmentation benchmarks. **Blue** indicates the best result, and **red** displays the second-best.

Methods	ISIC 2017				ISIC 2018				PH ²			
	DSC	SE	SP	ACC	DSC	SE	SP	ACC	DSC	SE	SP	ACC
U-Net [41]	0.8159	0.8172	0.9680	0.9164	0.8545	0.8800	0.9697	0.9404	0.8936	0.9125	0.9588	0.9233
Att-UNet [42]	0.8082	0.7998	0.9776	0.9145	0.8566	0.8674	0.9863	0.9376	0.9003	0.9205	0.9640	0.9276
DAGAN [32]	0.8425	0.8363	0.9716	0.9304	0.8807	0.9072	0.9588	0.9324	0.9201	0.8320	0.9640	0.9425
TransUNet [13]	0.8123	0.8263	0.9577	0.9207	0.8499	0.8578	0.9653	0.9452	0.8840	0.9063	0.9427	0.9200
MCGU-Net [1]	0.8927	0.8502	0.9855	0.9570	0.8950	0.8480	0.9860	0.9550	0.9263	0.8322	0.9714	0.9537
MedT [45]	0.8037	0.8064	0.9546	0.9090	0.8389	0.8252	0.9637	0.9358	0.9122	0.8472	0.9657	0.9416
FAT-Net [49]	0.8500	0.8392	0.9725	0.9326	0.8903	0.9100	0.9699	0.9578	0.9440	0.9441	0.9741	0.9703
TMU-Net [40]	0.9164	0.9128	0.9789	0.9660	0.9059	0.9038	0.9746	0.9603	0.9414	0.9395	0.9756	0.9647
Swin-Unet [10]	0.9183	0.9142	0.9798	0.9701	0.8946	0.9056	0.9798	0.9645	0.9449	0.9410	0.9564	0.9678
DeepLabv3+ (CNN) [16]	0.9162	0.8733	0.9921	0.9691	0.8820	0.8560	0.9770	0.9510	0.9202	0.8818	0.9832	0.9503
HiFormer-S	0.9238	0.9153	0.9832	0.9695	0.9079	0.8934	0.9801	0.9618	0.9455	0.9737	0.9604	0.9646
HiFormer-B	0.9253	0.9155	0.9840	0.9702	0.9102	0.9119	0.9755	0.9621	0.9460	0.9420	0.9772	0.9661
HiFormer-L	0.9225	0.9046	0.9856	0.9693	0.9053	0.8828	0.9820	0.9611	0.9451	0.9561	0.9691	0.9659

Table 4: Performance evaluation on the *SegPC* challenge.

Methods	mIOU
Frequency recalibration U-Net [3]	0.9392
XLAB Insights [7]	0.9360
DSC-IITISM [7]	0.9356
Multi-scale attention deeplabv3+ [7]	0.9065
U-Net [41]	0.7665
Contextual attention [40]	0.9395
HiFormer-S	0.9392
HiFormer-B	0.9406
HiFormer-L	0.9395

and depth (S and L) for both the small and large levels and the MLP expanding ratio in the MLP block of the transformer module (r). We observe that the pair of (2, 1) for (S, L) and six heads for both levels work best. As shown in row A, increasing the number of heads does not necessarily improve performance. Additionally, the expanding ratio (r) plays a significant role in the performance. Compared to row C, doubling r results in a 1.04% increase in DSC and a 1.82% drop in HD. More information regarding the

Table 5: Comparison of model parameters.

Model	# Params (M)	DSC \uparrow	HD \downarrow
TransUnet	105.28	77.48	31.69
Swin-Unet	27.17	79.13	21.55
LeVit-Unet	52.17	78.53	16.84
DeepLabv3+ (CNN)	59.50	77.63	39.95
HiFormer-S	23.25	80.29	18.85
HiFormer-B	25.51	80.39	14.70
HiFormer-L	29.52	80.69	19.14

Table 6: Impact of the DLF module on the *Synapse* dataset.

Model	DLF	DSC \uparrow	HD \downarrow
HiFormer-B	X	77.15	16.88
HiFormer-B	✓	80.39	14.70

technical design of the DLF module is provided in the SM. **Ablation on feature consistency.** We conduct two experiments to measure and demystify the feature consistency and discuss them in detail in the SM. First, we present the feature visualization of each level before and after involving the DLF module (SM, Fig. 1-2). The second experiment

