

Published in final edited form as:

IEEE Trans Med Imaging. 2021 February 01; 40(2): 699–711. doi:10.1109/TMI.2020.3035253.

CA-Net: Comprehensive Attention Convolutional Neural Networks for Explainable Medical Image Segmentation

Ran Gu,

School of Mechanical and Electrical Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China

Guotai Wang,

School of Mechanical and Electrical Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China

Tao Song,

SenseTime Research, Shanghai 200233, China

Rui Huang,

SenseTime Research, Shanghai 200233, China

Michael Aertsen,

Department of Radiology, University Hospitals Leuven, 3000 Leuven, Belgium

Jan Deprest,

School of Biomedical Engineering and Imaging Sciences, King's College London, London WC2R 2LS, U.K.

Department of Obstetrics and Gynaecology, University Hospitals Leuven, 3000 Leuven, Belgium

Institute for Women's Health, University College London, London WC1E 6BT, U.K.

Sébastien Ourselin,

School of Biomedical Engineering and Imaging Sciences, King's College London, London WC2R 2LS, U.K.

Tom Vercauteren,

School of Biomedical Engineering and Imaging Sciences, King's College London, London WC2R 2LS, U.K.

Member, IEEE, Shaoting Zhang

School of Mechanical and Electrical Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China

SenseTime Research, Shanghai 200233, China

Guotai Wang: guotai.wang@uestc.edu.cn

Abstract

Correspondence to: Guotai Wang, guotai.wang@uestc.edu.cn.

Corresponding author: Guotai Wang.

Accurate medical image segmentation is essential for diagnosis and treatment planning of diseases. Convolutional Neural Networks (CNNs) have achieved state-of-the-art performance for automatic medical image segmentation. However, they are still challenged by complicated conditions where the segmentation target has large variations of position, shape and scale, and existing CNNs have a poor explainability that limits their application to clinical decisions. In this work, we make extensive use of multiple attentions in a CNN architecture and propose a comprehensive attention-based CNN (CA-Net) for more accurate and explainable medical image segmentation that is aware of the most important spatial positions, channels and scales at the same time. In particular, we first propose a joint spatial attention module to make the network focus more on the foreground region. Then, a novel channel attention module is proposed to adaptively recalibrate channel-wise feature responses and highlight the most relevant feature channels. Also, we propose a scale attention module implicitly emphasizing the most salient feature maps among multiple scales so that the CNN is adaptive to the size of an object. Extensive experiments on skin lesion segmentation from ISIC 2018 and multi-class segmentation of fetal MRI found that our proposed CA-Net significantly improved the average segmentation Dice score from 87.77% to 92.08% for skin lesion, 84.79% to 87.08% for the placenta and 93.20% to 95.88% for the fetal brain respectively compared with U-Net. It reduced the model size to around 15 times smaller with close or even better accuracy compared with state-of-the-art DeepLabv3+. In addition, it has a much higher explainability than existing networks by visualizing the attention weight maps. Our code is available at <https://github.com/HiLab-git/CA-Net>.

Index Terms

Attention; convolutional neural network; medical image segmentation; explainability

I Introduction

AUTOMATIC medical image segmentation is important for facilitating quantitative pathology assessment, treatment planning and monitoring disease progression [1]. However, this is a challenging task due to several reasons. First, medical images can be acquired with a wide range of protocols and usually have low contrast and inhomogeneous appearances, leading to over-segmentation and undersegmentation [2]. Second, some structures have large variation of scales and shapes such as skin lesion in dermoscopic images [3], making it hard to construct a prior shape model. In addition, some structures may have large variation of position and orientation in a large image context, such as the placenta and fetal brain in Magnetic Resonance Imaging (MRI) [2], [4], [5]. To achieve good segmentation performance, it is highly desirable for automatic segmentation methods to be aware of the scale and position of the target.

With the development of deep Convolutional Neural Networks (CNNs), state-of-the-art performance has been achieved for many segmentation tasks [1]. Compared with traditional methods, CNNs have a higher representation ability and can learn the most useful features automatically from a large dataset. However, most existing CNNs are faced with the following problems: Firstly, by design of the convolutional layer, they use shared weights at different spatial positions, which may lead to a lack of spatial awareness and thus have

reduced performance when dealing with structures with flexible shapes and positions, especially for small targets. Secondly, they usually use a very large number of feature channels, while these channels may be redundant. Many networks such as the U-Net [6] use a concatenation of low-level and high-level features with different semantic information. They may have different importance for the segmentation task, and highlighting the relevant channels while suppressing some irrelevant channels would benefit the segmentation task [7]. Thirdly, CNNs usually extract multi-scale features to deal with objects at different scales but lack the awareness of the most suitable scale for a specific image to be segmented [8]. Last but not least, the decisions of most existing CNNs are hard to explain and employed in a black box manner due to their nested non-linear structure, which limits their application to clinical decisions.

To address these problems, attention mechanism is promising for improving CNNs' segmentation performance as it mimics the human behavior of focusing on the most relevant information in the feature maps while suppressing irrelevant parts. Generally, there are different types of attentions that can be exploited for CNNs, such as paying attention to the relevant spatial regions, feature channels and scales. As an example of spatial attention, the Attention Gate (AG) [9] generates soft region proposals implicitly and highlights useful salient features for the segmentation of abdominal organs. The Squeeze and Excitation (SE) block [7] is one kind of channel attention and it recalibrates useful channel feature maps related to the target. Qin [10] used an attention to deal with multiple parallel branches with different receptive fields for brain tumor segmentation, and the same idea was used in prostate segmentation from ultrasound images [11]. However, these works have only demonstrated the effectiveness of using a single or two attention mechanisms for segmentation that may limit the performance and explainability of the network. We assume that taking a more comprehensive use of attentions would boost the segmentation performance and make it easier to understand how the network works.

For artificial intelligence systems, the explainability is highly desirable when applied to medical diagnosis [12]. The explainability of CNNs has a potential for verification of the prediction, where the reliance of the networks on the correct features must be guaranteed [12]. It can also help human understand the model's weaknesses and strengths in order to improve the performance and discover new knowledge distilled from a large dataset. In the segmentation task, explainability helps developers interpret and understand how the decision is obtained, and accordingly modify the network in order to gain better accuracy. Some early works tried to understand CNNs' decisions by visualizing feature maps or convolution kernels in different layers [13]. Other methods such as Class Activation Map (CAM) [14] and Guided Back Propagation (GBP) [15] are mainly proposed for explaining decisions of CNNs in classification tasks. However, explainability of CNNs in the context of medical image segmentation has rarely been investigated [16], [17]. Schlemper *et al.* [16] proposed attention gate that implicitly learn to suppress irrelevant region while highlighting salient features. Furthermore, Roy *et al.* [17] introduced spatial and channel attention at the same time to boost meaningful features. In this work, we take advantages of spatial, channel and scale attentions to interpret and understand how the pixel-level predictions are obtained by our network. Visualizing the attention weights obtained by our network not only helps to

understand which image region is activated for the segmentation result, but also sheds light on the scale and channel that contribute most to the prediction.

To the best of our knowledge, this is the first work on using comprehensive attentions to improve the performance and explainability of CNNs for medical image segmentation. The contribution of this work is three-fold. First, we propose a novel Comprehensive Attention-based Network (i.e., CA-Net) in order to make a complete use of attentions to spatial positions, channels and scales. Second, to implement each of these attentions, we propose novel building blocks including a dual-pathway multi-scale spatial attention module, a novel residual channel attention module and a scale attention module that adaptively selects features from the most suitable scales. Thirdly, we use the comprehensive attention to obtain good explainability of our network where the segmentation result can be attributed to the relevant spatial areas, feature channels and scales. Our proposed CA-Net was validated on two segmentation tasks: binary skin lesion segmentation from dermoscopic images and multi-class segmentation of fetal MRI (including the fetal brain and the placenta), where the objects vary largely in position, scale and shape. Extensive experiments show that CA-Net outperforms its counterparts that use no or only partial attentions. In addition, by visualizing the attention weight maps, we achieved a good explainability of how CA-Net works for the segmentation tasks.

II Related Works

A CNNs for Image Segmentation

Fully Convolutional Network (FCN) [18] frameworks such as DeepLab [8] are successful methods for natural semantic image segmentation. Subsequently, an encoder-decoder network SegNet [19] was proposed to produce dense feature maps. DeepLabv3+ [20] extended DeepLab by adding a decoder module and using depth-wise separable convolution for better performance and efficiency.

In medical image segmentation, FCNs also have been extensively exploited in a wide range of tasks. U-Net [6] is a widely used CNN for 2D biomedical image segmentation. The 3D U-Net [21] and V-Net [22] with similar structures were proposed for 3D medical image segmentation. In [23], a dilated residual and pyramid pooling network was proposed for automated segmentation of melanoma. Some other CNNs with good performance for medical image segmentation include High-Res3DNet [24], DeepMedic [25], and H-DenseUNet [26], etc. However, these methods only use position-invariant kernels for learning, without focusing on the features and positions that are more relevant to the segmentation object. Meanwhile, they have a poor explainability as they provide little mechanism for interpreting the decision-making process.

B Attention Mechanism

In computer vision, there are attention mechanisms applied in different task scenarios [27]–[29]. Spatial attention has been used for image classification [27] and image caption [29], etc. The learned attention vector highlights the salient spatial areas of the sequence conditioned on the current feature while suppressing the irrelevant counter-parts, making the

prediction more contextualized. The SE block using a channel-wise attention was originally proposed for image classification and has recently been used for semantic segmentation [26], [28]. These ideas of attention mechanisms work by generating a context vector which assigns weights on the input sequence. In [30], an attention mechanism is proposed to learn to softly weight feature maps at multiple scales. However, this method feeds multiple resized input images to a shared deep network, which requires human expertise to choose the proper sizes and is not self-adaptive to the target scale.

Recently, to leverage attention mechanism for medical image segmentation, Oktay *et al.* [9] combined spatial attention with U-Net for abdominal pancreas segmentation from CT images. Roy *et al.* [17] proposed concurrent spatial and channel wise ‘Squeeze and Excitation’ (scSE) frameworks for whole brain and abdominal multiple organs segmentation. Qin *et al.* [10] and Wang *et al.* [11] got feature maps of different sizes from middle layers and recalibrate these feature maps by assigning an attention weight. Despite the increasing number of works leveraging attention mechanisms for medical image segmentation, they seldom pay attention to feature maps at different scales. What’s more, most of them focus on only one or two attention mechanisms, and to the best of our knowledge, the attention mechanisms have not been comprehensively incorporated to increase the accuracy and explainability of segmentation tasks.

