

Liu, Q., Kaul, C., Wang, J., Anagnostopoulos, C., Murray-Smith, R. and Deligianni, F. (2023) Optimizing Vision Transformers for Medical Image Segmentation. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2023), Rhodes, Greece, 4-10 June 2023, ISBN 9781728163277.



Copyright © 2023 IEEE. Reproduced under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

For the purpose of open access, the author(s) has applied a Creative Commons Attribution license to any Accepted Manuscript version arising.

<https://eprints.gla.ac.uk/292461/>

Deposited on: 20 February 2023

OPTIMIZING VISION TRANSFORMERS FOR MEDICAL IMAGE SEGMENTATION

Qianying Liu¹, Chaitanya Kaul¹, Jun Wang², Christos Anagnostopoulos¹,
Roderick Murray-Smith¹, Fani Deligianni¹

¹School of Computing Science, University of Glasgow ²University of Warwick

ABSTRACT

For medical image semantic segmentation (MISS), Vision Transformers have emerged as strong alternatives to convolutional neural networks thanks to their inherent ability to capture long-range correlations. However, existing research uses off-the-shelf vision Transformer blocks based on linear projections and feature processing which lack spatial and local context to refine organ boundaries. Furthermore, Transformers do not generalize well on small medical imaging datasets and rely on large-scale pre-training due to limited inductive biases. To address these problems, we demonstrate the design of a compact and accurate Transformer network for MISS, CS-Unet, which introduces convolutions in a multi-stage design for hierarchically enhancing spatial and local modeling ability of Transformers. This is mainly achieved by our well-designed Convolutional Swin Transformer (CST) block which merges convolutions with Multi-Head Self-Attention and Feed-Forward Networks for providing inherent localized spatial context and inductive biases. Experiments demonstrate CS-Unet without pre-training outperforms other counterparts by large margins on multi-organ and cardiac datasets with fewer parameters and achieves state-of-the-art performance. Our code is available at Github¹.

Index Terms— Medical Image Segmentation, Semantic Segmentation, Vision Transformer, Convolutions

1. INTRODUCTION

Medical image semantic segmentation (MISS), which classifies image pixels with semantic organ labels (e.g. Kidney and Liver) for various imaging modalities, is considered as one of the most fundamental problems in medical imaging. However, compared to natural scene images, MISS requires overcoming more challenges to create robust models. For instance, common benchmark datasets in MISS suffer from large deformation of organs under different image acquisition processes. In addition, shortage of costly pixel-level annotations is another problem leading to a performance gap. To achieve efficient and effective segmentation, models are not only required to have a better understanding of their local semantic features to capture more subtle organ structures, but also of global feature dependencies to capture the relationships among multiple organs.

UNet [1] and its variants [2][3][4][5][6] with Convolutional Neural Networks (CNNs) as the backbone have found huge success in MISS as they are good at modelling local attributes inside their receptive field. However, the inherent locality of convolution operations restricts their ability to model long-range semantic dependencies within the image, and as a result the challenging

boundaries of the whole organ may not be effectively segmented. Attention mechanisms alleviate this issue, but these tend to be 'single head' mechanisms that only calculate pixel-level similarities, and not 'multi head' with the ability to capture patch-level patterns.

For alleviating the inherent flaws of CNNs, there's a recent shift in the choice of architectures from CNNs to Vision Transformers (ViTs) due to their ability to model long range semantic attributes among input tokens (embeddings of image patches) via a linearly projected Multi-Head Self-Attention (MHSA) operation and a Feed-Forward Network (FFN). Most early works [7][8][9] treat CNNs as a backbone and exploit the Transformer's desirable characteristics in their encoder. They tend to have high complexity as they stack bulky Transformer blocks on top of convolutional feature extractors (large pretrained CNNs, e.g. ResNet). Recent research [10][11][12][13][14][15][16][17] has moved towards using Transformers as the main stem for building the entire segmentation architecture. Swin-UNet [12] is regarded as the first pure Transformer model. It keeps the familiar U-shape and adds hierarchical feature extraction using shifted windows proposed by the Swin Transformer [18]. This drastically reduces the quadratic complexity of traditional self-attention while achieving better performance.

