# TransResNet: Integrating the Strengths of ViTs and CNNs for High Resolution Medical Image Segmentation via Feature Grafting

Muhammad Hamza Sharif muhammad.sharif@mbzuai.ac.ae Dmitry Demidov dmitry.demidov@mbzuai.ac.ae Asif Hanif asif.hanif@mbzuai.ac.ae Mohammad Yaqub mohammad.yaqub@mbzuai.ac.ae

Min Xu 🖂 xumin100@gmail.com Mohamed Bin Zayed University of Artificial Intelligence, UAE

#### Abstract

High-resolution images are preferable in medical imaging domain as they significantly improve the diagnostic capability of the underlying method. In particular, high resolution helps substantially in improving automatic image segmentation. However, most of the existing deep learning-based techniques for medical image segmentation are optimized for input images having small spatial dimensions and perform poorly on high-resolution images. To address this shortcoming, we propose a parallel-in-branch architecture called TransResNet, which incorporates Transformer and CNN in a parallel manner to extract features from multi-resolution images independently. In TransResNet, we introduce Cross Grafting Module (CGM), which generates the grafted features, enriched in both global semantic and low-level spatial details, by combining the feature maps from Transformer and CNN branches through fusion and self-attention mechanism. Moreover, we use these grafted features in the decoding process, increasing the information flow for better prediction of the segmentation mask. Extensive experiments on ten datasets demonstrate that TransResNet achieves either state-of-the-art or competitive results on several segmentation tasks, including skin lesion, retinal vessel, and polyp segmentation. The source code and pre-trained models are available at https://github.com/Sharifmhamza/TransResNet.

## **1** Introduction

Segmentation is a fundamental problem in the domain of computer vision with numerous practical applications, particularly in biomedical imaging analysis. Medical segmented images can be used in a wide range of applications, such as disease localization [33], tissue volume estimation [40], and studying anatomical structure [29]. Accurate and precise segmentation of medical images is a challenging task due to the nature of the complexity of 2D

and 3D structures. Recent studies have demonstrated deep learning-based techniques as a powerful building block to accomplish this task accurately [2, 5, 32].

Resolution of an image in medical diagnosis plays an important role. In general, highresolution images improve the results of a diagnostic method to determine the presence of certain diseases. A high-resolution image contains rich semantic information and it provides better chances for the extraction of useful information for a downstream task e.g. segmentation [17]. Many deep learning-based approaches have been proposed to perform automatic medical image segmentation, such as segmenting organs [11], lesions [39], and tumors [15]. However, these existing techniques are primarily designed to segment low-resolution (small spatial dimension) images, and they do not provide favorable results on high-resolution images due to discrepancies between sampling depth and receptive field size. With rapid technological revolution, medical image-capturing devices have undergone extensive modifications and advancements in recent years. Compared to the previous appurtenances, these modern devices are capable of capturing images with higher resolution. This requires the demand for deep learning-based segmentation framework that can process high-resolution medical images efficiently and performs favorably.

Encoder-decoder based convolutional neural network (CNN) architectures have achieved unprecedented performance in medical image segmentation for low resolution input images [1, 27, 47]. Despite their impressive success, these approaches are still facing challenges to capture global context details due to narrow and fixed receptive field. Similarly, vision transformers (ViTs) [6, 21], which are efficient for modeling long-range dependencies and highly parallelizable, are computationally prohibitive and have to down-sample the image before processing. Because of the shortcomings of these architectures regarding high-resolution images, the substantial solution is to design a single architecture that collectively captures rich local and global information without increasing the computational complexity associated with high-resolution images and gives accurate segmentation results.

Inspired by deep learning-based high saliency object detection methods [36, 42, 44] for natural images, we propose an architecture for high-resolution segmentation of medical images named **TransResNet**, as shown in Fig. 1. In this paper, we use two encoder modules, one is CCN-based for extracting local feature details, the other is transformer-based for grasping global information. We introduce a *Cross Grafting Module (CGM)* to combine the features maps with similar spatial size from both encoder branches. CGM generates *grafted features* which are enriched in both local and global semantic cues. We use these grafted features in the decoding process for the prediction of segmentation masks. In summary, our main contributions are as follows:

- We propose a framework named **TransResNet** for efficient segmentation of high-resolution medical images by using two encoder backbones.
- We introduce *Cross-Grafting Module (CGM)*, to combine the low-level spatial features (from the CNN branch) and high level semantic information (from the Transformer branch) through fusion and self-attention mechanism.
- We have performed extensive experiments on ten datasets for three medical image segmentation tasks. Our experimental results demonstrate that the proposed approach outperforms state-of-the-art (SOTA) methods on high-resolution medical imaging datasets and competes comparably on datasets containing mixture of low and high resolution images.



Figure 1: An overview of the architecture of TransResNet for high-resolution medical image segmentation. Our TransResNet uses the parallel branches from Swin-transformer and Resnet-18 backbones as encoders. The core module of our architecture is Cross Grafting Module (CGM), explained briefly in the Fig. 2. The decoder module aggregates the flow of feature input maps from swin block, CGM block, and resnet block. D1, D2, and D3 are subblocks of the decoder with their structure on the right side.