III Methods

A Comprehensive-Attention CNN

The proposed CA-Net making use of comprehensive attentions is shown in Fig. 1, where we add specialized convolutional blocks to achieve comprehensive attention guidance with respect to the space, channel and scale of the feature maps simultaneously. Without loss of generality, we choose the powerful structure of the U-Net [6] as the backbone. The U-Net backbone is an end-to-end-trainable network consisting of an encoder and a decoder with shortcut connection at each resolution level. The encoder is regarded as a feature extractor that obtains high-dimensional features across multiple scales sequentially, and the decoder utilizes these encoded features to recover the segmentation target.

Our CA-Net has four spatial attention modules (SA_{1-4}), four channel attention modules (CA_{1-4}) and one scale attention module (LA), as shown in Fig. 1. The spatial attention is utilized to strengthen the region of interest on the feature maps while suppressing the potential background or irrelevant parts. Hence, we propose a novel multi-scale spatial attention module that is a combination of non-local block [31] at the lowest resolution level (SA_1) and dual-pathway AG [9] at the other resolution levels (SA_{2-4}). We call it as the joint spatial attention ($Js-A$) that enhances inter-pixel relationship to make the network better focus on the segmentation target. Channel attention (CA_{1-4}) is used to calibrate the concatenation of low-level and high-level features in the network so that the more relevant channels are weighted by higher coefficients. Unlike the SE block that only uses average-pooling to gain channel attention weight, we additionally introduce max-pooled features to exploit more salient information for channel attention [32]. Finally, we concatenate feature maps at multiple scales in the decoder and propose a scale attention module (LA) to

highlight features at the most relevant scales for the segmentation target. These different attention modules are detailed in the following.

1) Joint Spatial Attention Module: The joint spatial attention is inspired by the non-local network [31] and AG [9]. We use four attention blocks (SA_{1-4}) in the network to learn attention maps at four different resolution levels, as shown in Fig. 1. First, for the spatial attention at the lowest resolution level (SA_1), we use a non-local block that captures interactions between all pixels with a better awareness of the entire context. The detail of (SA_1) is shown in Fig. 2(a). Let x represent the input feature map with a shape of $256 \times H \times W$, where 256 is the input channel number, and H , W represent the height and width, respectively. We first use three parallel 1×1 convolutional layers with an output channel number of 64 to reduce the dimension of x , obtaining three compressed feature maps x' , x'' and x''' , respectively, and they have the same shape of $64 \times H \times W$. The three feature maps can then be reshaped into 2D matrices with shape of $64 \times HW$. A spatial attention coefficient map is obtained as:

$$\alpha_1 = \sigma(x'^T \cdot x''') \quad (1)$$

where T means matrix transpose operation. $\alpha_1 \in (0,1)^{HW \times HW}$ is a square matrix, and σ is a row-wise *Softmax* function so that the sum of each row equals to 1.0. α_1 is used to represent the feature of each pixel as a weighted sum of features of all the pixels, to ensure the interaction among all the pixels. The calibrated feature map in the reduced dimension is:

$$\hat{x} = \alpha_1 \cdot x'''^T \quad (2)$$

\hat{x} is then reshaped to $64 \times H \times W$, and we use Φ^{256} that is a 1×1 convolution with batch normalization and output channel number of 256 to expand x to match the channel number of \hat{x} . A residual connection is finally used to facilitate the information propagation during training, and the output of SA_1 is obtained as:

$$y_{SA1} = \Phi^{256}(\hat{x}) + x \quad (3)$$

Second, as the increased memory consumption limits applying the non-local block to feature maps with higher resolution, we extend AG to learn attention coefficients in SA_{2-4} . As a single AG may lead to a noisy spatial attention map, we propose a dual-pathway spatial attention that exploits two AGs in parallel to strengthen the attention to the region of interest as well as reducing noise in the attention map. Similarly to model ensemble, combining two AGs in parallel has a potential to improve the robustness of the segmentation. The details about a single pathway AG are shown in Fig. 2(b). Let x^l represent the low-level feature map at the scale s in the encoder, and x^h represent a high-level feature map up-sampled from the end of the decoder at scale $s+1$ with a lower spatial resolution, so that x^h and x^l have the same shape. In a single-pathway AG, the *query* feature x^h is used to calibrate the low-level *key* feature x^l . As shown in Fig. 2(b), x^h and x^l are compressed by a 1×1 convolution with an output channel number C (e.g., 64) respectively, and the results are summed and followed

by a *ReLU* activation function. Feature map obtained by the *ReLU* is then fed into another 1×1 convolution with one output channel followed by a *Sigmoid* function to obtain a pixel-wise attention coefficient $\alpha \in [0, 1]^{H \times W}$. x^l is then multiplied with α to be calibrated. In our dual-pathway AG, the spatial attention maps in the two pathways are denoted as $\hat{\alpha}$ and $\tilde{\alpha}$ respectively. As shown in Fig. 2(c), the output of our dual-pathway AG for SA_s ($s = 2, 3, 4$) is obtained as:

$$y_{SA_s} = \text{ReLU} \left[\Phi^C \left((x^l \cdot \hat{\alpha}) \odot (x^l \cdot \tilde{\alpha}) \right) \right] \quad (4)$$

where \odot means channel concatenation. Φ^C denotes 1×1 convolution with C output channels followed by batch normalization. Here C is 64, 32 and 16 for SA_2 , SA_3 and SA_4 , respectively.

2) Channel Attention Modules—In our network, channel concatenation is used to combine the spatial attention-calibrated low-level features from the encoder and higher-level features from the decoder as shown in Fig. 1. Feature channels from the encoder contain mostly low-level information, and their counterparts from the decoder contain more semantic information. Therefore, they may have different importance for the segmentation task. To better exploit the most useful feature channels, we introduce channel attention to automatically highlight the relevant feature channels while suppressing irrelevant channels. The details of proposed channel attention module (CA_{1-4}) is shown in Fig. 3.

Unlike previous SE block that only utilized average-pooled information to excite feature channels [7], we use max-pooled features additionally to keep more information [32]. Similarly, let x represent the concatenated input feature map with C channels, a global average pooling P_{avg} and a global maximal pooling P_{max} are first used to obtain the global information of each channel, and the outputs are represented as $P_{avg}(x) \in R^{C \times 1 \times 1}$ and $P_{max}(x) \in R^{C \times 1 \times 1}$, respectively. A multiple layer perception (*MLP*) \mathcal{M} is used to obtain the channel attention coefficient $\beta \in [0, 1]^{C \times 1 \times 1}$, and \mathcal{M} is implemented by two fully connected layers, where the first one has an output channel number of C/r followed by *ReLU* and the second one has an output channel number of C . We set $r = 2$ counting the trade-off of performance and computational cost [7]. Note that a shared \mathcal{M} is used for $P_{avg}(x)$ and $P_{max}(x)$, and their results are summed and fed into a *Sigmoid* to obtain β . The output of our channel attention module is obtained as:

$$y_{CA} = x \cdot \beta + x \quad (5)$$

where we use a residual connection to benefit the training. In our network, four channel attention modules (CA_{1-4}) are used (one for each concatenated feature), as shown in Fig. 1.

3) Scale Attention Modulue—The U-Net backbone obtains feature maps in different scales. To better deal with objects in different scales, it is reasonable to combine these features for the final prediction. However, for a given object, these feature maps at different scales may have different relevance to the object. It is desirable to automatically determine the scale-wise weight for each pixel, so that the network can be adaptive to corresponding

scales of a given input. Therefore, we propose a scale attention module to learn image-specific weight for each scale automatically to calibrate the features at different scales, which is used at the end of the network, as shown in Fig. 1.

Our proposed *LA* block is illustrated in Fig. 4. We first use bilinear interpolation to resample the feature maps F_s at different scales ($s = 1, 2, 3, 4$) obtained by the decoder to the original image size. To reduce the computational cost, these feature maps are compressed into four channels using 1×1 convolutions, and the compressed results from different scales are concatenated into a hybrid feature map \hat{F} . Similarly to our *CA*, we combine P_{avg} P_{max} with *MLP* to obtain a coefficient for each channel (i.e., scale here), as shown in Fig. 4. The scale coefficient attention vector is denoted as $\gamma \in [0, 1]^{4 \times 1 \times 1}$. To distribute multi-scale soft attention weight on each pixel, we additionally use a spatial attention block *LA** taking $\hat{F} \cdot \gamma$ as input to generate spatial-wise attention coefficient $\gamma^* \in [0, 1]^{1 \times H \times W}$, so that $\gamma \cdot \gamma^*$ represents a pixel-wise scale attention. *LA** consists of one 3×3 and one 1×1 convolutional layers, where the first one has 4 output channels followed by *ReLU*, and the second one has 4 output channels followed by *Sigmoid*. The final output of our *LA* module is:

$$y_{LA} = \hat{F} \cdot \gamma \cdot \gamma^* + \hat{F} \cdot \gamma + \hat{F} \quad (6)$$

where the residual connections are again used to facilitate the training, as shown in Fig. 4. Using scale attention module enables the CNN to be aware of the most suitable scale (how big the object is).

IV Experimental Results

We validated our proposed framework with two applications: (i) Binary skin lesion segmentation from dermoscopic images. (ii) Multi-class segmentation of fetal MRI, including the fetal brain and the placenta. For both applications, we implemented ablation studies to validate the effectiveness of our proposed CA-Net and compared it with state-of-the-art networks. Experimental results of these two tasks will be detailed in Section IV-B and Section IV-C, respectively.

A Implementation and Evaluation Methods

All methods were implemented in Pytorch framework.^{1,2} We used Adaptive Moment Estimation (Adam) for training with initial learning rate 10^{-4} , weight decay 10^{-8} , batch size 16, and iteration 300 epochs. The learning rate is decayed by 0.5 every 256 epochs. The feature channel number in the first block of our CA-Net was set to 16 and doubled after each down-sampling. In *MLPs* of our *CA* and *LA* modules, the channel compression factor r was 2 according to [7]. Training was implemented on one NVIDIA Geforce GTX 1080 Ti GPU. We used Soft Dice loss function for the training of each network and used the best performing model on the validation set among all the epochs for testing. We used 5-fold

¹ <https://pytorch.org/>

² <https://github.com/HiLab-git/CA-Net>

cross-validation for final evaluation. After training, the model was deployed at SenseCare platform to support clinic research [33].