However, most of these Transformers for MISS use off-the-shelf Transformer blocks from Computer Vision community and only model and extract linear semantic relations via MHSA and FFN, leading to the challenge of precisely delineating organ boundaries due to the lack of spatial and local information as shown in Figure 1.(d), although showing small influence on detection and classification tasks. Besides, these methods require a large dataset to compensate the lack of inductive biases such as translation equivariance [19], which may be defected or even lost when fine-tuning on downstream tasks, showing less robustness on small datasets.

Keeping the current state of the literature in mind, our paper highlights issues that today's Transformers in MISS face, followed by our contribution that helps alleviate those drawbacks. Most current Transformers are bulky and rely on pre-training weights from classical vision tasks to be adapted for MISS. To the best of our knowledge, no existing study explores the effects of adding spatial locality inside Transformer blocks via convolutions for medical imaging. To this end, we first propose an empirical analysis to show the need for spatial locality in pure Transformer based MISS. Our insights show the effects of introducing convolutions to Transformer blocks and multi-stage design of networks on segmentation performance. We call the final model resulting from our experiments, Convolutional Swin-Unet (CS-Unet), which is based on purely convolutional Transformer blocks created to make Transformers model local information better, segment organ boundaries more accurately, while maintaining a low computational complexity. Experiments on CT and MRI datasets show CS-Unet (24M parameters) trained from scratch outperforms pre-trained Swin-Unet (27M) on ImageNet by around 3% dice score, achieving state-of-the-art performance.

Thanks to EPSRC (EP/W01212X/1), Royal Society (RGS/R2/212199), UKRI 104690 (iCAIRD) and China Scholarship Councils for funding.

¹ <https://github.com/kathyliu579/CS-Unet>

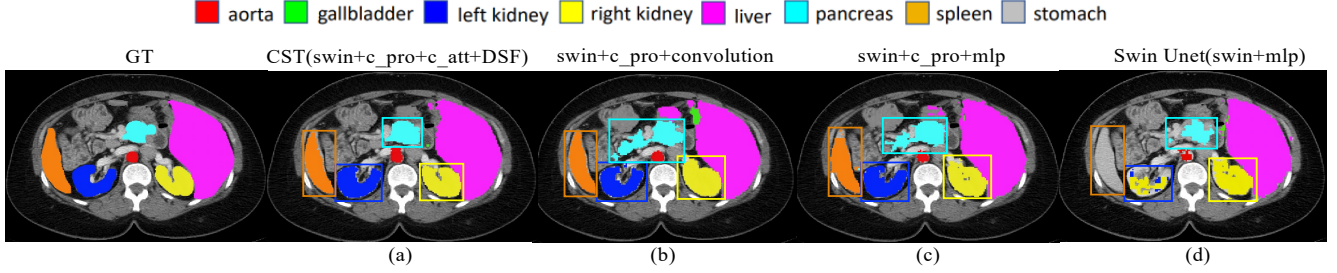


Fig. 1. Visualization of segmentation results of different methods trained from scratch on Synapse dataset.

2. METHOD

Most Transformer based methods in MISS, i.e., encoder-decoder models with a standard U-shape, use a standard Transformer block containing linear projections and linear FFNs, which are essentially MLPs, to process the data. Hence, to create effective image representations using such a regime requires huge amounts of data for training, as they lack local spatial information.

The first pure-Transformer based MISS model is the Swin-Unet [12] which adopts Swin Transformer blocks [18] to add locality information to Transformers. The data representation created here is still inherently linear as this block contains linear projections and feature processing. Figure 1 shows segmentation visualizations for the Synapse dataset. Swin-Unet trained from a random weight initialization (Figure 1.(d)) does not perform well. It fails to detect the spleen and misclassifies the left kidney as the right.