## 2 Related Work

Segmentation is not a trivial problem, especially in the biomedical imaging domain. A lot of studies have been proposed in the past that focus on low-level image information to predict the segmentation mask. The inefficiency of these methods in capturing rich semantics makes their performance inconsistent in complex settings. In this regard, we discuss a few works on biomedical image segmentation.

**Medical Image Segmentation with CNN-based Architectures:** Extensive studies have been proposed predominantly with various CNN-based architectures for medical image segmentation tasks. Some studies have designed the model architecture in an encoder-decoder style, for example, U-net [30] and its variants [1, 27, 47], while other works have integrated CNN extracted features with a self-attention mechanism in the decoder module to boost the network performance for capturing global interaction [7, 41]. Despite their excellent performance, these approaches are limited in their ability to capture global semantic information due to narrow receptive fields, as the kernel size of CNN-based techniques is typically fixed, making it more challenging to predict segmentation masks accurately. Our work proposes a method which allows the receptive field of the CNN backbone to capture rich semantic information from high-resolution input images.

**Medical Image Segmentation with Vision Transformers:** Transformers, which were primarily developed for the natural language processing task [37], have made marvelous achievements in the field of computer vision for downstream tasks [6, 21]. Despite transformer's powerful global modeling capabilities, they are limited in explicitly capturing rich semantic details that are essential for biomedical imaging analysis [4]. The first transformer-based medical image segmentation framework named TransUnet, proposed by Chen *et al.* [4], uses a transformer backbone in the U-Net style to extract global features in the encoder block and upsample these encoded features in the decoder block. In the TransFuse framework, Zhang *et al.* [46] combines CNN and Transformer in a parallel manner to grasp the local and global information for a similar task. Despite their success, these methods are suitable for low-resolution images and tend to ignore local semantic information in high-resolution images. Thus, our proposed architecture aims to mitigate this issue by incorporating novel *Cross Grafting Module (CGM)* which captures both global and rich semantic details in high-resolution medical images.

## 3 Methodology

The concrete architecture of proposed network is shown in the Fig. 1. The proposed network follows the encoder-decoder design architecture consisting of two encoders and one decoder. The encoder includes the ResNet-18 [16] and Swin-B [21] as backbones, while decoding phase comprises of three sub-stages. The feature maps from both encoders are grafted in a Cross Grafting Module (CGM), which emphasises on the salient regions, and make network learn precise pixel-level details. We will discuss each part in the following subsections.

#### 3.1 Encoder Module

As explained earlier, our proposed architecture is based upon two encoding streams: a CNN and a vision transformer (ViT). The main reason for using two encoding streams is to capture both local and global information, which makes the network learn salient features more accurately. The CNN-based encoder captures the low-level feature representations from high-resolution input images, while the ViT-based encoder is used to learn the global semantic information from low-resolution input images as demonstrated in Fig. 1. During the encoding phase, images  $(I \in \mathbb{R}^{H \times W \times C})$  having different spatial dimensions are passed to both encoders:  $I_R \in \mathbb{R}^{1024 \times 1024 \times 3}$  and  $I_S \in \mathbb{R}^{224 \times 224 \times 3}$  are passed through ResNet-18 and Swin-B encoders respectively. Using two encoder networks with multi-scale input size consumes massive amount of computational resources to generate feature maps. To handle this issue, we have discarded some layers from Resnet-18, and Swin-B network. With prior knowledge from the literature [16], we know that there are five feature maps generated by Resnet-18, denoted as  $\{R_{f_i} | i \in (1,2,3,4,5)\}$ . The first stage of Resnet-18 uses a large kernel size of  $7 \times 7$  to extract feature maps. In our case, the input  $(I_R \in \mathbb{R}^{1024 \times 1024 \times 3})$  passing through Resnet-18, is larger in size, thus resulting in huge demand of computational capacity to generate feature maps for this stage. We relinquish this stage while keeping the last four stages of Resnet-18,  $\{R_{f_i} | i \in (2,3,4,5)\}$ , which learn more complex features and become computationally less expensive due to gradual down-sampling of the feature maps at each stage, resulting in  $\{\mathbf{R}_{f_i} \in \mathbb{R}^{\frac{H}{2^i} \times \frac{W}{2^i} \times (32 \times 2^{i-1})}\}_{i=2}^5$ . In a similar fashion, we also follow similar approach with the Swin-B transformer. As there are four stages in Swin-B, we drop the last stage after the patch merging block. We utilize the feature maps generated from the output of first three stages and patch merging block of the fourth stage of Swin-B, denoted as  $\{S_{f_i} | i \in (1,2,3,4)\}$ , with their feature embedding dimension represented as  $\{\mathbf{S}_{f_i} \in \mathbb{R}^{\frac{56}{2^{l-1}} \times \frac{56}{2^{l-1}} \times (64 \times 2^l)}\}_{i=1}^3$ , and  $\mathbf{S}_{f_4} \in \mathbb{R}^{14 \times 14 \times 512}$ . As the spatial dimensions of feature maps  $R_{f_5}$  with dimension (32 × 32 × 512), and  $S_{f_2}$  with dimension (28 × 28 × 256), are very close to each other, we select these features for grafting in the CGM.

