

SA2-Net: Scale-aware Attention Network for Microscopic Image Segmentation

Mustansar Fiaz¹

mustansar.fiaz@mbzuai.ac.ae

Moein Heidari²

moeinheidari7829@gmail.com

Rao Muhammad Anwer¹

rao.anwer@mbzuai.ac.ae

Hisham Cholakkal¹

hisham.cholakkal@mbzuai.ac.ae

¹ Mohamed bin Zayed University of
Artificial Intelligence
Abu Dhabi, UAE

² Iran University of Science and Technol-
ogy, Iran

Abstract

Microscopic image segmentation is a challenging task, wherein the objective is to assign semantic labels to each pixel in a given microscopic image. While convolutional neural networks (CNNs) form the foundation of many existing frameworks, they often struggle to explicitly capture long-range dependencies. Although transformers were initially devised to address this issue using self-attention, it has been proven that both local and global features are crucial for addressing diverse challenges in microscopic images, including variations in shape, size, appearance, and target region density. In this paper, we introduce SA2-Net, an attention-guided method that leverages multi-scale feature learning to effectively handle diverse structures within microscopic images. Specifically, we propose a scale-aware attention (SA2) module designed to capture inherent variations in scales and shapes of microscopic regions, such as cells, for accurate segmentation. This module incorporates local attention at each level of multi-stage features, as well as global attention across multiple resolutions. Furthermore, we address the issue of blurred region boundaries (e.g., cell boundaries) by introducing a novel up-sampling strategy called the Adaptive Up-Attention (AuA) module. This module enhances discriminative ability for improved localization of microscopic regions using an explicit attention mechanism. Extensive experiments on five challenging datasets demonstrate the benefits of our SA2-Net model. Our source code is publicly available at <https://github.com/mustansarfiaz/SA2-Net>.

1 Introduction

Microscopic image segmentation assigns pixel-precise semantic labels to microscopic images and has demonstrated a wide range of applications in biological research and medical diagnosis [1]. For instance, it enables the identification and separation of cell regions or tissues within microscopic images, facilitating a more in-depth analysis of tissue characteristics such as size, shape, texture, and distribution [2]. Such information is crucial for

understanding cellular behavior, disease diagnosis, and drug discovery. However, microscopic image segmentation is challenging due to multiple reasons, such as the presence of noise, occlusions, and overlapping cells, as shown in Figure 1.

Existing microscopic image segmentation approaches are generally based on convolutional neural networks (CNNs) [11, 18]. In particular, the U-Net architecture [22], with its elegant U-shape encoder-decoder structure and the use of multi-scale features is effective in capturing complex image features. As a result, U-Net and its variants have become a cornerstone for various medical image segmentation tasks, including microscopic image segmentation. Despite the impressive performance of CNN-based methods, they often struggle to capture global dependencies due to their limited local receptive fields. The ability to capture long-range global dependencies is crucial for leveraging global information such as illumination levels, scale variations, and occlusions. To tackle this challenge, techniques such as dilated convolution [4, 40] and channel/spatial attention models [25] have shown promise. Nevertheless, these techniques continue to face difficulties in achieving comprehensive global dependencies [4, 6], which can potentially result in a decline in segmentation performance.

In contrast, transformers have showcased exceptional performance in capturing long-range dependencies, both in natural language processing (NLP)[9, 21] and computer vision[11, 33], thanks to their self-attention mechanism. Nonetheless, the hierarchical nature of these transformers can restrict their capacity to learn local contextual information, as global features tend to dominate. Consequently, suboptimal segmentation outcomes may arise. To address this challenge, hybrid CNN-transformer-based methods are introduced to leverage the strengths of both types. For instance, TransFuse [46] employs a parallel architecture that combines both CNN and transformer models, effectively merging multi-level features and resulting in an improved representation of global context. MBT-Net [45] employs a hybrid residual transformer module that harnesses the strengths of both CNNs and transformer blocks, enabling the capture of both local details and global semantics. Additionally, it incorporates a body-edge branch to enhance local consistency. However, directly integrating local representations into self-attention transformers can be challenging, as global features may dominate local ones. This can result in reduced performance, particularly when dealing with small or indistinct objects with blurred boundaries. Consequently, developing a model that can effectively capture both local and global representations simultaneously remains a significantly challenging task.