Next, we add convolutional projections to this Swin Transformer block structure. The projections follow the methodology proposed in [20] where tokens are first shaped into a 2D token map, then processed by a depth-wise separable convolution with kernel size s implemented by: Depth-wise Conv \rightarrow BatchNorm \rightarrow Point-wise Conv. Finally, the tokens are flattened into 1D token input $x_i^{q/k/v}$ for Q/K/V matrices. It can be formulated as: $x_i^{q/k/v} = Flatten(Conv(Reshape2D(x_i), s))$. Figure 1.(c) shows outputs of the resultant Unet trained with this block. It visually demonstrates how spatial locality is essential for low level pixel labelling tasks. It can be seen that although the convolutional projection alleviates a lot of the problems posed by the linearity of Swin-Unet, there are still severe over-segmentations on pancreas and liver and extremely rough boundaries of right kidney.

Following this, when a 3×3 convolution is used for FFNs instead of MLPs to introduce more spatial context, we see the full effects of adding complete spatial locality to Transformers through the boundaries of the left and right kidneys and spleen becoming greatly refined. The over-segmentation problem of the pancreas however gets worse (as shown in Figure 1.(b)). This is due to the limited receptive field not modeling the whole boundary of big organs effectively.

2.1. Convolutional Swin Transformer (CST) Layer

We propose a CST layer to fully explore spatial modeling ability of convolutions in MHSA and FFN. First, we propose a novel (shifted) window based convolutional multi-head self attention ((S)W-CMSA) to extract hierarchical semantic features while reducing computational costs, by combining a shifted windows mechanism and convolutional projection. Then, we replace the MLP with our novel depthwise separable feed-forward (DSF) module. From Figure 1.(a), we see the Transformer model based on CST handles challenging organ boundaries more efficiently. The CST layer is formulated as:

$$\hat{z}^l = W - CMSA(LN(z^{l-1})) + z^{l-1}, \quad (1)$$

$$z^l = DSF(\hat{z}^l) + \hat{z}^l, \quad (2)$$

$$\hat{z}^{l+1} = SW - CMSA(LN(z^l)) + z^l, \quad (3)$$

$$z^{l+1} = DSF(\hat{z}^{l+1}) + \hat{z}^{l+1} \quad (4)$$

where \hat{z}^l and z^l denote the outputs of (S)W-CMSA module and DSF of the l -th block, respectively.

(Shifted) Window based convolutional multi-head self attention




Fig. 2. Convolutional Swin Transformer (CST) Block.

As shown in Figure 2, once tokens enter (S)W-CMSA, they are reshaped into a 2D token map, and partitioned into windows. For each window, we use three depth-wise convolutions with kernel size s of 3×3 , padding of 1 and stride of 1 to create our Q, K and V vectors via: $Flatten(DepthConv(Window(Reshape(x_i)), s))$.

CST is different from [20] as we create a projection based on windows rather than the whole image, leading to more refined local features as now the kernels learnt on each window are different. In order to better adapt to medical images with smaller data volumes, point-wise convolutions are removed to avoid over-fitting. Furthermore, we replace Batch Normalization with Layer Normalization (LN), providing a performance boost. The token vectors are fed to MHSA as:

$$MHSA(x_i^q, x_i^k, x_i^v) = SoftMax\left(\frac{x_i^q(x_i^k)^T}{\sqrt{d}} + B\right)x_i^v \quad (5)$$

Here d represents the dimension of the query and key. The values in B are the bias.

Then, we replace the linear layer and feed the attention output to a 3×3 depth-wise convolution for fine-tuning for more spatial information. We follow this by reversing the windows to 2D token maps, resulting in more robust estimations compared to Swin Transformer [18] removing our dependence on positional encoding.

Depthwise separable feed-forward (DSF) module After computing (S)W-CMSA, the feature maps are fed into a FFN. Existing Transformers implement this module as an MLP: $\text{LN}, d \rightarrow \text{Linear}, 4 \times d \rightarrow \text{GELU} \rightarrow \text{Linear}, d \rightarrow \text{RC}$. The d denotes the number of channels of a reshaped feature map and RC denotes the residual connection. We propose a DSF module as a choice of FFN which provides adding spatial context. We use three depth-wise convolutions instead of two linear layers for utilizing the features between channels, C . In addition, we found that adding LN after convolution gives better segmentation results. The DSF is implemented as: 7×7 Depth-wise Conv, $d \rightarrow \text{LN}, d \rightarrow \text{Point-wise Conv}, 4 \times d \rightarrow \text{GELU} \rightarrow \text{Point-wise Conv}, d \rightarrow \text{RC}$.