Quantitative evaluation of segmentation accuracy was based on: (i) The Dice score between a segmentation and the ground truth, which is defined as:

$$\text{Dice} = \frac{2|R_a \cap R_b|}{|R_a| + |R_b|} \quad (7)$$

where R_a and R_b denote the region segmented by algorithm and the ground truth, respectively. (ii) Average symmetric surface distance (ASSD). Supposing S_a and S_b represent the set of boundary points of the automatic segmentation and the ground truth respectively, the ASSD is defined as:

$$\text{ASSD} = \frac{1}{|S_a| + |S_b|} \times \left(\sum_{a \in S_a} d(a, S_b) + \sum_{b \in S_b} d(b, S_a) \right) \quad (8)$$

where $d(v, S_a) = \min_{w \in S_a} (\|v - w\|)$ denotes the minimum Euclidean distance from point v to all the points of S_a .

B Lesion Segmentation From Dermoscopic Images

With the emergence of automatic analysis algorithms, it becomes possible that accurate automatic skin lesion boundary segmentation helps dermatologists for early diagnosis and screening of skin diseases quickly. The main challenge for this task is that the skin lesion areas have various scales, shapes and colors, which requires automatic segmentation methods to be robust against shape and scale variations of the lesion [34].

1) Dataset—For skin lesion segmentation, we used the public available training set of ISIC 2018³ with 2594 images and their ground truth. We randomly split the dataset into 1816, 260 and 518 for training, validation and testing respectively. The original size of the skin lesion segmentation dataset ranged from 720×540 to 6708×4439 , and we resized each image to 256×342 and normalized it by the mean value and standard deviation. During training, random cropping with a size of 224×300 , horizontal and vertical flipping, and random rotation with a angle in $(-\pi/6, \pi/6)$ were used for data augmentation.

2) Comparison of Spatial Attention Methods—We first investigated the effectiveness of our spatial attention modules without using the channel attention and scale attention modules. We compared different variants of our proposed multi-level spatial attention: 1) Using standard single-pathway AG [9] at the position of SA_{1-4} , which is referred to as s-AG; 2) Using the dual-pathway AG at the position of SA_{1-4} , which is referred to as t-AG; 3) Using the non-local block of SA_1 only, which is referred to as n-Local [31]. Our proposed joint attention method using non-local block in SA_1 and dual-pathway AG in SA

³ <https://challenge2018.isic-archive.com/>

2_{-4} is denoted as Js-A. For the baseline U-Net, the skip connection was implemented by a simple concatenation of the corresponding features in the encoder and the decoder [6]. For other compared variants that do not use SA_{2-4} , their skip connections were implemented as the same as that of U-Net. Table I shows a quantitative comparison between these methods. It can be observed that all the variants using spatial attention lead to higher segmentation accuracy than the baseline. Also, we observe that dual-pathway spatial AG is more effective than single-pathway AG, and our joint spatial attention block outperforms the others. Compared with the standard AG [9], our proposed spatial attention improved the average Dice from 88.46% to 90.83%.

Fig. 5(a) visualizes the spatial attention weight maps obtained by s-AG, t-AG and our Js-A. It can be observed that single-pathway AG pays attention to almost every pixel, which means it is dispersive. The dual-pathway AG is better than the single-pathway AG but still not self-adaptive enough. In comparison, our proposed Js-A pays a more close attention to the target than the above methods.

Fig. 5(b) presents some examples of qualitative segmentation results obtained by the compared methods. It can be seen that introducing spatial attention block in neural network largely improves the segmentation accuracy. Furthermore, the proposed Js-A gets better result than the other spatial attention methods in both cases. In the second case where the lesion has a complex shape and blurry boundary, our proposed Js-A keeps a better result.

We observed that there may exist skew between original annotation and our cognition in ISIC 2018, as shown in Fig. 5. This is mainly because that the image contrast is often low along the true boundary, and the exact lesion boundary requires some expertise to delineate. The ISIC 2018 dataset was annotated by experienced dermatologists, and some annotations may be different from what a non-expert thinks.

3) Comparison of Channel Attention Methods—In this comparison, we only introduced channel attention modules to verify the effectiveness of our proposed method. We first investigate the effect of position in the network the channel attention module plugged in: 1) the encoder, 2) the decoder, 3) both the encoder and decoder. These three variants are referred to as C-A (Enc), C-A (Dec) and C-A (Enc&Dec) respectively. We also compared the impact of using and not using max pooling for the channel attention module.

Table II shows the quantitative comparison of these variants, which demonstrates that channel attention blocks indeed improve the segmentation performance. Moreover, channel attention block with additional max-pooled information generally performs better than those using average pooling only. Additionally, we find that channel attention block plugged in the decoder performs better than plugged into the encoder or both the encoder and decoder. The C-A (Dec) achieved an average Dice score of 91.68%, which outperforms the others.

Fig. 6 shows the visual comparison of our proposed channel attention and its variants. The baseline U-Net has a poor performance when the background has a complex texture, and the channel attention methods improve the accuracy for these cases. Clearly, our proposed channel attention module C-A (Dec) obtains higher accuracy than the others.

4) Comparison of Scale Attention Methods—In this comparison, we only introduced scale attention methods to verify the effectiveness of our proposed scale attention. Let L-A (1-K) denote the scale attention applied to the concatenation of feature maps from scale 1 to K as shown in Fig. 1. To investigate the effect of number of feature map scales on the segmentation, we compared our proposed method with $K = 2, 3, 4$ and 5 respectively.

Table III shows the quantitative comparison results. We find that combining features for multiple scales outperforms the baseline. When we concatenated features from scale 1 to 4, the Dice score and ASSD can get the best values of 91.58% and 0.66 pixels respectively. However, when we combined features from all the 5 scales, the segmentation accuracy is decreased. This suggests that the feature maps at the lowest resolution level is not suitable for predicting pixel-wise label in details. As a result, we only fused the features from scale 1 to 4, as shown in Fig. 1 in the following experiments. Fig. 7 shows a visualization of skin lesion segmentation based on different scale attention variants.

Fig. 8 presents the visualization of pixel-wise scale attention coefficient $\gamma \cdot \gamma^*$ at different scales, where the number under each picture denotes the scale-wise attention coefficient γ . This helps to better understand the importance of feature at different scales. The two cases show a large and a small lesion respectively. It can be observed that the large lesion has higher global attention coefficients γ in scale 2 and 3 than the small lesion, and γ in scale 1 has a higher value in the small lesion than the large lesion. The pixel-wise scale attention maps also show that the strongest attention is paid to scale 2 in the first row, and scale 1 in the second row. This demonstrates that the network automatically leans to focus on the corresponding scales for segmentation of lesions at different sizes.

5) Comparison of Partial and Comprehensive Attention—To investigate the effect of combining different attention mechanisms, we compared CA-Net with six variants of different combinations of the three basic spatial, channel and scale attentions. Here, SA means our proposed multi-scale joint spatial attention and CA represents our channel attention used only in the decoder of the backbone.

Table IV presents the quantitative comparison of our CA-Net and partial attention methods for skin lesion segmentation. From Table IV, we find that each of SA, CA and LA obtains performance improvement compared with the baseline U-Net. Combining two of these attention methods outperforms the methods using a single attention. Furthermore, our proposed CA-Net outperforms all other variants both in Dice score and ASSD, and the corresponding values are 92.08% and 0.58 pixels, respectively.

6) Comparison With the State-of-the-art Frameworks—We compared our CA-Net with three state-of-the-art methods: 1) DenseASPP [35] that uses DenseNet-121 [36] as the backbone; 2) RefineNet [37] that uses Resnet101 [38] as the backbone; 3) Two variants of DeepLabv3+ [20] that use Xception [39] and Dilated Residual Network (DRN) [40] as feature extractor, respectively. We retrained all these networks on ISIC 2018 and did not use their pre-trained models.

Quantitative comparison results of these methods are presented in Table V. It shows that all state-of-the-art methods have good performance in terms of Dice score and ASSD. Our CA-Net obtained a Dice score of 92.08%, which is a considerable improvement compared with the U-Net whose Dice is 87.77%. Though our CA-Net has a slightly lower performance than DeepLabv3+, the difference is not significant ($p\text{-value}=0.46 > 0.05$), and our CA-Net has a model size that is 15 times smaller with better explainability. For skin lesion segmentation, the average inference time per image for our CA-Net was 2.1ms, compared with 4.0ms and 3.4ms by DeepLabv3+ [20] and RefineNet [37], respectively. Fig. 9 shows the visual comparison of different CNNs dealing with skin lesion segmentation task.

C Segmentation of Multiple Organs From Fetal MRI

In this experiment, we demonstrate the effectiveness of our CA-Net in multi-organ segmentation, where we aim to jointly segment the placenta and the fetal brain from fetal MRI slices. Fetal MRI has been increasingly used to study fetal development and pathology, as it provides a better soft tissue contrast than more widely used prenatal sonography [4]. Segmentation of some important organs such as the fetal brain and the placenta is important for fetal growth assessment and motion correction [41]. Clinical fetal MRI data are often acquired with a large slice thickness for good contrast-to-noise ratio. Moreover, movement of the fetus can lead to inhomogeneous appearances between slices. Hence, 2D segmentation is considered more suitable than direct 3D segmentation from motion-corrupted MRI slices [2].

1) Dataset—The dataset consists of 150 stacks with three views (axial, coronal, and sagittal) of T2-weighted fetal MRI scans of 36 pregnant women in the second trimester with Single-shot Fast-Spin echo (SSFSE) with pixel size 0.74 to 1.58 mm and inter-slice spacing 3 to 4 mm. The gestational age ranged from 22 to 29 weeks. 8 of the fetuses were diagnosed with spinal bifida and the others had no fetal pathologies. All the pregnant women were above 18 years old, and the use of data was approved by the Research Ethics Committee of the hospital.

As the stacks contained an imbalanced number of slices covering the objects, we randomly selected 10 of these slices from each stack for the experiments. Then, we randomly split the slices at patient level and assigned 1050 for training, 150 for validation, and 300 for testing. The test set contains 110 axial slices, 80 coronal slices, and 110 sagittal slices. Manual annotations of the fetal brain and placenta by an experienced radiologist were used as the ground truth. We trained a multi-class segmentation network for simultaneous segmentation of these two organs. Each slice was resized to 256×256 . We randomly flipped in x and y axis and rotated with an angle in $(-\pi/6, \pi/6)$ for data augmentation. All the images were normalized by the mean value and standard deviation.

2) Comparison of Spatial Attention Methods—In parallel to section IV-B.2, we compared our proposed Js-A with: (1) the single-pathway AG (s-AG) only, (2) the dual-pathway AG (t-AG) only, (3) the non-local block (n-local) only.