Despite several attempts to preserve the structure and size during segmentation, accurate localization remains a challenging task. UCTransNet [29] employs multiscale channel-wise




Figure 1: Challenges in microscopic image segmentation across different datasets. *First row:* MoNuSeg dataset [16] sample image (Column 1), ground-truth mask (Column 2), and ground-truth overlay showcasing dense, crowded, and chromatic-sparse nuclei (Column 3). *Second row:* SegPC [43] sample image (Column 1) with corresponding masks and ground-truth overlay showing clustered cells, variations in cell shape, and scale (Columns 2, 3). *Third row:* Additional challenges in SegPC.

fusion with cross-attention, albeit with quadratic computational complexity, and lacks explicit consideration of local features, which are crucial for capturing scale variations, particularly in small regions. In contrast, we propose a novel approach called SA2-Net, a scale-aware attention network designed to effectively manage the diverse and intricate structures of microscopic objects, such as cells. Our primary objective is to incorporate a scale-aware module capable of effectively handling the various scales and shapes of cells or regions present in the input image.

Contributions: The key contributions of this paper are the following: (i) a microscopic image segmentation network known as SA2-Net, consisting of an encoder and a decoder; (ii) in the decoder, the introduction of the scale-aware attention (SA2) module, which effectively captures scale and shape variations through local scale attention at each stage and global scale attention across scales using multi-resolution features; (iii) progressive refinement and upsampling of scale-enhanced features via the adaptive up-attention (AuA) module to generate the final segmentation mask; (iv) comprehensive experiments performed on five challenging datasets, including four microscopic segmentation datasets, demonstrate the merits of the proposed SA2-Net.

2 Related Work

CNN-based Segmentation Networks: Inspired by the success of seminal U-Net [22], it has been incorporated for diverse medical image segmentation tasks. Zhou et al. [69] introduced a nested U-Net (Unet++) which adds a series of nested and skip pathways in the original U-Net, enabling the capability of feature sharing between each sub-network. Numerous attempts have been undertaken, demonstrating encouraging results, including the utilization of dilated convolution [0, 50] and attention models [25]. Moreover, a line of research has strived for the expedient design of skip-connections which is crucial for accurate segmentation as it can assist feature fusion between the expanding and contracting paths [70, 52].

Transformers-based Segmentation Networks: Recently, transformers have shown the ability to extract global information, which is the shortcomings of CNN [24]. Dosovitskiy et al. [53] propose the pioneering vision transformer (ViT), which achieved SOTA performance on image classification tasks by utilizing self-attention mechanisms to retain global information. In the medical domain, Swin-UNet [4] and DS-TransUNet [18] propose models with a U-shaped architecture that relies on the Swin Transformer [19] for 2D image segmentation. TransUNet [6] as a pioneer work in the direction of hybrid models, takes advantage of both CNNs and transformers to grasp both low-level and high-level features.

Multi-Scale Feature Learning: Multi-scale feature representations have lately demonstrated powerful performance in medical image segmentation due to their nature with variations in size and shape. Xu et al. [53] propose MSAU-Net which is a Multi-scale Attention network based on U-Net for nuclei segmentation. CoTr [52] uses deformable attention to fuse flattened multi-scale feature maps from a CNN-based encoder. Gao et al. [12] proposes a multi-scale feature fusion approach that integrates spatial and semantic information globally with a linear-complexity attention mechanism. Inspired by the pyramid structure in CNNs, PVT [11] presents a pyramid vision transformer. Later, Swin Transformer [19] proposed a hierarchical vision transformer that employs shifted window for self-attention. CrossViT [9], suggests a new architecture that utilizes a dual-branch vision transformer in conjunction with a cross-attention module.

Existing methods still struggle to localize the objects due to unclear boundaries between

tissues and surrounding areas. In contrast, we propose SA2 module that strives to capture explicit local information and implicit global structural information without employing self-attention to handle complex shapes and sizes of the microscopic regions such as cells. Furthermore, our A2A learns salient features of the microscopic regions during the upsampling.