2.2. Overall Structure Design

CS-Unet keeps a symmetrical UNet shape. The input of our model is a 2D image slice with the resolution of $H \times W \times 3$ sampled from a 3D volume of images. H , W and 3 denote the height, width and number of channels of each input. The input images on entering the encoder are passed through the convolutional token embedding to create a sequence embedding on overlapping patches of the image, following which CST and patch merging layers are applied. Extracted features are then processed by the model’s bottleneck that consists of two CST blocks. A symmetrical decoder then creates the final segmentation marks. In addition, skip convolution (SC) modules are added between corresponding feature pyramids of the encoder and decoder to compensate for the missing information caused by down-sampling. The overall architecture of the proposed CS-Unet is presented in Figure 3.

2.3. Encoder

The input image is first passed through the convolutional token embedding layer to create a sequence embedding with the resolution of $\frac{H}{4} \times \frac{W}{4} \times C$ ($C = 96$ in experiments). This embedding is fed to three main CST layers and a patch merging module which downsamples the image and doubles the number of channels. For example, at the first patch merging module, an input with size $\frac{H}{4} \times \frac{W}{4} \times C$ is divided into four parts and concatenated along the C dimension to create a feature map of size $\frac{H}{8} \times \frac{W}{8} \times 4C$. Then a linear layer is applied to this map to reduce the C dimension by a factor of 2.

Convolutional Token Embedding layer Existing models use a linear layer to split images into non-overlapping patches and reduce the size of the image drastically, e.g. by 75%, while increasing the channel dimension C . However, as the images’ highest resolution is $H \times W$ at the encoder, using a linear layer to compress these features not only loses high-quality spatial and local information, but also increases model size. Our embedding layer, is implemented as four convolutions with overlapping patches to compress features in stages, helping to introduce more spatial dependency between, and inside the patches, while greatly reducing the parameters (by 6M. See Ablation 3, Method 1). Specifically, this layer is implemented as follows: 3×3 $s=1$ Conv, $d/2 \rightarrow \text{GELU} \rightarrow 3 \times 3$ $s=2$ Conv, $d/2 \rightarrow \text{GELU}+\text{LN} \rightarrow 3 \times 3$ $s=1$ Conv, $d \rightarrow \text{GELU} \rightarrow 3 \times 3$ $s=2$ Conv, $d \rightarrow \text{GELU}$. Here, s is stride, the input dimension is 3, and $d = C$. In the end, 2D reshaped token maps with resolution $\frac{H}{4} \times \frac{W}{4} \times C$ are outputted.

Bottleneck The bottleneck contains two CST blocks, based on W-CMSA. The feature map size here remains unchanged.