#### **3.2** Cross Grafting Module (CGM)

We propose a new module, *Cross Grafting Module (CGM)*, in such way that network effectively adapts the local and global semantic representations. To accomplish this task, we select the features maps  $R_{f_5} \in \mathbb{R}^{H' \times W' \times C'}$ , and  $S_{f_2} \in \mathbb{R}^{H'' \times W'' \times C''}$  extracted from the Resnet-18 and Swin-B encoders respectively. As transformers have shown excellence performance at modeling long range relationships [6],  $S_{f_2}$  is responsible for providing global semantic detail, while  $R_{f_5}$  contributes towards local information due to CNN's excellent low-level



Figure 2: An overview of the architecture of the proposed Cross Grafting Module (CGM). The CGM module takes dual input i.e., the feature maps from Swin-transformer, and Resnet branches, and outputs the grafted features through fusion and self-attention mechanism. These grafted features are used in the decoding process. The module also generates a cross-transposed attention matrix (CTAM), which is used in the objective function.

feature learning capabilities. But the major issue is the discrepancy between receptive field of feature maps, as  $R_{f_5}$  is down-sampled to match with receptive dimension of  $S_{f_2}$ , thus produces the noisy output.

To alleviate this issue, we apply operator  $\mathcal{A}$  (defined below) on the feature maps ( $\mathbf{R}_{f_5} \in \mathbb{R}^{H' \times W' \times C'}$ ,  $\mathbf{S}_{f_2} \in \mathbb{R}^{H'' \times W'' \times C''}$ ) and obtain ( $\hat{\mathbf{R}}_f \in \mathbb{R}^{1 \times H' W' C'}$ ,  $\hat{\mathbf{S}}_f \in \mathbb{R}^{1 \times H'' W'' C''}$ ) (see Eq. 1). Next, we apply layer normalization on ( $\hat{\mathbf{R}}_f$ ,  $\hat{\mathbf{S}}_f$ ) and generate *query* (Q), *key* (K), and *value* (V) projections from transformer and CNN branches separately (see Eq. 2 and Eq. 3). By fusing (element-wise addition) the tensors in each tuple i.e. ( $Q_R, Q_S$ ), ( $K_R, K_S$ ), ( $V_R, V_S$ ), we achieve the resultant tensors with enriched local details, thus mitigating the effects of noise. For learning of global semantic information, we apply the self-attention (SA) mechanism that efficiently calculates the point-wise relationship between these resultant tensors. Grafted features (Z) i.e. the output of CGM, is used in the decoding process as shown in Fig. 1, but spatial dimension of SA output (X) has contradiction with the input dimension of decoder flow. To make it identical, we apply linear projection layer to SA output, reshape it to original size and feed it to convolution layer. During the whole process of grafting, the spatial dimension flow and facilitate the training process, we use two skip connections before final output (Z) as shown in the Fig. 2. Overall, the whole grafting process with intermediate steps is expressed as follows:

$$\hat{\mathbf{R}}_{f} = \mathcal{A}(\mathbf{R}_{f_{5}}) \quad ; \quad \hat{\mathbf{S}}_{f} = \mathcal{A}(\mathbf{S}_{f_{2}}), 
\hat{\mathbf{R}} = LN(\hat{\mathbf{R}}_{f}) \quad ; \quad \hat{\mathbf{S}} = LN(\hat{\mathbf{S}}_{f}),$$
(1)

$$Q_R = W^Q \hat{R} \quad ; \quad K_R = W^K \hat{R} \quad ; \quad V_R = W^V \hat{R}, Q_S = W^Q \hat{S} \quad ; \quad K_S = W^K \hat{S} \quad ; \quad V_S = W^V \hat{S},$$
(2)

$$Q = Q_R + Q_S$$
;  $K = K_R + K_S$ ;  $V = V_R + V_S$ , (3)

$$\begin{aligned} \mathbf{X} &= \mathbf{V} \cdot softmax \left( \mathbf{Q} \cdot \mathbf{K}^{\top} / \boldsymbol{\alpha} \right), \\ \mathbf{Y} &= linear(\mathbf{X}) + pool(\hat{\mathbf{R}}_{f} \oplus \hat{\mathbf{S}}_{f}), \\ \mathbf{Z} &= \mathbf{Y} + conv(\mathbf{Y}), \end{aligned} \tag{4}$$