Table VI presents quantitative comparison results between these methods. From Table VI, we observe that all the variants of spatial attention modules led to better Dice and ASSD

scores. It can be observed that dual-pathway AG performs more robustly than the single-pathway AG, and Js-A module can get the highest scores, with Dice of 95.47% and ASSD of 0.30 pixels, respectively. Furthermore, in placenta segmentation which has fuzzy tissue boundary, our model still maintains encouraging segmentation performance, with Dice score of 85.65%, ASSD of 0.58 pixels, respectively.

Fig. 10 presents a visual comparison of segmentation results obtained by these methods as well as their attention weight maps. From Fig. 10(b), we find that spatial attention has reliable performance when dealing with complex object shapes, as highlighted by red arrows. Meanwhile, with visualizing their spatial attention weight maps as shown in Fig. 10(a), our proposed Js-A has a greater ability to focus on the target areas compared with the other methods as it distributes a higher and closer weight on the target of our interest.

3) Comparison of Channel Attention Methods—We compared the proposed channel attention method with the same variants as listed in section IV-B.3 for fetal MRI segmentation. The comparison results are presented in Table VII. It shows that channel attention plugged in decoder brings noticeably fewer parameters and still maintains similar or higher accuracy than the other variants. We also compared using and not using max-pooling in the channel attention block. From Table VII, we can find that adding extra max-pooled information indeed increases performance in terms of Dice and ASSD, which proves the effectiveness of our proposed method.

4) Comparison of Scale Attention Methods—In this comparison, we investigate the effect of concatenating different number of feature maps from scale 1 to K as described in section IV-B.4, and Table VIII presents the quantitative results. Analogously, we observe that combining features from multiple scales outperforms the baseline. When we concatenate features from scale 1 to 4, we get the best results, and the corresponding Dice values for the placenta and the fetal brain are 86.21% and 95.18%, respectively. When the feature maps at the lowest resolution is additional used, i.e., L-A (1-5), the Dice scores are slightly reduced.

Fig. 11 shows the visual comparison of our proposed scale attention and its variants. In the second row, the placenta has a complex shape with a long tail, and combining features from scale 1 to 4 obtained the best performance. Fig. 12 shows the visual comparison of scale attention weight maps on fetal MRI. From the visualized pixel-wise scale attention maps, we observed that the network pays much attention to scale 1 in the first row where the fetal brain is small, and to scale 2 in the second row where the fetal brain is larger.

5) Comparison of Partial and Comprehensive Attention—Similar to Section IV-B.5, we compared comprehensive attentions with partial attentions in the task of segmenting fetal brain and placenta from fetal MRI. From Table IX, we find that models combining two of the three attention mechanism basically outperform variants that using a single attention mechanism. SA + CA gets the highest scores among three binary-attention methods, which can achieve Dice scores of 86.68% for the placenta and 95.42% for the fetal brain. Furthermore, our proposed CA-Net outperforms all these binary-attention methods,

achieving Dice scores of 87.08% for the placenta and 95.88% for the fetal brain, respectively. The ASSD value of CA-Net is lower than those of other methods.

6) Comparison of State-of-the-art Frameworks—We also compared our CA-Net with the state-of-the-art methods and their variants as implemented in section IV-B.5. The segmentation performance on images in axial, sagittal and coronal views was measured respectively. A quantitative evaluation of these methods for fetal MRI segmentation is listed in Table X. We observe that our proposed CA-Net obtained better Dice scores than the others in all the three views. Our CA-Net can improve the Dice scores by 2.35%, 1.78%, and 2.60% for placenta segmentation and 3.75%, 0.85%, and 2.84% for fetal brain segmentation in three views compared with U-Net, respectively, surpassing the existing attention method and the state-of-the-art segmentation methods. In addition, for the average Dice and ASSD values across the three views, CA-Net outperformed the others. Meanwhile, CA-Net has a much smaller model size compared with RefineNet [37] and DeepLabv3+ [20], which leads to lower computational cost for training and inference. For fetal MRI segmentation, the average inference time per image for our CA-Net was 1.5ms, compared with 3.4ms and 2.2ms by DeepLabv3+ and RefineNet, respectively. Qualitative results in Fig. 13 also show that CA-Net performs noticeably better than the baseline and the other methods for fetal MRI segmentation. In dealing with the complex shapes as shown in the first and fifth rows of Fig. 13, as well as the blurry boundary in the second row, CA-Net performs more closely to the authentic boundary than the other methods. Note that visualization of the spatial and scale attentions as show in Fig. 10 and Fig. 12 helps to interpret the decision of our CA-Net, but such explainability is not provided by DeepLabv3+, RefineNet and DenseASPP.

V Discussion and Conclusion

For a medical image segmentation task, some targets such as lesions may have a large variation of position, shape and scale, enabling the networks to be aware of the object's spatial position and size is important for accurate segmentation. In addition, convolution neural networks generate feature maps with a large number of channels, and concatenation of feature maps with different semantic information or from different scales are often used. Paying an attention to the most relevant channels and scales is an effective way to improve the segmentation performance. Using scale attention to adaptively make use of the features in different scales would have advantages in dealing with objects with a variation of scales. To take these advantages simultaneously, we take a comprehensive use of these complementary attention mechanisms, and our results show that CA-Net helps to obtain more accurate segmentation with only few parameters.

For explainable CNN, previous works like CAM [14] and GBP [15] mainly focused on image classification tasks and they only consider the spatial information for explaining the CNN's prediction. In addition, they are post-hoc methods that require additional computations after a forward pass prediction to interpret the prediction results. Differently from these methods, CA-Net gives a comprehensive interpretation of how each spatial position, feature map channel and scale is used for the prediction in segmentation tasks. What's more, we obtain these attention coefficients in a single forward pass and require no additional computations. By visualizing the attention maps in different aspects as show in

Fig. 5 and Fig. 8, we can better understand how the network works, which has a potential to help us improve the design of CNNs.

We have done experiments on two different image domains, i.e., RGB image and fetal MRI. These two are representative image domains, and in both cases our CA-Net has a considerable segmentation improvement compared with U-Net, which shows that the CA-Net has competing performance for different segmentation tasks in different modalities. It is of interest to apply our CA-Net to other image modalities such as the Ultrasound and other anatomies in the future.

In this work, we have investigated three main types of attentions associated with segmentation targets in various positions and scales. Recently, some other types of attentions have also been proposed in the literature, such as attention to parallel convolution kernels [42]. However, using multiple parallel convolution kernels will increase the model complexity.

Most of the attention blocks in our CA-Net are in the decoder. This is mainly because that the encoder acts as a feature extractor that is exploited to obtain enough candidate features. Applying attention at the encoder may lead some potentially useful features to be suppressed at an early stage. Therefore, we use the attention blocks in the decoder to highlight relevant features from all the candidate features. Specifically, following [9], the spatial attention is designed to use high-level semantic features in the decoder to calibrate low-level features in the encoder, so they are used at the skip connections after the encoder. The scale attention is designed to better fuse the raw semantic predictions that are obtained in the decoder, which should naturally be placed at the end of the network. For channel attentions, we tried to place them at different positions of the network, and found that placing them in the decoder is better than in the encoder. As shown in Table II, all the channel attention variants outperformed the baseline U-Net. However, using channel attention only in the decoder outperformed the variants with channel attention in the encoder. The reason may be that the encoding phase needs to maintain enough feature information, which confirms our assumption that suppressing some features at an early stage will limit the model's performance. However, some other attentions [42] might be useful in the encoder, which will be investigated in the future.

Differently from previous works that mainly focus on improving the segmentation accuracy while hard to explain, we aim to design a network with good comprehensive property, including high segmentation accuracy, efficiency and explainability at the same time. Indeed, the segmentation accuracy of our CA-Net is competing: It leads to a significant improvement of Dice compared with the U-Net (92.08% VS 87.77%) for skin lesion. Compared with state-of-the-art DeepLabv3+ and RefineNet, our CA-Net achieved very close segmentation accuracy with around 15 times fewer parameters. What's more, CA-Net is easy to explain as shown in Fig. 5, 8, 10, and 12, but DeepLabv3+ and RefineNet have poor explainability on how they localize the target region, recognize the scale and determine the useful features. Meanwhile, in fetal MRI segmentation, experimental results from Table X shows that our CA-Net has a considerable improvement compared with U-Net (Dice was 87.08% VS 84.79%), and it outperforms the state-of-the-art methods in all the three views.

Therefore, the superiority of our CA-Net is that it could achieve high explainability and efficiency than state-of-the-art methods while maintaining comparable or even better accuracy.

In the skin lesion segmentation task, we observe that our CA-Net leads to slightly inferior performance than Deeplabv3+, which is however without significant difference. We believe the reason is that Deeplabv3+ is mainly designed for natural image segmentation task, and the dermoscopic skin images are color images, which has a similar distribution of intensity to natural images. However, compared to Deeplabv3+, our CA-Net can achieve comparable performance, and it has higher explainability and 15 times fewer parameters, leading to higher computational efficiency. In the fetal MRI segmentation task, our CA-Net has distinctly higher accuracy than those state-of-the-art methods, which shows the effectiveness and good explainability of our method.

In conclusion, we propose a comprehensive attention-based convolutional neural network (CA-Net) that learns to make a comprehensive use of multiple attentions for better performance and explainability of medical image segmentation. We enable the network to adaptively pay attention to spatial positions, feature channels and object scales at the same time. Motivated by existing spatial and channel attention methods, we make further improvements to enhance the network's ability to focus on areas of interest. We propose a novel scale attention module implicitly emphasizing the most salient scales to obtain multiple-scale features. Experimental results show that compared with the state-of-the-art semantic segmentation models like Deeplabv3+, our CA-Net obtains comparable and even higher accuracy for medical image segmentation with a much smaller model size. Most importantly, CA-Net gains a good model explainability which is important for understanding how the network works, and has a potential to improve clinicians' acceptance and trust on predictions given by an artificial intelligence algorithm. Our proposed multiple attention modules can be easily plugged into most semantic segmentation networks. In the future, the method can be easily extended to segment 3D images.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 81771921 and Grant 61901084, in part by the Key Research and Development Project of Sichuan, China, under Grant 20ZDYF2817, in part by the Wellcome Trust under Grant WT101957 and Grant 203148/Z/16/Z, and in part by the Engineering and Physical Sciences Research Council (EPSRC) under Grant NS/A000027/1 and Grant NS/A000049/1. The work of Tom Vercauteren was supported by the Medtronic/Royal Academy of Engineering Research Chair under Grant RCSR18194.