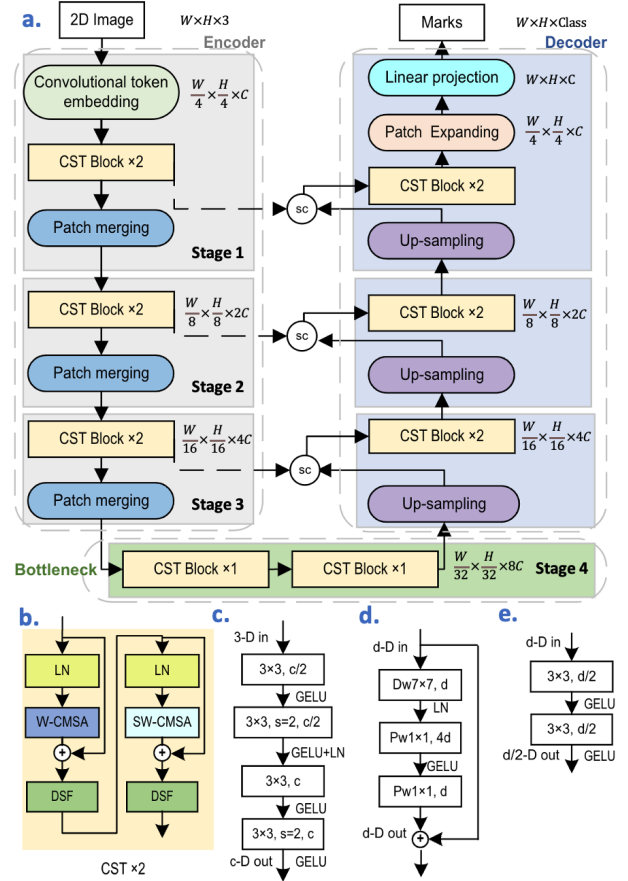


Fig. 3. (a) Overall architecture of CS-Unet, (b) one CST layer, (c) convolutional token embedding, (d) DSF and (e) skip convolutions. d is the current number of channels, c is an arbitrary dimension.

2.4. Decoder

Our decoder is symmetric to the encoder. Feature representation is created by enlarging the feature volume through a convolutional up-sampling module and then passing it through a SC module to compensate for the information lost due to down-sampling. A CST layer then provides spatial context to the upsampled features. After repeating the above process three times, the features are fed into the patch expansion layer which up-samples by $4 \times$, followed by a linear projection to fine tune the final segmentation prediction. Specifically, convolutional up-sampling module employs strided deconvolution to $2 \times$ up-sample feature maps and halves the channel dimension as: $\text{LN}, d \rightarrow 2 \times 2$ $s=2$ ConvTranspose, $d/2 \rightarrow \text{GELU}$.

Skip Convolutions (SC) module The outputs of high-resolution feature maps created from up-sampling are concatenated with shallow feature representations from the encoder, and then merged by a SC module. It further enriches both spatial and fine-grained information, while compensating for the missing information caused by down-sampling. It is implemented as: 3×3 $s=1$ Conv, $d/2 \rightarrow \text{GELU} \rightarrow 3 \times 3$ $s=1$ Conv, $d/2 \rightarrow \text{GELU}$.

3. EXPERIMENTS

We use two publicly available datasets to benchmark our method.

Synapse multiorgan segmentation (Synapse): This dataset [21] contains abdominal CT scans from 30 subjects. Following [7],

Methods	DSC	HD	Aorta	Gallb	Kid.L	Kid.R	Liver	Pancr	Spleen	Stom
R50 UNet [7]	74.68	36.87	84.18	62.84	79.19	71.29	93.35	48.23	84.41	73.92
R50 AttnUNet [7]	75.57	36.97	55.92	63.91	79.20	72.71	93.56	49.37	87.19	74.95
UNet [1]	76.85	39.70	<u>89.07</u>	<u>69.72</u>	77.77	68.60	93.43	53.98	86.67	75.58
AttnUNet [4]	77.77	36.02	89.55	68.88	77.98	71.11	93.57	58.04	87.30	75.75
R50 ViT [7]	71.29	32.87	73.73	55.13	75.80	72.20	91.51	45.99	81.99	73.95
TransUnet [7]	77.48	31.69	87.23	63.13	81.87	77.02	94.08	55.86	85.08	75.62
Swin-Unet [12]	<u>79.13</u>	21.55	85.47	66.53	<u>83.28</u>	79.61	<u>94.29</u>	56.58	<u>90.66</u>	<u>76.60</u>
MT-UNet [16]	78.59	<u>26.59</u>	87.92	64.99	81.47	77.29	93.06	<u>59.46</u>	87.75	76.81
Ours	82.21	27.02	88.40	72.59	85.28	<u>79.52</u>	94.35	70.12	91.06	75.72

Table 1. Comparison with different models on Synapse. Gallbladder, left Kidney, right Kidney, Pancreas and Stomach are abbreviated as Gallb, Kid.L, Kid.R, Pancr and Stom.

18 cases (2212 axial slices) are extracted for training, while other 12 cases are used for testing. We report the model performance evaluated with the average Dice score Coefficient (DSC) and average Hausdorff Distance (HD) on eight abdominal organs.

Automatic Cardiac Diagnosis Challenge (ACDC): ACDC [22] contains MRI images from 100 patients, with right ventricle (RV), left ventricle (LV) and myocardium (MYO) to be segmented. Using data splits proposed in [16], the dataset is split into 70 (1930 axial slices), 10 and 20 for training, validation and testing, respectively. Evaluation metrics used are average DSC (%) and HD (mm).