where  $R_{f_5}$  and  $S_{f_2}$  are the input feature maps and  $\mathcal{A} = \text{Flatten} \leftarrow \text{GELU} \leftarrow \text{BN} \leftarrow \text{Conv}(\cdot)$ is an operator that takes an input and performs convolution, batch normalization, GELU activation and flattening layer sequentially. Grafted feature (Z) is the output of CGM used in the decoding process. Here,  $(W^Q, W^K, W^V)$  are the weights of linear layers,  $\alpha$  is a scaling parameter used to control the magnitude of dot product of Q and K tensors, and (*LN*, *linear*, *pool*) represent layer-normalization, linear and pooling layers respectively.  $\oplus$  denotes the concatenation used to amalgamate the flatten output of transformer and resnet branches. Additionally, the dot product interaction of query and key projections generates the attention matrix  $\mathbf{A} = \text{softmax} (\mathbf{Q} \cdot \mathbf{K}^{\top} / \alpha)$ , which is transposed and added to the itself to form a cross transposed attention matrix which is defined by Eq. 5 as follows:

$$\mathbf{CTAM} = \mathbf{GELU}(\mathbf{BN}(\mathbf{Conv}(\mathbf{A} + \mathbf{A}^{\top}))).$$
(5)

CTAM matrix is used in the objective function (see Sec. 3.4).

#### **3.3 Decoder Module**

The flow of feature maps in the decoder block of our proposed architecture is illustrated in Fig. 1. The decoder module first receives the flow of features from the Swin-B branch, followed by the feature grafting module, and finally the ResNet-18 branch in a staggered pattern. The decoder block is divided into chunks of three different sub-blocks, denoted as **D1**, **D2**, and **D3**, which aggregate the flow of features from these ResNet-18, Swin-B and CGM branches.

#### **3.4 Objective Function**

The entire network is trained end-to-end with the joint objective function, which includes the segmentation loss  $L_{seg}$  for segmentation maps, attention loss  $L_{att}$  for the cross-transposed attention matrix (**CTAM**) map , and auxiliary loss  $L_{aux}$  for deep supervision to improve the gradient flow by supervising the ResNet-18 and Swin-B transformer branches, which is defined as follows:

$$L_{seg} = \frac{1}{2} \{ \phi_{bce}(\hat{M}_{pm}, M_{gt}) + \phi_{iou}(\hat{M}_{pm}, M_{gt}) \},$$
(6)

$$L_{att} = \phi_{bce}^{w} (\mathbf{CTAM}_{map}, M_{gt_{map}}), \tag{7}$$

$$L_{aux} = \frac{1}{2} \{ \phi_{bce}(\hat{M}_R, M_{gt}) + \phi_{iou}(\hat{M}_R, M_{gt}) \} + \frac{1}{2} \{ \phi_{bce}(\hat{M}_S, M_{gt}) + \phi_{iou}(\hat{M}_S, M_{gt}) \},$$
(8)

$$L_{total} = L_{seg} + L_{att} + \lambda L_{aux}, \tag{9}$$

where  $\phi_{bce}$ ,  $\phi_{iou}$ , and  $\phi_{bce}^w$  denote the binary cross entropy function, intersection-over-union function, and weighted binary cross entropy function respectively. Here,  $M_{gt}$ ,  $\hat{M}_{pm}$ ,  $(\hat{M}_R, \hat{M}_S)$ , and  $M_{gt_{map}}$  are ground truth mask, predicted segmentation mask, salient prediction maps extracted from the transformer, and resnet branches which are used in the grafting module, and attention matrix map generated from ground truth. Ground truth attention matrix map  $M_{gt_{map}}$  is achieved by matching the shape of the ground truth mask  $(M_{gt})$  to that of **CTAM** via down-sampling, flattening, and taking the self dot-product of flattened vector. The  $\lambda$  is the weight parameter used to balance the auxiliary loss  $L_{aux}$  calculated from two encoder branches.

### 4 **Experiments**

We evaluate the effectiveness of our proposed model TransResNet on the three different segmentation tasks: (a) skin lesion segmentation (2 datasets), (b) retinal vessel segmentation

(3 datasets), and (c) polyp segmentation (5 datasets). More details on datasets, training settings, and additional quantitative and visual results have been presented in the following subsections. We have highlighted the best and the second-best highest scores in the result Sec. 4.4 with different evaluation metrics.

#### 4.1 Datasets

**Skin Lesion Segmentation:** In order to segment skin lesions, we have used two publicly available benchmark datasets with the majority of high-resolution images (resolution > 1K): ISIC-2016 [14], and PH2 [24]. The ISIC-2016 has the train-validation split of samples: 900/379, while the PH2 database includes 200 samples. As most of the SOTA methods use the same split, therefore, we also keep the same sample size for the fair evaluation of our model. To test the generalization ability and robustness of our method, we use PH2 dataset. **Retinal Vessel Segmentation:** The proposed method is evaluated using three publicly accessible retinal fundus imaging datasets: HRF [26], IOSTAR [45], and CHASE\_DB1 [9]. The HRF database has a total of 45 samples, each image with a resolution of  $3504 \times 2306$ , whereas IOSTAR contains 30, and CHASE\_DB1 has 28 image samples, respectively. For fair comparison, we also follow the same number of train-test split as mentioned in these papers [10, 19, 20, 25].