References

- [1]. Litjens G, et al. A survey on deep learning in medical image analysis. *Med Image Anal.* 2017 Dec.42:60–88. [PubMed: 28778026]
- [2]. Wang G, et al. Interactive medical image segmentation using deep learning with image-specific fine tuning. *IEEE Trans Med Imag.* 2018 Jul; 37(7):1562–1573.
- [3]. Codella N, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC). *arXiv:1902.03368.* 2019
- [4]. Wang G, et al. DeepIGeoS: A deep interactive geodesic framework for medical image segmentation. *IEEE Trans Pattern Anal Mach Intell.* 2019 Jul; 41(7):1559–1572. [PubMed: 29993532]

- [5]. Salehi, SSM; , et al. Real-time automatic fetal brain extraction in fetal MRI by deep learning; Proc IEEE 15th Int Symp Biomed Imag (ISBI); 2018 Apr. 720–724.
- [6]. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation. Proc MICCAI. 2015 Oct.:234–241.
- [7]. Hu, J; Shen, L; Sun, G. Squeeze- and-excitation networks; Proc IEEE/CVF Conf Comput Vis Pattern Recognit; 2018 Jun. 7132–7141.
- [8]. Chen L-C, Papandreou G, Kokkinos I, Murphy K, Yuille AL. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. IEEE Trans Pattern Anal Mach Intell. 2018 Apr; 40(4):834–848. [PubMed: 28463186]
- [9]. Oktay O, et al. Attention U-Net: Learning where to look for the pancreas. Proc MIDL. 2018 Jul.:1–10.
- [10]. Qin Y, et al. Autofocus layer for semantic segmentation. Proc MICCAI. 2018 Sep.:603–611.
- [11]. Wang Y, et al. Deep attentional features for prostate segmentation in ultrasound. Proc MICCAI. 2018 Sep.:523–530.
- [12]. Samek W, Wiegand T, Müller K-R. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. arXiv:1708.08296. 2017
- [13]. Chen, R; Chen, H; Huang, G; Ren, J; Zhang, Q. Explaining neural networks semantically and quantitatively; Proc IEEE/CVF Int Conf Comput Vis (ICCV); 2019 Oct. 9187–9196.
- [14]. Zhou, B; Khosla, A; Lapedriza, A; Oliva, A; Torralba, A. Learning deep features for discriminative localization; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2016 Jun. 2921–2929.
- [15]. Tobias Springenberg J, Dosovitskiy A, Brox T, Ried-miller M. Striving for simplicity: The all convolutional net. arXiv:1412.6806. 2014
- [16]. Schlemper J, et al. Attention gated networks: Learning to leverage salient regions in medical images. Med Image Anal. 2019 Apr;53:197–207. [PubMed: 30802813]
- [17]. Roy AG, Navab N, Wachinger C. Concurrent spatial and channel ‘squeeze & excitation’ in fully convolutional networks. Proc MICCAI. 2018 Sep.:421–429.
- [18]. Long, J; Shelhamer, E; Darrell, T. Fully convolutional networks for semantic segmentation; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2015 Jun. 3431–3440.
- [19]. Badrinarayanan V, Kendall A, Cipolla R. SegNet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE Trans Pattern Anal Mach Intell. 2017 Dec; 39(12):2481–2495. [PubMed: 28060704]
- [20]. Chen L-C, Zhu Y, Papandreou G, Schroff F, Adam H. Encoderdecoder with atrous separable convolution for semantic image segmentation. Proc ECCV. 2018 Sep.:801–818.
- [21]. Çiçek O, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O. 3D U-Net: learning dense volumetric segmentation from sparse annotation. Proc MICCAI. 2016:424–432.
- [22]. Milletari, F; Navab, N; Ahmadi, S-A. V-Net: Fully convolutional neural networks for volumetric medical image segmentation; Proc 4th Int 3D Conf Vis (3DV); 2016 Oct. 565–571.
- [23]. Sarker MMK, et al. SLSDeep: Skin lesion segmentation based on dilated residual and pyramid pooling networks. Proc MICCAI. 2018 Sep.:21–29.
- [24]. Li W, Wang G, Fidon L, Ourselin S, Cardoso MJ, Vercauteren T. On the compactness, efficiency, and representation of 3D convolutional networks: Brain parcellation as a pretext task. Proc IPMI. 2017 Jun.:348–360.
- [25]. Kamnitsas K, et al. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med Image Anal. 2017 Feb;36:61–78. [PubMed: 27865153]
- [26]. Li, K; Wu, Z; Peng, K-C; Ernst, J; Fu, Y. Tell me where to look: Guided attention inference network; Proc IEEE/CVF Conf Comput Vis Pattern Recognit; 2018 Jun. 9215–9223.
- [27]. Wang, F; , et al. Residual attention network for image classification; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2017 Jul. 3156–3164.
- [28]. Fu, J; , et al. Dual attention network for scene segmentation; Proc IEEE/CVF Conf Comput Vis Pattern Recognit (CVPR); 2019 Jun. 3146–3154.

- [29]. Lu, J; Xiong, C; Parikh, D; Socher, R. Knowing when to look: Adaptive attention via a visual sentinel for image captioning; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2017 Jul. 375–383.
- [30]. Chen, L-C; Yang, Y; Wang, J; Xu, W; Yuille, AL. Attention to scale: Scale-aware semantic image segmentation; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2016 Jun. 3640–3649.
- [31]. Wang, X; Girshick, R; Gupta, A; He, K. Non-local neural networks; Proc IEEE/CVF Conf Comput Vis Pattern Recognit; 2018 Jun. 7794–7803.
- [32]. Woo S, Park J, Lee J-Y, So Kweon I. CBAM: Convolutional block attention module. Proc ECCV. 2018 Sep.:3–19.
- [33]. Duan Q, et al. SenseCare: A research platform for medical image informatics and interactive 3D visualization. arXiv:2004.07031. 2020
- [34]. Bi L, Kim J, Ahn E, Kumar A, Fulham M, Feng D. Dermoscopic image segmentation via multistage fully convolutional networks. IEEE Trans Biomed Eng. 2017 Sep; 64(9):2065–2074. [PubMed: 28600236]
- [35]. Yang, M; Yu, K; Zhang, C; Li, Z; Yang, K. DenseA-SPP for semantic segmentation in street scenes; Proc IEEE/CVF Conf Comput Vis Pattern Recognit; 2018 Jun. 3684–3692.
- [36]. Huang, G; Liu, Z; Van Der Maaten, L; Weinberger, KQ. Densely connected convolutional networks; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2017 Jul. 4700–4708.
- [37]. Lin, G; Milan, A; Shen, C; Reid, I. RefineNet: Multi-path refinement networks for high-resolution semantic segmentation; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2017 Jul. 1925–1934.
- [38]. He, K; Zhang, X; Ren, S; Sun, J. Deep residual learning for image recognition; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2016 Jun. 770–778.
- [39]. Chollet, F. Xception: Deep learning with depthwise separable convolutions; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2017 Jul. 1251–1258.
- [40]. Yu, F; Koltun, V; Funkhouser, T. Dilated residual networks; Proc IEEE Conf Comput Vis Pattern Recognit (CVPR); 2017 Jul. 472–480.
- [41]. Torrents-Barrena J, et al. Segmentation and classification in MRI and US fetal imaging: Recent trends and future prospects. Med Image Anal. 2019 Jan.51:61–88. [PubMed: 30390513]
- [42]. Chen, Y; Dai, X; Liu, M; Chen, D; Yuan, L; Liu, Z. Dynamic convolution: Attention over convolution kernels; Proc IEEE/CVF Conf Comput Vis Pattern Recognit (CVPR); 2020 Jun. 11030–11039.

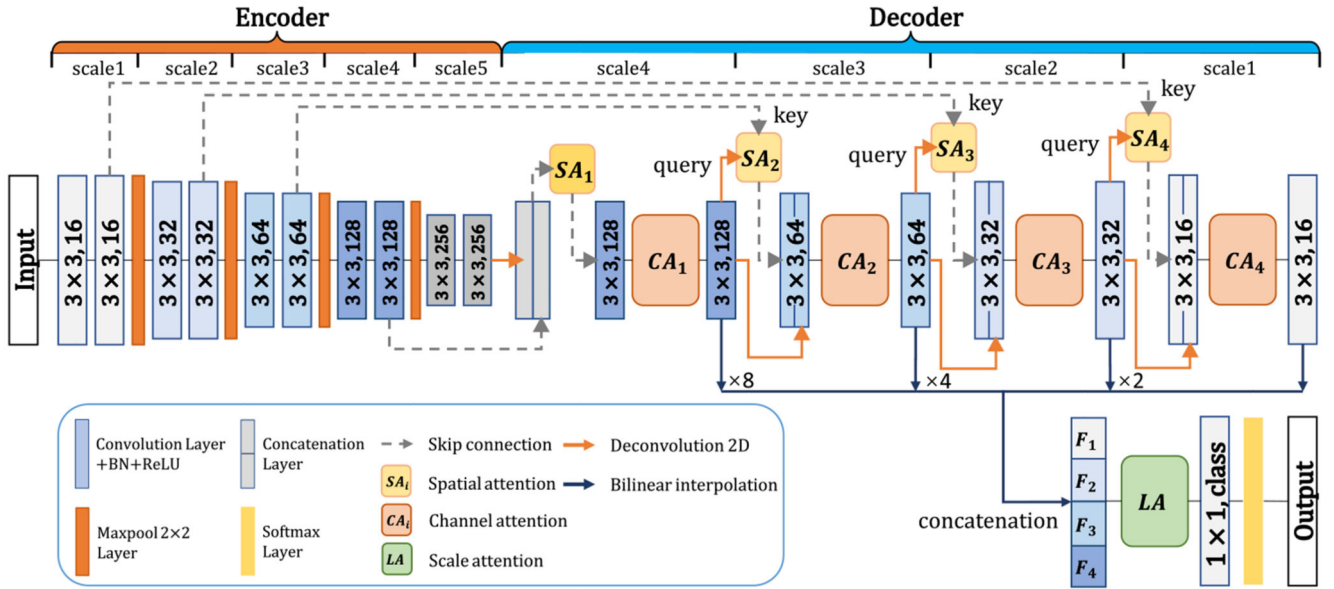


Fig. 1. Our proposed comprehensive attention CNN (CA-Net). Blue rectangles with 3×3 or 1×1 and numbers (16, 32, 64, 128, and 256, or class) correspond to the convolution kernel size and the output channels. We use four spatial attentions (SA_1 to SA_4), four channel attentions (CA_1 to CA_4) and one scale attention (LA). F_{1-4} means the resampled version of feature maps that are concatenated as input of the scale attention module.




Fig. 2. Details of our proposed joint spatial attention block.