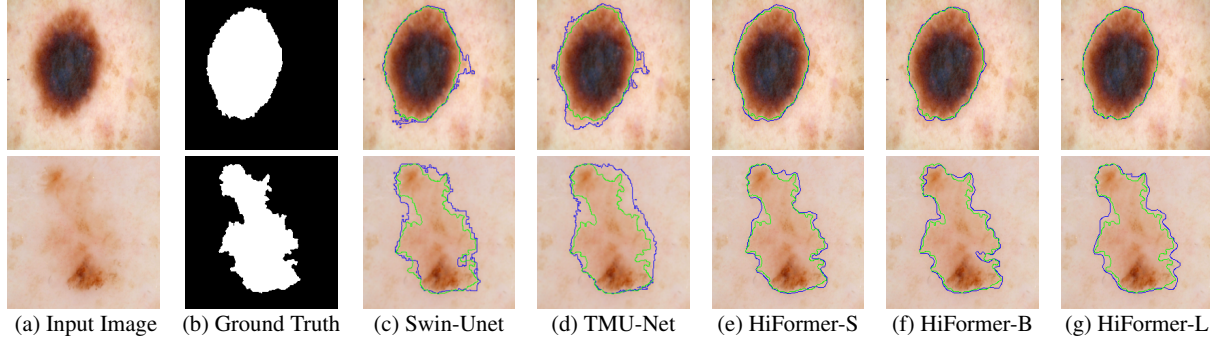


Figure 3: Visual comparisons of different methods on the *ISIC2017* skin lesion segmentation dataset. Ground truth boundaries are shown in green, and predicted boundaries are shown in blue.




Figure 4: Visual representation of the proposed method on the *SegPC* cell segmentation dataset.

Table 7: Comparison of different backbones for the CNN module on the *Synapse* dataset. Except for the CNN module, all configurations are identical to HiFormer-B.

Model	# Params (M)	DSC \uparrow	HD \downarrow
HiFormer+ResNet18	19.36	77.15	16.88
HiFormer+ResNet34	24.75	79.39	22.71
HiFormer+ResNet50	25.51	80.39	14.70
HiFormer+ResNet101	44.50	79.42	17.18
HiFormer+DenseNet121	23.92	78.65	16.18
HiFormer+DenseNet169	29.55	78.73	15.94
HiFormer+DenseNet201	35.36	79.08	21.30

proves how applying each module aids with feature consistency (SM, Table 3). Overall, the contribution of each module in providing more consistent features can be inferred from the results.

6. Discussion

Our comprehensive experiments on different medical image segmentation datasets demonstrate the effectiveness of our proposed HiFormer model compared to CNN and transformer-based approaches. The key advances of our approach are two folds. The first rationality of its design is combining CNN and transformer both in the shallow layers of the network. Second, the skip-connection module provides feature reusability and blends CNN local features with global features provided by the transformer module. The quantitative view of the HiFormer network on five challeng-

Table 8: Ablation study for the DLF module with different parameters on the *Synapse* dataset. For a fair comparison, ResNet-50 has been used as the CNN module in all the configurations, and r denotes the MLP expanding ratio used in the transformer block of the DLF module.

Model	Dimension P ^s	P ^l	S	L	r	# Heads P ^s	P ^l	DSC \uparrow	HD \downarrow	Params (M)
HiFormer-B	384	96	2	1	2	6	6	80.39	14.70	25.51
A	384	96	2	1	2	12	6	79.00	15.81	25.51
B	384	96	2	1	2	3	3	77.95	19.11	25.51
C	384	96	2	1	1	6	6	79.35	16.52	24.90
D	384	96	2	1	3	6	6	79.22	17.96	26.12
E	384	96	1	1	2	6	6	79.48	20.15	24.33
F	384	96	2	2	2	6	6	78.86	19.75	25.59

ing datasets reveals that it can perform segmentation well, surpassing the SOTA methods in most cases. From the perspective of visual analysis, Fig. 2 illustrates noise-less segmentation of organs such as the Liver and Kidney, which is also consistent with quantitative benchmarks. In contrast, our model acquires failure cases in some cases (e.g., Aorta), which again agrees with numerical results. Moreover, it is perceived that the low-contrast skin images still bring great difficulties for our model. In general, HiFormer has shown the potential to learn the critical anatomical relationships represented in medical images effectively. In terms of model parameters, HiFormer is a lightweight model compared with other complex models, which impose serious problems in medical image segmentation.

7. Conclusions

In this paper, we introduce HiFormer, a novel hybrid CNN-transformer-based method for medical image segmentation. Specifically, we combine the global features obtained from a Swin Transformer module with local representations of a CNN-based encoder. Then, using a DLF module, we attain a finer fusion of features derived from the aforementioned representations. We achieve superior performance over CNN-based, vanilla transformer-based, and hybrid models indicating that our methodology secures the balance in keeping the details of low-level features and modeling the long-range interactions.