**Polyp Segmentation:** A total of five polyp segmentation benchmark datasets are used to evaluate the performance of TransResNet: Kvasir [18], CVC-ClinicDB [3], CVC-ColonDB [35], Endoscene [38], and ETIS [34]. To ensure fairness, we keep the same number of training and testing images as used in [7, 41, 46], i.e. 1450 training images from Kvasir, and CVC-ClinicDB, while 798 testing images from all of the five datasets. These datasets are highly versatile benchmark datasets with mixture of low and high resolution images.

#### 4.2 Implementation Details

We have implemented TransResNet using the Pytorch [28] framework and NVIDIA A100-SXM4 GPU with a maximum of 36GB of memory to accelerate the smooth training pipeline. All the input images have been resized to  $1024 \times 1024$ , and various data augmentations have been applied, including horizontal flip, vertical flip, rotation, and random brightness to increase the data diversity, volume, and avoid overfitting. The entire network is trained end-to-end with a Stochastic Gradient Optimizer (SGD) [31] algorithm with the initial learning rate of 0.03, which gradually decreases with cosine annealing [22]. We use different hyper-parameter settings of weight decay ( $5e^{-2}$  to  $7e^{-5}$ ), and momentum (0.9 to 0.999) for different datasets. Due to the scarcity of training data samples, we train the network for a large number of epochs, for example, for the retina segmentation task, the network is trained for 3000 epochs and 200 epochs for other tasks with batch size of 8 and 16. We also use the Probability Correction Strategy (PCS) [41] during inference to improve the final prediction. The detail of Probability Correction Strategy (PCS) have been provided in Appendix B.

### 4.3 Evaluation Metric

We evaluate the performance of our best model using standard medical image segmentation metrics, i.e. mean dice coefficient (mDice), mean Intersection-over-Union (mIoU), and mean F1 (mF1) scores.

## 4.4 Quantitative Results

We have evaluated TransResNet on three different segmentation tasks as discussed in Sec. 4.1, to demonstrate its effectiveness. We have compared our method with six SOTA methods for skin lesion segmentation, four for retinal vessel segmentation, and seven for polyp segmentation tasks.

**Results of Skin lesion Segmentation:** We report our results for the skin lesion segmentation task on two datasets and compare them with six SOTA methods. Table 1, shows that our model has achieved the highest performance on both validation ISIC-2016, and test-PH2 datasets on both evaluation metrics. Since these datasets have higher resolution images, the model has captured rich semantic information, resulting in enhanced performance. In addition, the samples from the PH2 dataset are not part of the training phase, which indicates that our model has better generalization ability and robustness towards skin lesion segmentation task than other SOTA approaches.

**Results of Retinal Vessel Segmentation:** We compare TransResNet with four SOTA methods on three high-resolution fundus imaging datasets for the retinal vessel segmentation task. Table 2, demonstrates that our method has surpassed all the other SOTA methods by a margin of 1.4%, 1.0%, and 0.9% in terms of mean F1 (mF1) score on HRF, IOSTAR, and CHASE\_DB1 datasets respectively, without applying any pre-processing technique on these datasets.

Methods	ISIC	-2016	test-PH2			
	$mIoU\uparrow$	mDice↑	$mIoU^{\uparrow}$	mDice↑		
U-Net [30]	0.825	0.878	0.739	0.836		
U-Net++ [47]	0.818	0.889	0.812	0.889		
Attn U-Net [27]	0.797	0.874	0.695	0.805		
CE-Net [13]	0.842	0.905	0.824	0.894		
CA-Net [12]	0.807	0.881	0.751	0.846		
TransFuse [46]	0.840	0.900	0.823	0.897		
TransResNet	0.843	0.907	0.831	0.905		

Table 1: Quantitative results on skin lesion segmentation datasets compared with six SOTA methods. The red and green color cells represent the highest and the second highest scores respectively. Performance is measured by mean Dice and mean IoU scores.

Methods	HRF	IOSTAR	CHASE	
	mF1↑	mF1↑	$mF1\uparrow$	
DRIU [23]	0.783	0.825	0.810	
HED [43]	0.783	0.825	0.810	
M2U-Net [20]	0.780	0.817	0.802	
U-Net [30]	0.788	0.812	0.812	
TransResNet	0.802	0.835	0.821	

Table 2: Quantitative results on retinal vessel segmentation datasets compared with four SOTA methods. The red and green color cells represent the highest and the second highest scores respectively. Performance is measured by the mean F1 score.

**Results of Polyp Segmentation:** The performance of TransResNet for polyp segmentation has been evaluated and compared with seven SOTA methods across five different benchmark datasets. The quantitative results are shown in Table 4. As highlighted from the scores, our proposed architecture did not surpass some SOTA methods on polyp segmentation tasks except for the ClinicDB dataset. In addition, the mean dice scores of our method on Kvasir and EndoScence datasets are also very close to the highest and second highest SOTA methods, i.e., SANet, and TransFuse, with a minor difference of 0.037 on Kvasir, and 0.028 on ColonDB datasets respectively. We also find that our method performs unfavorably and does not generalize on ColonDB and ETIS datasets because these datasets have lower image resolution. Our method is better suited for high resolution images. We have provided detailed information and analysis in Appendix A about datasets.