(a) SA_1 is a non-local block used at the lowest resolution level. (b) single-pathway spatial attention block (SA). (c) SA_{2-4} are the dual-pathway attention blocks used in higher resolution levels. The *query* feature x^h is used to calibrate the low-level *key* feature x^l . δ means the *Sigmoid* function.

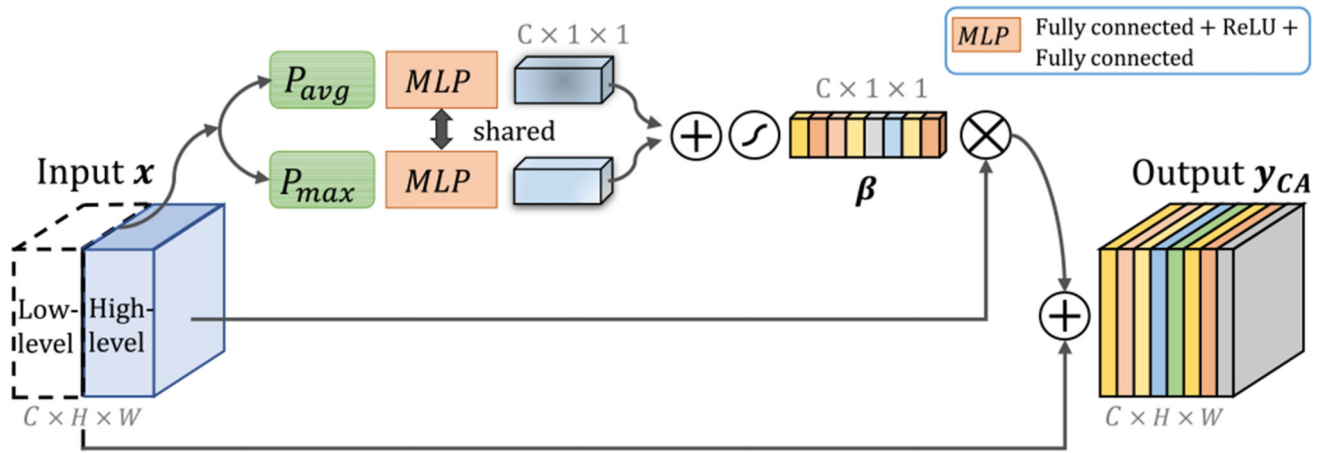


Fig. 3. Structure of our proposed channel attention module with residual connection. Additional global max-pooled features are used in our module. β means the channel attention coefficient.

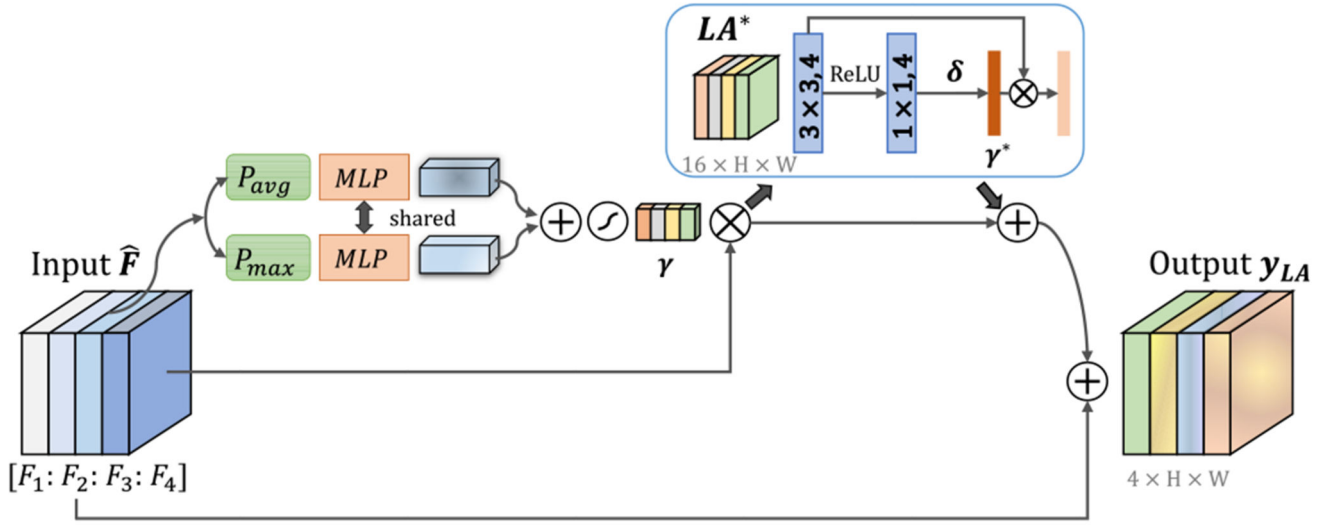


Fig. 4. Structure of our proposed scale attention module with residual connection. Its input is the concatenation of interpolated feature maps at different scales obtained in the decoder. γ means scale-wise attention coefficient. We additionally use a spatial attention block LA^* to gain pixel-wise scale attention coefficient γ^* .

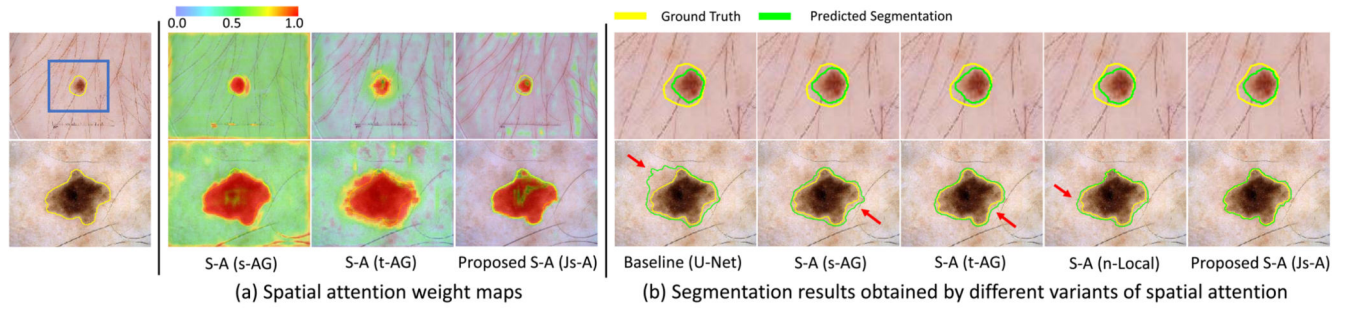
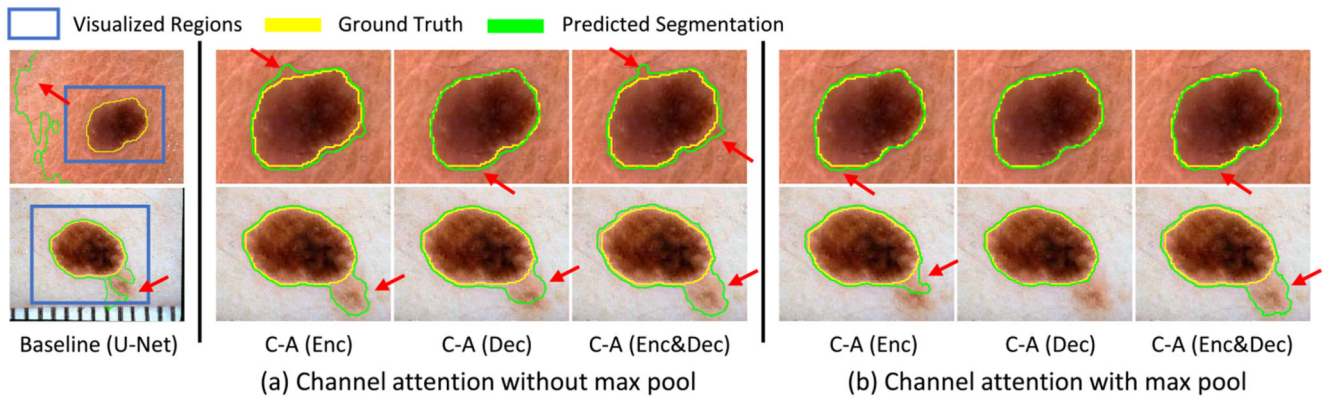


Fig. 5. Visual comparison between different spatial attention methods for skin lesion segmentation.

(a) is the visualized attention weight maps of single-pathway, dual-pathway and our proposed spatial attention. (b) shows segmentation results, where red arrows highlight some missegmentations. For better viewing of the segmentation boundary of the small target lesion, the first row of (b) shows the zoomed-in version of the region in the blue rectangle in (a).

**Fig. 6.**

Visual comparison of our proposed channel attention method with different variants. Our proposed attention block is CA (Dec) where the channel attention module uses an additional max pooling and is plugged in the decoder. The red arrows highlight some mis-segmentations.

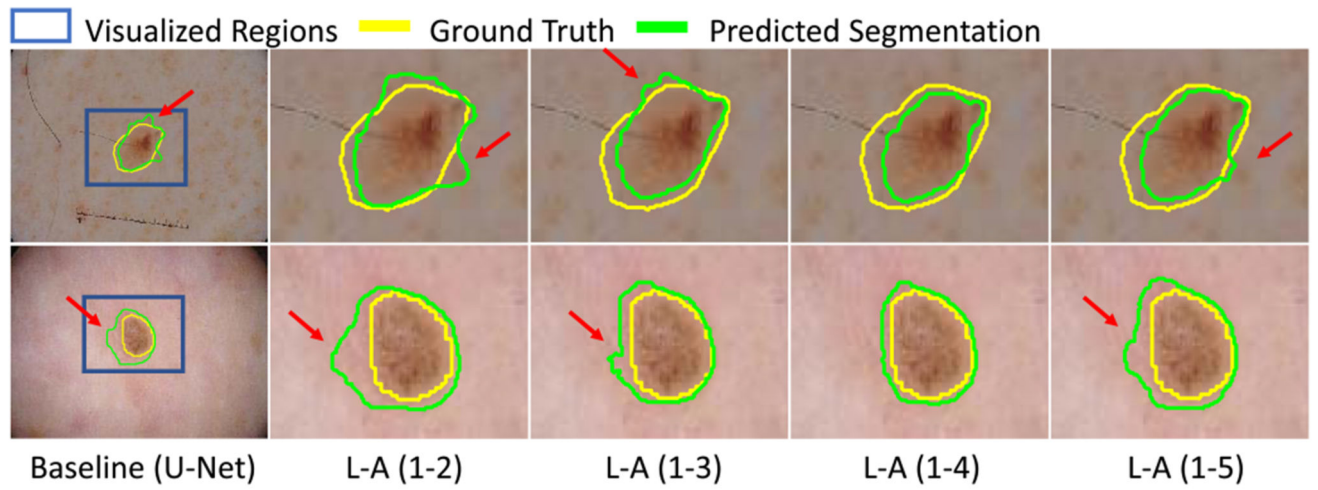


Fig. 7. Visual comparison of segmentation obtained by scale attention applied to concatenation of features from different scales.