3 Methodology

Motivation: We distinguish two important characteristics that need to be taken into account when designing an encoder-decoder framework for microscopic image segmentation. The first characteristic is multi-scale feature representation learning, which is necessary due to the significant variations in size and shape of microscopic cellular regions across different images. To address this challenge, it is important to capture features at multiple scales that can effectively handle irregular cell shapes and sizes. The second characteristic is improved up-sampled feature discriminability for accurate localization. The lack of clear boundaries between cells and their surroundings makes it difficult to accurately segment from heavily cluttered backgrounds. To overcome this issue, an explicit attention mechanism is required during the upsampling and refinement of the features to enhance the discriminative ability for better cell localization. Additionally, these upsampled and refined features should be enriched with high-level semantic information through deep supervision.

3.1 Overall Method

The proposed framework, named SA2-Net, is illustrated in Figure 2, which consists of an encoder that takes an input image (I) and a decoder that predicts the mask (M) for microscopic image segmentation. The encoder generates multi-scale features for four stages with varying resolutions, which are fed into a convolution layer to generate features (F) with 64 channels, striking a balance between performance and efficiency. In the decoder, these multi-resolution features are input to the scale-aware attention (SA2) module (Fig. 2-e), which captures shape and scale variations of cells at each stage as well as jointly across all scales to produce enhanced multi-scale features for better segmentation. The enhanced multi-scale features are upsampled and refined using the adaptive up-attention (AuA) module (Fig. 2-d) to generate the prediction map (M). Additionally, deep supervision is employed at the outputs of the AuA module to accurately capture the varying shapes and sizes of foreground regions such as cells.

3.2 Scale-aware Attention (SA2)

As discussed, the diverse shapes and sizes of cells necessitate the use of multi-scale features to enhance segmentation accuracy. However, we contend that directly combining these features could lead to redundancy and inconsistency, which might result in a sub-optimal segmentation task. Therefore, we present a novel scale-aware attention (SA2) module to adequately capture the size and shape discrepancies of the cells. Figure 2-(e) shows the architecture of our proposed SA2 module. The multi-resolution features from different encoder stages F^i are input to the SA2 module. Within the SA2 module, a local scale attention (LSA) is performed for each stage, and a global scale attention is introduced across all stages.

Local scale attention (LSA): Our LSA captures the scale variations at each feature resolution using different kernels, as shown in 2-(c). It splits the input features F^i and applies