#### Model Performance on different Image Resolu-

**tions:** From the segmentation results of our proposed method, we have analyzed four cases regarding the model performance based on input images during training and inference, which is summarized in Table 3. Given the results, one can conclude that image resolution is a major issue in the context of model performance. Our model design is highly suitable for high resolution images.

Train Image	Test Image	Model	
Resolution	Resolution	Performance	
Higher	Higher	Increases	
Lower & upscaled	Higher	Decreases	
Higher	Lower & upscaled	Increases	
Lower & upscaled	Lower & upscaled	Slightly decreases	

Table 3: Analysis of model performance based on the image resolution during training and inference. Lower resolution images are upscaled to  $1024 \times 1024$ .

#### SHARIF ET AL: TRANSRESNET

Table 4: Quantitative results on polyp segmentation datasets compared with seven SOTA methods. The red and green color cells represent the highest and the second highest scores respectively. Performance is measured by mean Dice and mean IoU scores. "-" indicates results are not available.

Methods	Kva	sir ClinicDB		cDB	ColonDB		EndoScene		ETIS	
	<i>mDice</i> ↑	$mIoU\uparrow$	<i>mDice</i> ↑	$mIoU\uparrow$	mDice↑	$mIoU\uparrow$	mDice↑	$mIoU\uparrow$	mDice↑	$mIoU\uparrow$
U-Net [30]	0.818	0.746	0.823	0.750	0.512	0.444	0.710	0.627	0.398	0.335
U-Net++ [47]	0.821	0.743	0.794	0.729	0.483	0.410	0.707	0.624	0.401	0.344
ResUNet++ [1]	0.813	0.793	0.796	0.796	-	-	-	-	-	-
SFA [8]	0.723	0.611	0.700	0.607	0.469	0.347	0.467	0.329	0.297	0.217
PraNet [7]	0.898	0.840	0.899	0.849	0.712	0.640	0.871	0.797	0.628	0.567
SANet [41]	0.904	0.847	0.916	0.859	0.753	0.670	0.888	0.815	0.750	0.654
TransFuse [46]	0.918	0.868	0.918	0.868	0.773	0.696	0.902	0.833	0.733	0.659
TransResNet	0.881	0.824	0.917	0.861	0.685	0.604	0.874	0.804	0.564	0.493

#### 4.5 Qualitative Results

The Fig. 3 shows predicted segmentation masks for some input images for each segmentation task. From the predicted masks, we can easily conclude that our method has not only performed an accurate prediction but also suppressed background noise. Additional visual results have been provided in Appendix C.



Figure 3: Qualitative results on all three segmentation tasks. The figure shows an example image, ground truth (GT) and predicted (PRED) segmentation mask for the skin lesion segmentation task (row 1), the polyp segmentation task (row 2) and retinal vessel segmentation task (row 3).

### 4.6 Ablation Study

In order to evaluate the effectiveness of the proposed method, we analyze our approach using two ablative studies, which are as follows: (a) Effect of network performance by using different feature map pairs in Cross Grafting Module (CGM) and (b) Effect of network performance by eliminating each module from the proposed architecture.

Ablation Study for Grafted Features: For better design of Cross Grafting Module (CGM), we conduct an ablation study, in which we change the feature map pairs used in the grafted module and study its impact on the performance of the network. In Table 5, we present the quantitative results for two datasets: ISIC-2016 and PH2. From the Table 5, we observe that the performance gradually increases and then tends to decrease, indicating that the pair ( $\mathbf{R}_{f_5}$ ,  $\mathbf{S}_{f_2}$ ) is suitable for grafting. The main reason is that the spatial sizes of both feature maps are very close, and the information captured by each model corresponds to each other, increasing the network performance through grafting.

Ablation Study for Elimination each Module: The proposed framework comprises of Swin-B transformer, Resnet-18 backbone, and Cross Grafting Module (CGM) that are used

for learning both local and global features. To experimentally evaluate the effect and contribution of each module in the generalization performance, we selectively remove one of the modules, as shown in Table 6. The qualitative finding suggests that removing any of the modules from the architecture results in a performance degradation.

Feature Pairs	Spatial Dimensions	val-ISIC-2016		test-PH2	
		$mIoU\uparrow$	mDice↑	$mIoU\uparrow$	mDice↑
$(\mathbf{R}_{f_5}, \mathbf{S}_{f_1})$	(32×32, 56×56)	0.829	0.881	0.824	0.881
$(\mathbf{R}_{f_{5}}, \mathbf{S}_{f_{2}})$	(32×32, 28×28)	0.843	0.907	0.831	0.905
$(\mathbf{R}_{f_5}, \mathbf{S}_{f_3})$	(32×32, 14×14)	0.836	0.892	0.825	0.896
$(\mathbf{R}_{f_5}, \mathbf{S}_{f_4})$	(32×32, 14×14)	0.819	0.862	0.812	0.874

 Table 5: An ablation study for TransResNet on ISIC-2016, and PH2 datasets for selection of feature pairs from Swin-transformer and ResNet for grafting.