3.1. Implementation details

We train our models on a single Nvidia RTX3090 GPU with 24GB memory. We use flipping and rotation augmentations on the training data. The input image size is 224x224. Pre-trained weights are used for other methods if provided, while our model is trained from scratch for 300 epochs from a randomly initialized set of weights. A batch size of 24 and a combination of cross entropy and dice loss are used. Our model is optimized by AdamW [23] with a weight decay of 5E-4 for both datasets. The learning rates for Synapse and ACDC are 1e-3 and 5e-3, respectively. We start with a 10-epoch linear warmup. Layer Scale [24] of initial value 1e-6 is applied.

3.2. Experimental Results

As shown in Table 1 and Table 2, our model consistently surpasses a variety of convolution-based and Transformer-based methods. CS-Unet outperforms Swin-Unet by 3.08% and 3.3% DSC on Synapse and ACDC, respectively. In addition, our method gets the highest DSC for five and two organs of Synapse and ACDC respectively, especially providing large boosts for challenging organs like gallbladder, pancreas and RV. Overall, compared to pretrained Swin-Unet (27 M), nnFormer(158 M) and TransUnet (96 M), CS-Unet achieves the best performance without pretraining while being lightweight (24 M) via introducing more local perception and inductive bias.

Figure 4 visualizes segmentation results. In case 1, our method has overwhelming advantage on segmenting the pancreas, stomach and liver. CS-Unet is also more discriminative on the complex shape of RV than other Transformer-based models in case 2 due to its better ability of spatial context modelling.

3.3. Ablation study

We explore the influence of proposed modules on the performance on Synapse as shown in Table 3. The Swin-Unet trained from scratch is treated as the baseline (method 0) which cannot adapt to small datasets. Adding convolutional token embedding (method 1) and convolutional projections (method 2), we observe large improvements of 8% and 9% on DSC which is competitive with pre-trained

Methods	DSC	RV	Myo	LV
R50 UNet	87.60	84.62	84.52	93.68
R50 AttnUNet	86.90	83.27	84.33	93.53
R50 ViT	86.19	82.51	83.01	93.05
TransUnet	89.71	86.67	87.27	95.18
Swin-Unet	88.07	85.77	84.42	94.03
MT-UNet	<u>90.43</u>	86.64	<u>89.04</u>	95.62
Ours	91.37	89.20	89.47	<u>95.42</u>

Table 2. Experimental results of ACDC.

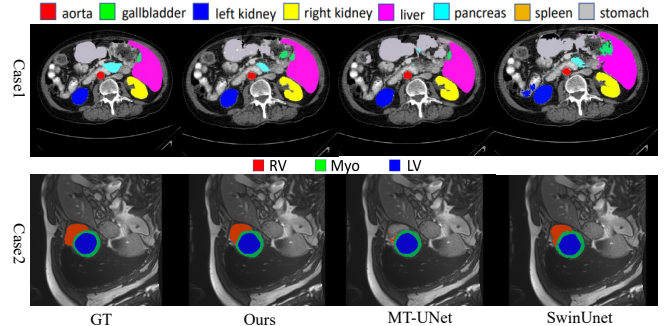


Fig. 4. Visualization of segmentation results on two datasets.

Swin-Unet emphasizing the importance of adding local modeling ability to Transformers. Removing the position embedding in early stages and using a convolution instead of a linear layer to fine-tune the attention computation (method 3) leads to a slight increase in performance and parameter reduction. Method 4 combines the CST block with the DSF module leading to an improved DSC and HD without extra parameters. After utilizing convolutional up-sampling and feature fusion module, SC, for merging information during skip connection, our best performing model method 6 achieves 3.17% improvements on DSC. Comparison with method 5, shows that fully convolutional pure Transformers can track the position of pixels better without needing an extra positional embedding, and that spatial feature extraction, is in fact, a necessity for Transformers.