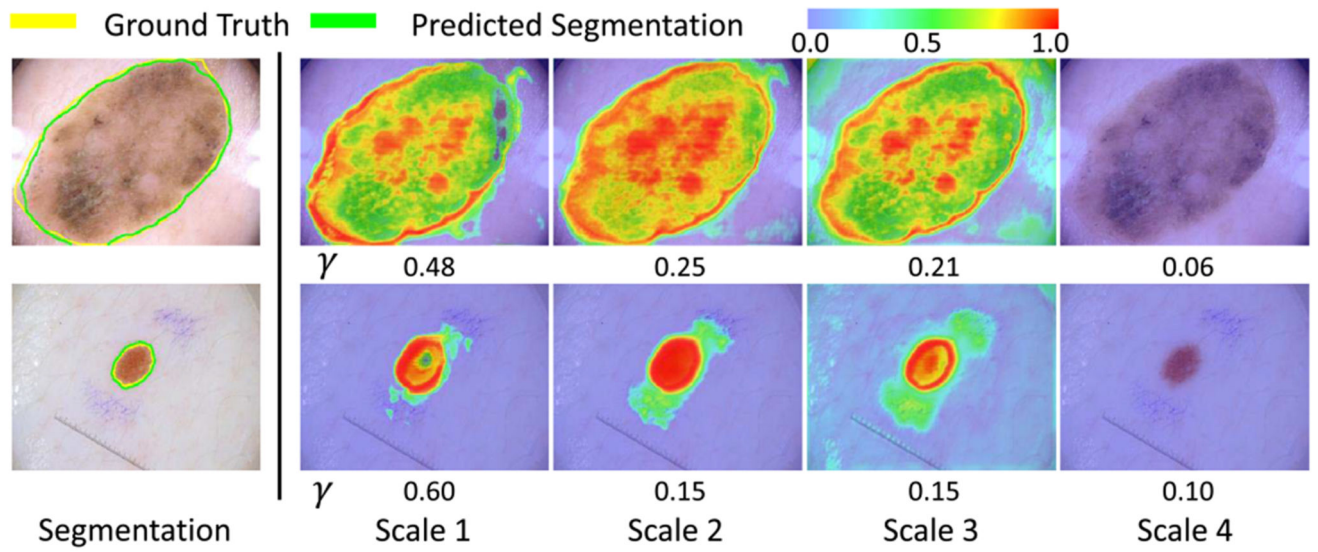


Fig. 8. Visualization of scale attention on dermoscopic images. warmer color represents higher attention coefficient values. γ means the global scale-wise attention coefficient.

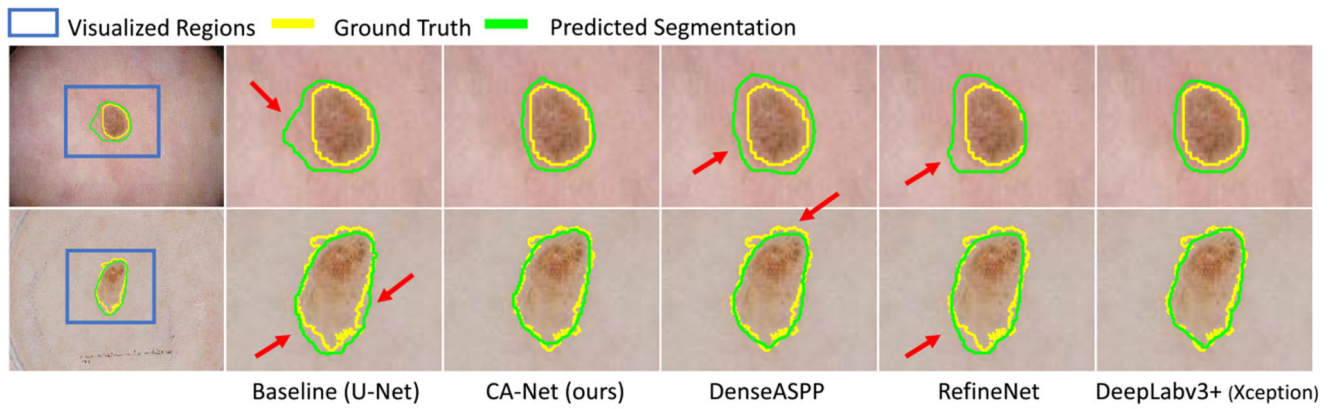


Fig. 9. Visual comparison between CA-Net and state-of-the-art networks for skin lesion segmentation. Red arrows highlight some mis-segmentation. DeepLabv3+ has similar performance to ours, but our CA-Net has fewer parameters and better explainability.

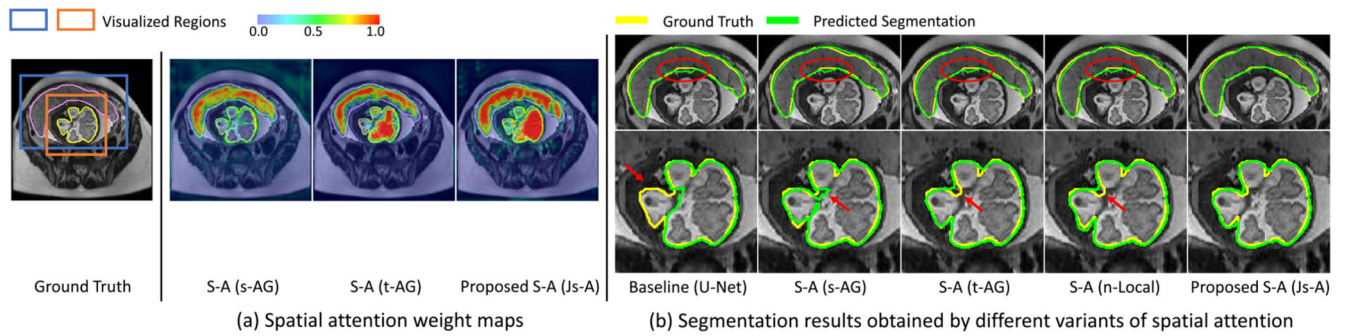


Fig. 10. Visual comparison between different spatial attention methods for fetal MRI segmentation.

(a) is the visualized attention weight maps of single-pathway, dual-pathway and proposed spatial attention. (b) shows segmentation results, where red arrows and circles highlight missegmentations.

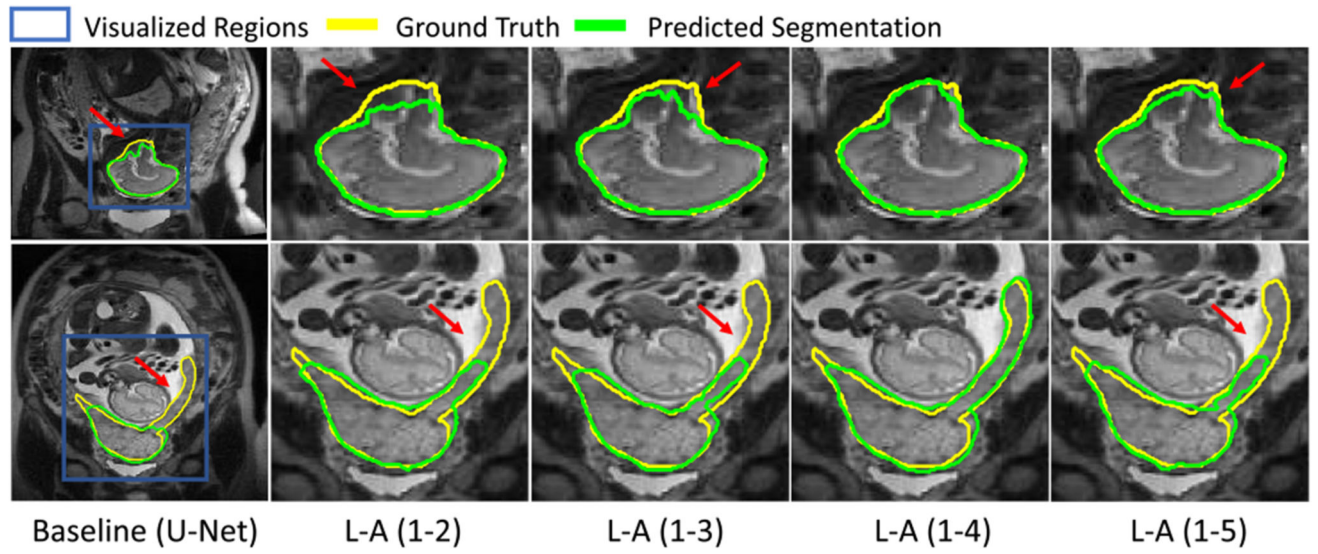


Fig. 11. Visual comparison of proposed scale attention method applied to concatenation of features form different scales.