ResNet-18	Swin-B	CGM	val-ISI	C-2016	test-PH2		
			$mDice\uparrow$	$mIoU\uparrow$	$mDice^{\uparrow}$	$mIoU\uparrow$	
1	X	X	0.879	0.819	0.879	0.814	
x	1	x	0.881	0.821	0.884	0.806	
1	1	X	0.889	0.832	0.900	0.821	
1	1	1	0.907	0.843	0.905	0.831	

Table 6: An ablation study to analyze the effect on the overall performance of the proposed method by eliminating each module.

## 5 Conclusion

In this paper, we present the TransResNet architecture for the segmentation of high-resolution medical images. A key component of TransResNet is the *Cross Grafting Module (CGM)*, which is used to learn *grafted features* with rich semantic and global information, allowing accurate prediction of segmentation masks during decoding. Extensive experiments on ten different datasets for three medical segmentation tasks indicate that our architecture shows better results on high-resolution images. One of the main limitations of our architecture, is that it is computationally expensive. With our research work, we intend to introduce the scientific community with AI-based model for high-resolution medical image segmentation. This will open new directions for conducting research on this problem as the demand for learning-based models with capability of efficiently processing the high-resolution images is expected to rise. Future direction in this line of research includes extending the proposed method to multi-class medical image segmentation and making it computationally less expensive.

## References

- Md Zahangir Alom, Mahmudul Hasan, Chris Yakopcic, Tarek M Taha, and Vijayan K Asari. Recurrent residual convolutional neural network based on u-net (r2u-net) for medical image segmentation. *arXiv preprint arXiv:1802.06955*, 2018.
- [2] Saeid Asgari Taghanaki, Kumar Abhishek, Joseph Paul Cohen, Julien Cohen-Adad, and Ghassan Hamarneh. Deep semantic segmentation of natural and medical images: a review. *Artificial Intelligence Review*, 54(1):137–178, 2021.
- [3] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilariño. Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized medical imaging and graphics*, 43:99–111, 2015.
- [4] 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.
- [5] Nanqing Dong, Min Xu, Xiaodan Liang, Yiliang Jiang, Wei Dai, and Eric Xing. Neural architecture search for adversarial medical image segmentation. In *International*

*Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 828–836. Springer, 2019.

- [6] 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.
- [7] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Pranet: Parallel reverse attention network for polyp segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 263–273. Springer, 2020.
- [8] Yuqi Fang, Cheng Chen, Yixuan Yuan, and Kai-yu Tong. Selective feature aggregation network with area-boundary constraints for polyp segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 302–310. Springer, 2019.
- [9] M. M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A. R. Rudnicka, C. G. Owen, and S. A. Barman. An ensemble classification-based approach applied to retinal blood vessel segmentation. *IEEE Transactions on Biomedical Engineering*, 59(9):2538– 2548, Sep. 2012. ISSN 0018-9294. doi: 10.1109/TBME.2012.2205687.
- [10] Muhammad Moazam Fraz, Paolo Remagnino, Andreas Hoppe, Bunyarit Uyyanonvara, Alicja R Rudnicka, Christopher G Owen, and Sarah A Barman. An ensemble classification-based approach applied to retinal blood vessel segmentation. *IEEE Transactions on Biomedical Engineering*, 59(9):2538–2548, 2012.
- [11] Eli Gibson, Francesco Giganti, Yipeng Hu, Ester Bonmati, Steve Bandula, Kurinchi Gurusamy, Brian Davidson, Stephen P Pereira, Matthew J Clarkson, and Dean C Barratt. Automatic multi-organ segmentation on abdominal ct with dense v-networks. *IEEE transactions on medical imaging*, 37(8):1822–1834, 2018.
- [12] Ran Gu, Guotai Wang, Tao Song, Rui Huang, Michael Aertsen, Jan Deprest, Sébastien Ourselin, Tom Vercauteren, and Shaoting Zhang. Ca-net: Comprehensive attention convolutional neural networks for explainable medical image segmentation. *IEEE transactions on medical imaging*, 40(2):699–711, 2020.
- [13] Zaiwang Gu, Jun Cheng, Huazhu Fu, Kang Zhou, Huaying Hao, Yitian Zhao, Tianyang Zhang, Shenghua Gao, and Jiang Liu. Ce-net: Context encoder network for 2d medical image segmentation. *IEEE transactions on medical imaging*, 38(10):2281–2292, 2019.
- [14] David Gutman, Noel CF Codella, Emre Celebi, Brian Helba, Michael Marchetti, Nabin Mishra, and Allan Halpern. Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (isbi) 2016, hosted by the international skin imaging collaboration (isic). arXiv preprint arXiv:1605.01397, 2016.
- [15] Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger Roth, and Daguang Xu. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. *arXiv preprint arXiv:2201.01266*, 2022.