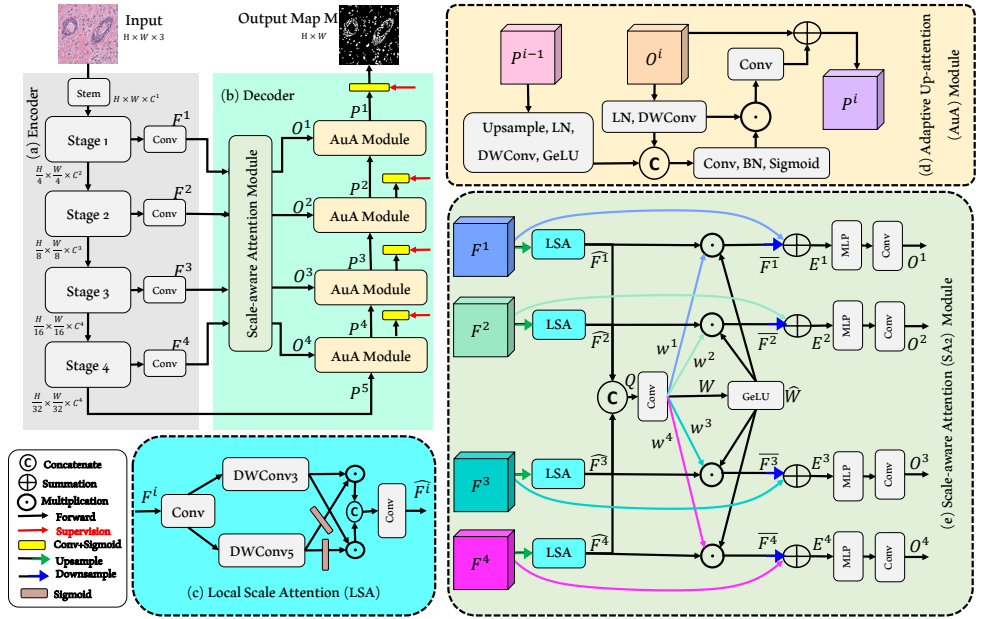


Figure 2: Our proposed microscopic image segmentation network, called SA2-Net, is composed of an encoder (a) and a decoder (b). The SA2-Net takes input and uses an encoder to produce multi-resolution features at various stages. These features with different scales are then fed into a decoder, which generates the final segmentation map (M). The proposed decoder comprises scale-aware attention (SA2) and adaptive up-attention modules. The SA2 module (e) is designed to compute attention both for individual scales and across all scales. The proposed local scale attention (LSA) module (c) first computes attention locally for each scale. These locally attended features are then fused to produce weights for each scale (w^i) and global features (W), enabling joint capture of scale information across all scales. Finally, the adaptive up-attention-(d) module is utilized to progressively refine and upsample the scale-attended features, ultimately leading to the final mask prediction.

an attention mechanism formulated as the element-wise product of two parallel paths of depthwise convolutional layers, one of which is activated with the Sigmoid function. These attended features are fused via concatenation and a convolution layer to generate \hat{F}^i features. We include depthwise convolution layers in LSA to encode information from spatially neighboring pixel positions, to capture the local image structure for effective scale information learning. Furthermore, we intend to enhance these locally attended features to across scales via computing weights for each scale. We argue that each resolution feature assists in better segmentation. To do so, locally attended features \hat{F}^i from all scales are fused and input to a convolution layer to compute single channel weights (w^i) for each scale, and global features. These global features are input to a GeLU to generate global weights (\hat{W}).

$$\begin{aligned}
 Q &= \text{Concat}(\hat{F}^1, \hat{F}^2, \hat{F}^3, \hat{F}^4) \\
 \{w^1, w^2, w^3, w^4\} &= \text{Conv}(Q) \\
 \hat{W} &= \text{GeLU}(\text{Conv}(Q))
 \end{aligned} \tag{1}$$

These global weights \hat{W} are multiplied with \hat{F}^i to enhance relevant features while suppressing irrelevant features based on weights across scales. Furthermore, we perform element-wise product between these weighted features as the following attention formulation:

$$\begin{aligned}\bar{F}^1 &= w^1 \cdot (\hat{F}^1 \cdot \hat{W}), & \bar{F}^2 &= w^2 \cdot (\hat{F}^2 \cdot \hat{W}), \\ \bar{F}^3 &= w^3 \cdot (\hat{F}^3 \cdot \hat{W}), & \bar{F}^4 &= w^4 \cdot (\hat{F}^4 \cdot \hat{W}).\end{aligned}\quad (2)$$

These attended multi-scale features \bar{F}^i are passed to the MLP block after a residual connection, which are then taken as input to a convolution layer to obtain final scale-aware attended features O^i for each scale as shown in Fig. 2-(e). Here, the MLP block has LayerNorm (LN), two convolutional layers, a GeLU, and a residual depthwise convolution to incorporate the positional information as below:

$$F_{out}^i = \text{Conv}(\text{GeLU}(\text{Conv}(\text{DWConv}(\text{LN}(F_{in}^i)) + \text{LN}(F_{in}^i)))) + F_{in}^i. \quad (3)$$

3.3 Adaptive up-attention (AuA)

Generally, in the decoder, the features are upsampled via upsampling and convolution layers which results in suboptimal prediction masks. On the contrary, we propose an adaptive up-attention module (AuA) that upsamples and presents an attention mechanism to enrich each stage feature and strives to learn the salient features of the cells using deep supervision. As shown in Fig. 2-(d), AuA takes input from the previous stage features (P^{i-1}) and current stage scale-aware attended (O^i), which upsamples the previous stage features and strives to capture the salient features for the current stage via Sigmoid function. Furthermore, the outputs from each stage (P^i) are supervised to capture the boundaries of the cells, shown in figure 2-(b). During the inference, the prediction from the last stage P^1 is used to obtain the final segmentation mask M .