Methods	DSC	HD	Emb	Proj	Pos	Att	DSF	SC	#param
0 (Base)	60.80	54.35			✓				27.15
1	68.57	51.02	✓		✓				21.55
2	77.47	18.54	✓	✓	✓				21.55
3	78.32	25.43	✓	✓	×	✓			19.63
4	79.04	22.96	✓	✓	×	✓	✓		19.84
5	81.93	24.59	✓	✓	✓	✓	✓	✓	24.68
6	82.21	27.02	✓	✓	×	✓	✓	✓	24.68

Table 3. Ablation study on modules used in CS-Unet on Synapse.

4. CONCLUSIONS

In this work, we presented the effects of introducing convolutions to Transformer blocks and to a multi-stage Transformer network to alleviate limitations of non-locality and need for extensive pre-training that Transformers in MISS face. Extensive experiments demonstrated that merging Convolutions with MHSAs and FFNs to create our CST layer, provided inherent local context inside Transformer blocks. Based on CST, our compact, accurate and pure Transformer architecture, CS-Unet, achieved superior performance without pre-training while maintaining less parameters.

5. REFERENCES

- [1] 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*. Springer, 2015, pp. 234–241.
- [2] 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*, pp. 3–11. Springer, 2018.
- [3] 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)*. IEEE, 2020, pp. 1055–1059.
- [4] 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*, vol. 53, pp. 197–207, 2019.
- [5] Chaitanya Kaul, Suresh Manandhar, and Nick Pears, “Focusnet: An attention-based fully convolutional network for medical image segmentation,” in *2019 IEEE 16th international symposium on biomedical imaging (ISBI 2019)*. IEEE, 2019, pp. 455–458.
- [6] Chaitanya Kaul, Nick Pears, Hang Dai, Roderick Murray-Smith, and Suresh Manandhar, “Focusnet++: Attentive aggregated transformations for efficient and accurate medical image segmentation,” in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2021, pp. 1042–1046.
- [7] 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.
- [8] 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*. Springer, 2021, pp. 14–24.
- [9] Bingzhi Chen, Yishu Liu, Zheng Zhang, Guangming Lu, and David Zhang, “Transattunet: Multi-level attention-guided unet with transformer for medical image segmentation,” *arXiv preprint arXiv:2107.05274*, 2021.
- [10] 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*. Springer, 2021, pp. 36–46.
- [11] Davood Karimi, Serge Didenko Vasylechko, and Ali Gholipour, “Convolution-free medical image segmentation using transformers,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2021, pp. 78–88.
- [12] 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.
- [13] 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.
- [14] Xiaohong Huang, Zhifang Deng, Dandan Li, and Xueguang Yuan, “Missformer: An effective medical image segmentation transformer,” *arXiv preprint arXiv:2109.07162*, 2021.
- [15] Hong-Yu Zhou, Jiansen Guo, Yinghao Zhang, Lequan Yu, Liansheng Wang, and Yizhou Yu, “nnformer: Interleaved transformer for volumetric segmentation,” *arXiv preprint arXiv:2109.03201*, 2021.
- [16] Hongyi Wang, Shiao Xie, Lanfen Lin, Yutaro Iwamoto, Xian-Hua Han, Yen-Wei Chen, and Ruofeng Tong, “Mixed transformer u-net for medical image segmentation,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 2390–2394.
- [17] Athanasios Tragakis, Chaitanya Kaul, and Husmeier Dirk Murray-Smith Roderick, “The fully convolutional transformer for medical image segmentation,” in <https://arxiv.org/abs/2206.00566>, 2022.
- [18] 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*, 2021, pp. 10012–10022.
- [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,” in *International Conference on Learning Representations*, 2020.
- [20] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang, “Cvt: Introducing convolutions to vision transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 22–31.
- [21] Bennett Landman, Zhoubing Xu, J Igelsias, Martin Styner, T Langerak, and Arno Klein, “Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge,” in *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge*, 2015, vol. 5, p. 12.
- [22] Olivier Bernard, Alain Lalonde, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al., “Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved?,” *IEEE transactions on medical imaging*, vol. 37, no. 11, pp. 2514–2525, 2018.
- [23] Ilya Loshchilov and Frank Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.
- [24] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou, “Going deeper with image transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 32–42.