- [16] 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.
- [17] Jithin Saji Isaac and Ramesh Kulkarni. Super resolution techniques for medical image processing. In 2015 International Conference on Technologies for Sustainable Development (ICTSD). IEEE, 2015.
- [18] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *International Conference on Multimedia Modeling*, pages 451–462. Springer, 2020.
- [19] Qiangguo Jin, Zhaopeng Meng, Tuan D Pham, Qi Chen, Leyi Wei, and Ran Su. Dunet: A deformable network for retinal vessel segmentation. *Knowledge-Based Systems*, 178: 149–162, 2019.
- [20] Tim Laibacher, Tillman Weyde, and Sepehr Jalali. M2u-net: Effective and efficient retinal vessel segmentation for real-world applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [21] 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.
- [22] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [23] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc Van Gool. Deep retinal image understanding. In *International conference on medical image computing and computer-assisted intervention*, pages 140–148. Springer, 2016.
- [24] 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.
- [25] Maria Ines Meyer, Pedro Costa, Adrian Galdran, Ana Maria Mendonça, and Aurélio Campilho. A deep neural network for vessel segmentation of scanning laser ophthalmoscopy images. In *International Conference Image Analysis and Recognition*, pages 507–515. Springer, 2017.
- [26] Jan Odstrcilik, Radim Kolar, Attila Budai, Joachim Hornegger, Jiri Jan, Jiri Gazarek, Tomas Kubena, Pavel Cernosek, Ondrej Svoboda, and Elli Angelopoulou. Retinal vessel segmentation by improved matched filtering: evaluation on a new high-resolution fundus image database. *IET Image Processing*, 7(4):373–383, 2013.
- [27] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.

- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [29] Dzung L Pham, Chenyang Xu, and Jerry L Prince. A survey of current methods in medical image segmentation. *Annual review of biomedical engineering*, 2(3):315–337, 2000.
- [30] 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.
- [31] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv* preprint arXiv:1609.04747, 2016.
- [32] Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu. Transformers in medical imaging: A survey. arXiv preprint arXiv:2201.09873, 2022.
- [33] Neeraj Sharma and Lalit M Aggarwal. Automated medical image segmentation techniques. *Journal of medical physics/Association of Medical Physicists of India*, 35(1): 3, 2010.
- [34] Juan Silva, Aymeric Histace, Olivier Romain, Xavier Dray, and Bertrand Granado. Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer. *International journal of computer assisted radiology and surgery*, 9(2):283– 293, 2014.
- [35] Nima Tajbakhsh, Suryakanth R Gurudu, and Jianming Liang. Automated polyp detection in colonoscopy videos using shape and context information. *IEEE transactions on medical imaging*, 35(2):630–644, 2015.
- [36] Lv Tang, Bo Li, Yijie Zhong, Shouhong Ding, and Mofei Song. Disentangled high quality salient object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3580–3590, 2021.
- [37] 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.
- [38] David Vázquez, Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Antonio M López, Adriana Romero, Michal Drozdzal, and Aaron Courville. A benchmark for endoluminal scene segmentation of colonoscopy images. *Journal of healthcare engineering*, 2017, 2017.
- [39] Jiacheng Wang, Lan Wei, Liansheng Wang, Qichao Zhou, Lei Zhu, and Jing Qin. Boundary-aware transformers for skin lesion segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 206– 216. Springer, 2021.

- [40] Yue Wang, Tülay Adali, Sun-Yuan Kung, and Zsolt Szabo. Quantification and segmentation of brain tissues from mr images: A probabilistic neural network approach. *IEEE transactions on image processing*, 7(8):1165–1181, 1998.
- [41] Jun Wei, Yiwen Hu, Ruimao Zhang, Zhen Li, S Kevin Zhou, and Shuguang Cui. Shallow attention network for polyp segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 699–708. Springer, 2021.
- [42] Chenxi Xie, Changqun Xia, Mingcan Ma, Zhirui Zhao, Xiaowu Chen, and Jia Li. Pyramid grafting network for one-stage high resolution saliency detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11717–11726, 2022.
- [43] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*, pages 1395–1403, 2015.
- [44] Yi Zeng, Pingping Zhang, Jianming Zhang, Zhe Lin, and Huchuan Lu. Towards highresolution salient object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7234–7243, 2019.
- [45] J. Zhang, B. Dashtbozorg, E. Bekkers, J. P. W. Pluim, R. Duits, and B. M. ter Haar Romeny. Robust retinal vessel segmentation via locally adaptive derivative frames in orientation scores. *IEEE Transactions on Medical Imaging*, 35(12):2631–2644, Dec 2016. ISSN 0278-0062.
- [46] Yundong Zhang, Huiye Liu, and Qiang Hu. Transfuse: Fusing transformers and cnns for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 14–24. Springer, 2021.
- [47] 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.