4 Experiments

4.1 Datasets

We perform experiments on five medical/microscopic image segmentation benchmark datasets, including MoNuSeg [16], SegPC-2021 [13], GlaS [26], ISIC-2018 [8], and ACDC [6]. The Multiorgan Nucleus Segmentation (MoNuSeg) dataset [16] has been generated utilizing H&E-stained cell images. The training set comprises 30 high-resolution images, which have undergone manual annotation with approximately 22,000 nuclear boundary annotations. In comparison, the test set contains 14 images labelled with over 7000 nuclear boundary annotations. Our methodology is also evaluated on the SegPC2021 dataset [13] for the segmentation of multiple myeloma

Method	SegPC21		ISIC2018	
	Dice (%)	IoU (%)	Dice (%)	IoU (%)
U-Net [16]	88.08	88.2	86.71	84.91
UNet++ [16]	91.02	90.92	88.22	86.51
AttUNet [26]	91.58	91.44	88.20	86.49
MultiResUNet [16]	86.49	86.76	86.94	85.37
TransUNet [9]	82.33	83.38	84.99	83.65
ResidualUNet [16]	84.79	85.41	86.89	85.09
MissFormer[16]	80.82	82.09	86.57	84.84
UCTransNet [16]	91.74	91.59	88.98	87.29
SA2-Net (Ours)	92.41	92.23	88.88	87.21

Table 1: Comparison of our method over SegPC21 and ISIC2018 dataset with SOTA methods using Dice and IoU metrics. The best results are in bold.

Method	GlaS		MoNuSeg	
	Dice (%)	IoU (%)	Dice (%)	IoU (%)
U-Net [14]	85.45 ± 1.3	74.78 ± 1.7	76.45 ± 2.6	62.86 ± 3.0
UNet++ [19]	87.56 ± 1.2	79.13 ± 1.7	77.01 ± 2.1	63.04 ± 2.5
AttUNet [14]	88.80 ± 1.1	80.69 ± 1.7	76.67 ± 1.1	63.47 ± 1.2
MultiResUNet [14]	88.73 ± 1.2	80.89 ± 1.7	78.22 ± 2.5	64.83 ± 2.9
TransUNet [9]	88.40 ± 0.7	80.40 ± 1.0	78.53 ± 1.1	65.05 ± 1.3
MedT [14]	85.93 ± 2.9	75.47 ± 3.5	77.46 ± 2.4	63.37 ± 3.1
Swin-Unet [9]	89.58 ± 0.6	82.07 ± 0.7	77.69 ± 0.9	63.77 ± 1.2
UCTransNet [14]	90.18 ± 0.7	82.96 ± 1.1	79.08 ± 0.7	65.50 ± 0.9
SA2-Net (Ours)	91.38 ± 0.4	84.90 ± 0.6	81.34 ± 0.5	68.70 ± 0.7

Table 2: Comparison of our method over GlaS and MoNuSeg dataset with SOTA methods using Dice and IoU metrics. The best results are in bold.

cells. The dataset consists of a training set comprising 290 samples, and two other sets used for validation and testing containing 200 and 277 samples respectively. To evaluate our method on this dataset, we follow the same strategy presented in [9]. The Gland Segmentation in Colon Histology Images (GlaS) [26] dataset includes high-resolution microscopic images of Hematoxylin and Eosin (H&E) stained slides and corresponding ground truth annotations which comprise a total of 165 images, including 85 images for training and 80 images for testing purposes. Moreover, we used the ISIC-2018 dataset [8], collected by the International Skin Imaging Collaboration (ISIC) as a large-scale dataset of 2594 dermoscopy images along with their corresponding ground truth annotations and followed [9] in splitting and pre-processing the data. Lastly, we used the ACDC dataset containing 100 MRI scans, where multiple patients were scanned, and each scan was labelled for three different organs which are the left ventricle (LV), right ventricle (RV), and myocardium (MYO). The training, validation, and testing process follow literature work [9], where 70 cases are used for training, 10 for validation, and 20 for testing purposes.