References

- [1] 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.
- [2] Reza Azad, Maryam Asadi-Aghbolaghi, Mahmood Fathy, and Sergio Escalera. Bi-directional convlstm u-net with densley connected convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, pages 0–0, 2019.
- [3] Reza Azad, Afshin Bozorgpour, Maryam Asadi-Aghbolaghi, Dorit Merhof, and Sergio Escalera. Deep frequency recalibration u-net for medical image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3274–3283, 2021.
- [4] 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.
- [5] Reza Azad, Nika Khosravi, and Dorit Merhof. Smu-net: Style matching u-net for brain tumor segmentation with missing modalities. *arXiv preprint arXiv:2204.02961*, 2022.
- [6] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.
- [7] Afshin Bozorgpour, Reza Azad, Eman Showkatian, and Alaa Sulaiman. Multi-scale regional attention deeplab3+: Multiple myeloma plasma cells segmentation in microscopic images. *arXiv preprint arXiv:2105.06238*, 2021.
- [8] Sijing Cai, Yunxian Tian, Harvey Lui, Haishan Zeng, Yi Wu, and Guannan Chen. Dense-unet: a novel multiphoton in vivo cellular image segmentation model based on a convolutional neural network. *Quantitative imaging in medicine and surgery*, 10(6):1275, 2020.
- [9] Zhaowei Cai, Quanfu Fan, Rogerio S Feris, and Nuno Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. In *European conference on computer vision*, pages 354–370. Springer, 2016.
- [10] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. *arXiv preprint arXiv:2105.05537*, 2021.
- [11] 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.
- [12] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 357–366, 2021.
- [13] 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.
- [14] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.
- [15] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [16] 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.
- [17] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kallou, 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.
- [18] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kallou, 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.
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [20] Yuxin Fang, Bencheng Liao, Xinggang Wang, Jiemin Fang, Jiyang Qi, Rui Wu, Jianwei Niu, and Wenyu Liu. You only look at one sequence: Rethinking transformer in vision through object detection. *Advances in Neural Information Processing Systems*, 34, 2021.
- [21] Abdur R Feyjie, Reza Azad, Marco Pedersoli, Claude Kauffmann, Ismail Ben Ayed, and Jose Dolz. Semi-supervised few-shot learning for medical image segmentation. *arXiv preprint arXiv:2003.08462*, 2020.
- [22] Shuhao Fu, Yongyi Lu, Yan Wang, Yuyin Zhou, Wei Shen, Elliot Fishman, and Alan Yuille. Domain adaptive relational reasoning for 3d multi-organ segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 656–666. Springer, 2020.
- [23] Jiaqi Gu, Hyounjun Kwon, Dilin Wang, Wei Ye, Meng Li, Yu-Hsin Chen, Liangzhen Lai, Vikas Chandra, and David Z Pan. Multi-scale high-resolution vision transformer for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12094–12103, 2022.

- [24] Ruohao Guo, Dantong Niu, Liao Qu, and Zhenbo Li. Sotr: Segmenting objects with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7157–7166, 2021.
- [25] Anubha Gupta, Ritu Gupta, Shiv Gehlot, and Shubham Goswami. Segpc-2021: Segmentation of multiple myeloma plasma cells in microscopic images, 2021.
- [26] Anubha Gupta, Pramit Mallick, Ojaswa Sharma, Ritu Gupta, and Rahul Duggal. Pcseg: Color model driven probabilistic multiphase level set based tool for plasma cell segmentation in multiple myeloma. *PLoS one*, 13(12):e0207908, 2018.
- [27] 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.
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [29] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [30] Huimin Huang, Lanfen Lin, Ruofeng Tong, Hongjie Hu, Qiaowei Zhang, Yutaro Iwamoto, Xianhua Han, Yen-Wei Chen, and Jian Wu. 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)*, pages 1055–1059. IEEE, 2020.
- [31] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6399–6408, 2019.
- [32] 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.
- [33] Ailiang Lin, Bingzhi Chen, Jiayu Xu, Zheng Zhang, Guangming Lu, and David Zhang. Ds-transunet: Dual swin transformer u-net for medical image segmentation. *IEEE Transactions on Instrumentation and Measurement*, 2022.
- [34] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [35] 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.
- [36] 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.
- [37] 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.
- [38] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016.
- [39] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015.
- [40] Azad Reza, Heidari Moein, Wu Yuli, and Merhof Dorit. Contextual attention network: Transformer meets u-net. *arXiv preprint arXiv:2203.01932*, 2022.
- [41] 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.
- [42] 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.
- [43] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7262–7272, 2021.
- [44] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- [45] 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.
- [46] Jeya Maria Jose Valanarasu, Vishwanath A Sindagi, Ilker Hacihaliloglu, and Vishal M Patel. Kiu-net: Towards accurate segmentation of biomedical images using over-complete representations. In *International conference on medical image computing and computer-assisted intervention*, pages 363–373. Springer, 2020.
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [48] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 568–578, 2021.
- [49] Huisi Wu, Shihuai Chen, Guilian Chen, Wei Wang, Baiying Lei, and Zhenkun Wen. Fat-net: Feature adaptive transform-

- ers for automated skin lesion segmentation. *Medical Image Analysis*, 76:102327, 2022.
- [50] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
 - [51] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4794–4803, 2022.
 - [52] 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.
 - [53] Songfan Yang and Deva Ramanan. Multi-scale recognition with dag-cnns. In *Proceedings of the IEEE international conference on computer vision*, pages 1215–1223, 2015.
 - [54] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
 - [55] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 558–567, 2021.
 - [56] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
 - [57] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 3–11. Springer, 2018.
 - [58] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.