4.2 Implementation Details

Our model was developed using PyTorch on a single NVIDIA RTX A6000 GPU card with 48 GB memory. The input resolution was set to 224×224 , and we opted not to use pre-trained weights for training. During training, random flipping and rotation were employed for data augmentation, and the combined weighted IoU and weighted BCE loss function was used. Similar to UCTransNet [14], for GlaS and MoNuSeg datasets, a batch size of 4 was set, and the Adam optimizer was used with an initial learning rate of 0.001. To ensure the reliability of the results on small datasets, we conducted three times 5-fold cross-validation, totaling 15 cross-validations. For the GlaS and MoNuSeg datasets, an ensemble technique was employed by obtaining predictions from all five models and taking the mean of these predictions as the final prediction. Similar to [9], for SegPC21 and ISIC2018, a batch size of 16 was set, and training was done for 100 epochs using the Adam optimizer with an initial learning rate of 0.0001. For ACDC, a batch size of 12 was set, and training was done for 150 epochs using the Adam optimizer with an initial learning rate of 0.0001. Sigmoid activation was applied for the binary segmentation mask and Softmax activation for multi-class segmentation. Our method was evaluated using dice and IoU metrics.

Method	Dice Avg	RV	Myo	LV
U-Net [14]	87.55	87.10	80.63	94.92
R50+AttnUNet [9]	86.75	87.58	79.20	93.47
Vit+CUP [9]	81.45	81.46	70.71	92.18
R50+ViT+CUP [9]	87.57	86.07	81.88	94.75
TransUNet [9]	89.71	86.67	87.27	95.18
SwinUnet [9]	90.00	88.55	85.62	95.83
SA2Net (Ours)	92.30	90.33	90.39	96.18

Table 3: Comparison of our approach over ACDC dataset reported an average Dice score metric and Dice score for each organ. The best results are in bold.




Figure 3: Qualitative results on Glas and MoNuSeg datasets. The red box in the top row highlights regions where exactly SA2-Net performs better than the other methods, while in the bottom row green indicates true positive and red shows prediction errors.




Figure 4: Qualitative comparison of the proposed method on the *SegPC-2021* dataset for segmentation of Multiple Myeloma plasma cells. For a fair comparison, we follow [14] and perform the cytoplasm segmentation (red mask) for a given input nucleus mask (blue).




Figure 5: Feature visualization of our the AuA modules output (P^4 to P^1) using Grad-CAM [23] over GlaS (first row) and MoNuSeg (second row) images. The results show that with the help of the SA2 module, AuA modules, and deep supervision, features that are robust to the shape, size, and density variations of cells are learned at the decoder.

4.3 Quantitative Results

The comparison of the proposal with previous SOTA methods is presented in Table 2, 1, and 3 respectively. Table 2 shows that the proposed SA2-Net has significant improvements over prior arts for GlaS and MoNuSeg datasets, e.g., our method gains a dice score of 1.20% and 2.67% over best-performing UTransNet [29]. A similar trend can be also witnessed using IoU metric, where our method outperforms SOTA over GlaS and MoNuSeg by a large margin and achieves 84.90 and 68.70 scores respectively. Table 1 presents the results for SegPC and ISIC datasets, which highlights that our method significantly outperforms other methods for SegPC and obtains comparable results compared to UTransNet [29] over ISIC2018 dataset. Also, our method achieves SOTA over the ACDC dataset shown in table 3.

4.4 Qualitative Results

We also provide qualitative comparison results on three datasets to further prove the generalization ability of our SA2-Net as shown in Figures 3 and 4. Looking at Figure 3, our predictions on Glas and MoNuSeg datasets adjust well to the provided GT masks indicat-

ing that our method has the capability to create multi-scale representation and can filter out background noise which is common in datasets that have overlapping backgrounds. Furthermore, the SegPC dataset’s image outputs (Figure 4) illustrate that SA2-Net has the ability to grasp intricate details, create precise contours, and exhibit strong performance in areas where boundaries can be unclear due to interference, showing robustness to noisy items adapting to the given GT masks. This can be attributed to the expedient combination of attention-based modules and CNN for modeling global connections and local representations.

4.5 Ablation Study

We performed thorough ablation study further to validate the performance of each module in SA2-Net under different settings. Specifically, we investigate the influence of each contribution and the results are summarized in Table 4.

The combination of SA2 and LSA in SA2-Net shows that the two attention mechanisms are complementary and can collaborate to provide better segmentation predictions. However, an attention-based signal for our decoder using deep supervision produces the best performance in all metrics as it captures local contextual relations among pixels. Lastly, we present the feature visualization using Grad-CAM [23] of each level of our AuA module (Figure 5). As illustrated, after applying the attention mechanism at each level, more emphasis is drawn to the desired region and is more highlighted than the surroundings, demonstrating that the SA2 module, AuA modules, and deep supervision learn features that are robust to the shape, size, and density variations of cells. Furthermore, each level serves a complimentary function, with the larger level providing fine-grained features and the smaller level attempting to give extra information. As a result, all levels are required for the model to function effectively.

Model	SA2	LSA	AuA	Deep Supervision	Dice	IoU
Baseline (UNet)	✗	✗	✗	✗	76.55	62.76
SA2-Net	✓	✗	✗	✗	80.25	68.96
SA2-Net	✗	✓	✗	✗	79.98	68.84
SA2-Net	✓	✓	✗	✗	81.53	69.12
SA2-Net	✓	✓	✓	✗	81.65	69.20
SA2-Net	✓	✓	✓	✓	81.75	69.25

Table 4: Ablation experiments on MoNuSeg dataset: impact of individual contributions on segmentation performance of SA2-Net.

5 Conclusion

In this paper, we present SA2-Net, a novel framework for microscopic image segmentation. The proposed approach effectively captures variations in the size and shape of microscopic regions such as cells through multi-scale feature learning. Furthermore, we incorporate local scale attention at each stage and global scale attention across scales to address the challenges posed by diverse cell structures. Overall, the proposed framework leverages multi-scale features, scale-aware attention, adaptive up-attention, and deep supervision for accurate microscopic image segmentation. Experiments across five diverse datasets show the superiority of the proposed method.

Acknowledgement

This work is partially supported by the MBZUAI-WIS Joint Program for AI Research (Project grant number- WIS P008).

References

- [1] Shumoos Al-Fahdawi, Rami Qahwaji, Alaa S Al-Waisy, Stanley Ipson, Maryam Ferdousi, Rayaz A Malik, and Arun Brahma. A fully automated cell segmentation and morphometric parameter system for quantifying corneal endothelial cell morphology. *Computer methods and programs in biomedicine*, 160:11–23, 2018.
- [2] Maryam Asadi-Aghbolaghi, Reza Azad, Mahmood Fathy, and Sergio Escalera. Multi-level context gating of embedded collective knowledge for medical image segmentation. *arXiv preprint arXiv:2003.05056*, 2020.
- [3] Reza Azad, Ehsan Khodapanah Aghdam, Amelie Rauland, Yiwei Jia, Atlas Haddadi Avval, Afshin Bozorgpour, Sanaz Karimijafarbigloo, Joseph Paul Cohen, Ehsan Adeli, and Dorit Merhof. Medical image segmentation review: The success of u-net. *arXiv preprint arXiv:2211.14830*, 2022.
- [4] 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. In *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*, pages 205–218. Springer, 2023.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kaloo, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [10] 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.
- [11] Huazhu Fu, Yanwu Xu, Stephen Lin, Damon Wing Kee Wong, Baskaran Mani, Meenakshi Mahesh, Tin Aung, and Jiang Liu. Multi-context deep network for angle-closure glaucoma screening in anterior segment oct. In *Medical Image Computing*

- and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11*, pages 356–363. Springer, 2018.
- [12] Yunhe Gao, Mu Zhou, Di Liu, Zhennan Yan, Shaoting Zhang, and Dimitris N Metaxas. A data-scalable transformer for medical image segmentation: architecture, model efficiency, and benchmark. *arXiv preprint arXiv:2203.00131*, 2022.
- [13] Anubha Gupta, Ritu Gupta, Shiv Gehlot, and Shubham Goswami. Segpc-2021: Segmentation of multiple myeloma plasma cells in microscopic images. *IEEE Dataport*, 1(1):1, 2021.
- [14] Xiaohong Huang, Zhifang Deng, Dandan Li, and Xueguang Yuan. Missformer: An effective medical image segmentation transformer. *arXiv preprint arXiv:2109.07162*, 2021.
- [15] Nabil Ibtehaz and M Sohel Rahman. Multiresunet: Rethinking the u-net architecture for multimodal biomedical image segmentation. *Neural networks*, 121:74–87, 2020.
- [16] Neeraj Kumar, Ruchika Verma, Sanuj Sharma, Surabhi Bhargava, Abhishek Vahadane, and Amit Sethi. A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE transactions on medical imaging*, 36(7):1550–1560, 2017.
- [17] Markus Leygeber, Dorina Lindemann, Christian Carsten Sachs, Eugen Kaganovitch, Wolfgang Wiechert, Katharina Nöh, and Dietrich Kohlheyer. Analyzing microbial population heterogeneity—expanding the toolbox of microfluidic single-cell cultivations. *Journal of molecular biology*, 431(23):4569–4588, 2019.
- [18] 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*, 71:1–15, 2022.
- [19] 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.
- [20] 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.
- [21] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.

- [23] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [24] Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu. Transformers in medical imaging: A survey. *Medical Image Analysis*, page 102802, 2023.
- [25] Ashish Sinha and Jose Dolz. Multi-scale self-guided attention for medical image segmentation. *IEEE journal of biomedical and health informatics*, 25(1):121–130, 2020.
- [26] Korsuk Sirinukunwattana, Josien PW Pluim, Hao Chen, Xiaojuan Qi, Pheng-Ann Heng, Yun Bo Guo, Li Yang Wang, Bogdan J Matuszewski, Elia Bruni, Urko Sanchez, et al. Gland segmentation in colon histology images: The glas challenge contest. *Medical image analysis*, 35:489–502, 2017.
- [27] Jeya Maria Jose Valanarasu, Poojan Oza, Ilker Hacihaliloglu, and Vishal M Patel. Medical transformer: Gated axial-attention for medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, pages 36–46. Springer, 2021.
- [28] Juan P Viguera-Guillén, Busra Sari, Stanley F Goes, Hans G Lemij, Jeroen van Rooij, Koenraad A Vermeer, and Lucas J van Vliet. Fully convolutional architecture vs sliding-window cnn for corneal endothelium cell segmentation. *BMC Biomedical Engineering*, 1(1):1–16, 2019.
- [29] Haonan Wang, Peng Cao, Jiaqi Wang, and Osmar R Zaiane. Uctransnet: rethinking the skip connections in u-net from a channel-wise perspective with transformer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 2441–2449, 2022.
- [30] Shuhang Wang, Szu-Yeu Hu, Eugene Cheah, Xiaohong Wang, Jingchao Wang, Lei Chen, Masoud Baikpour, Arinc Ozturk, Qian Li, Shinn-Huey Chou, et al. U-net using stacked dilated convolutions for medical image segmentation. *arXiv preprint arXiv:2004.03466*, 2020.
- [31] 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.
- [32] Yutong Xie, Jianpeng Zhang, Chunhua Shen, and Yong Xia. Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24*, pages 171–180. Springer, 2021.
- [33] Qing Xu and Wenting Duan. An automatic nuclei image segmentation based on multi-scale split-attention u-net. In *MICCAI Workshop on Computational Pathology*, pages 236–245. PMLR, 2021.

- [34] Qing Xu, Zhicheng Ma, HE Na, and Wenting Duan. Dcsau-net: A deeper and more compact split-attention u-net for medical image segmentation. *Computers in Biology and Medicine*, 154:106626, 2023.
- [35] Yinglin Zhang, Risa Higashita, Huazhu Fu, Yanwu Xu, Yang Zhang, Haofeng Liu, Jian Zhang, and Jiang Liu. A multi-branch hybrid transformer network for corneal endothelial cell segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, pages 99–108. Springer, 2021.
- [36] Yundong Zhang, Huiye Liu, and Qiang Hu. Transfuse: Fusing transformers and cnns for medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, pages 14–24. Springer, 2021.
- [37] Zhengxin Zhang, Qingjie Liu, and Yunhong Wang. Road extraction by deep residual u-net. *IEEE Geoscience and Remote Sensing Letters*, 15(5):749–753, 2018.
- [38] Zhuangzhuang Zhang and Weixiong Zhang. Pyramid medical transformer for medical image segmentation. *arXiv preprint arXiv:2104.14702*, 2021.
- [39] 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: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018.