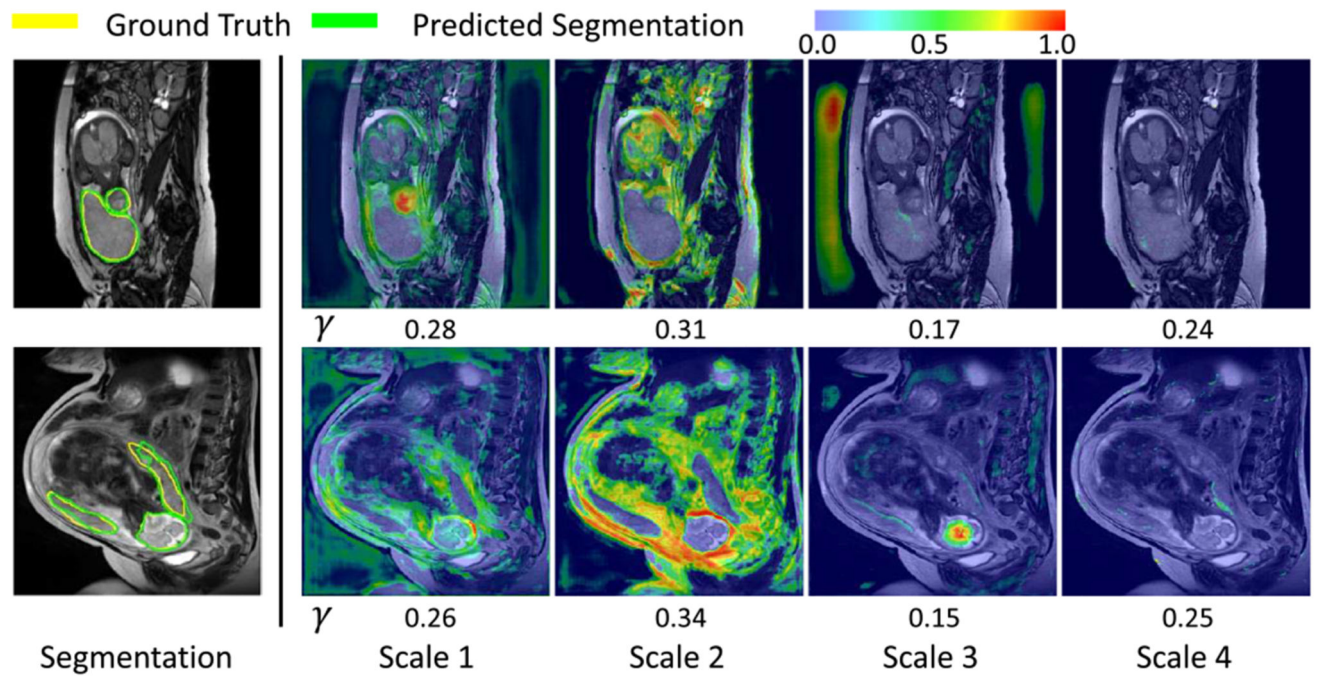


Fig. 12. Visualization of scale attention weight maps on fetal MRI. Warmer color represents higher attention coefficient values. γ means the global channel-wise scale attention coefficient.

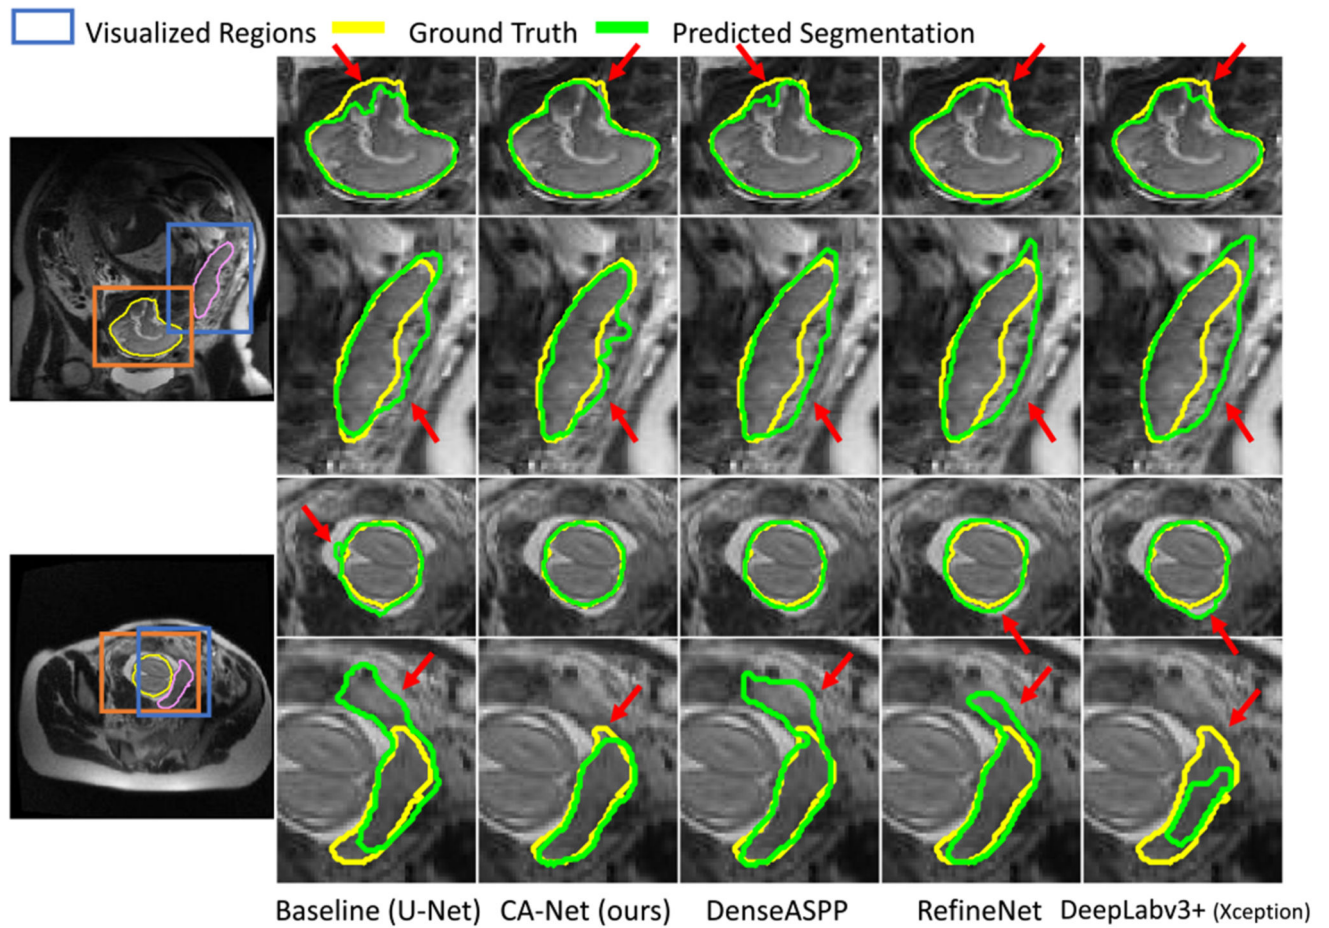


Fig. 13. Visual comparison of our proposed CA-Net with the state-of-the-art segmentation methods for fetal brain and placenta segmentation from MRI. Red arrows highlight the mis-segmented regions.

Table I

Quantitative Evaluation of Different Spatial Attention Methods for Skin Lesion Segmentation. (s-AG) Means Single-Pathway AG, (T-AG) Means Dual-Pathway AG, (n-Local) Means non-Local Networks. Js-A Is Our Proposed Multi-Scale Spatial Attention That Combines Non-Local Block And Dual-Pathway AG

| Network | Para | Dice(%) | ASSD(pix) |
|---------------------------|------|-------------------|------------------|
| Baseline(U-Net [6]) | 1.9M | 87.77±3.51 | 1.23±1.07 |
| S-A(s-AG) [9] | 2.1M | 88.46±3.37 | 1.18±1.24 |
| S-A(t-AG) | 2.4M | 89.18±3.29 | 0.90±0.50 |
| S-A(n-Local) | 1.9M | 90.15±3.21 | 0.65±0.72 |
| Proposed S-A(Js-A) | 2.0M | 90.83±3.31 | 0.81±1.06 |

Table II

Quantitative Comparison of Different Channel Attention Methods for Skin Lesion Segmentation. Enc, Dec and Enc&Dec Means Channel Attention Blocks Are Plugged in the Encoder, the Decoder and Both Encoder and Decoder, Respectively

| Network | P_{\max} | Para | Dice(%) | ASSD(pixel) |
|-----------------|------------|------|-------------------|------------------|
| Baseline | - | 1.9M | 87.77±3.51 | 1.23±1.07 |
| C-A(Enc) | × | 2.7M | 91.06±3.17 | 0.73±0.56 |
| C-A(Enc) | | 2.7M | 91.36±3.09 | 0.74±0.55 |
| C-A(Dec) | × | 2.7M | 91.56±3.17 | 0.64±0.45 |
| C-A(Dec) | | 2.7M | 91.68±2.97 | 0.65±0.54 |
| C-A(Enc&Dec) | × | 3.4M | 90.85±3.42 | 0.92±1.40 |
| C-A(Enc&Dec) | | 3.4M | 91.63±3.50 | 0.58±0.41 |

Table III

Quantitative Evaluation of Different Scale-Attention Methods for Skin Lesion Segmentation. L-A (1-K)
Represents the Features From Scale I to K Were Concatenated for Scale Attention

| Network | Para | Dice(%) | ASSD(pix) |
|-----------------|------|-------------------|------------------|
| Baseline | 1.9M | 87.77±3.51 | 1.23±1.07 |
| L-A(1-2) | 2.0M | 91.21±3.33 | 1.00±1.36 |
| L-A(1-3) | 2.0M | 91.53±2.52 | 0.70±0.61 |
| L-A(1-4) | 2.0M | 91.58±2.48 | 0.66±0.47 |
| L-A(1-5) | 2.0M | 89.67±3.40 | 0.82±0.50 |

Table IV

Comparison Between Partial and Comprehensive Attention Methods for Skin Lesion Segmentation. SA, CA and LA Represent Our Proposed Spatial, Channel and Scale Attention Modules Respectively

| Network | Para | Dice(%) | ASSD(pixel) |
|---------------------|------|-------------------|------------------|
| Baseline | 1.9M | 87.77±3.51 | 1.23±1.07 |
| SA | 2.0M | 90.83±3.31 | 0.81±1.06 |
| LA | 2.0M | 91.58±2.48 | 0.66±0.47 |
| CA | 2.7M | 91.68±2.97 | 0.65±0.54 |
| SA+LA | 2.1M | 91.62±3.13 | 0.70±0.48 |
| CA+LA | 2.7M | 91.75±2.87 | 0.67±0.48 |
| SA+CA | 2.8M | 91.87±3.00 | 0.73±0.69 |
| CA-Net(Ours) | 2.8M | 92.08±2.67 | 0.58±0.39 |

Table V

Comparison of the State-of-the-art Methods and Our Proposed CA-NET For Skin Lesion Segmentation. INF-T Means the Inference Time for a Single Image. E-Able Means the Method Is Explainable

| Network | Para/Inf-T | E-able | Dice(%) | ASSD(pix) |
|---------------------|-------------|--------|-------------|------------|
| Baseline(U-Net [6]) | 1.9M/1.7ms | × | 87.77±3.51 | 1.23±1.07 |
| Attention U-Net [9] | 2.1M/1.8ms | | 88.46±3.37 | 1.18±1.24 |
| DenseASPP [35] | 8.3M/4.2ms | × | 90.80±3.81 | 0.59±0.70 |
| DeepLabv3+(DRN) | 40.7M/2.2ms | × | 91.79±3.39 | 0.54±0.64 |
| RefineNet [37] | 46.3M/3.4ms | × | 91.55±2.11 | 0.64±0.77 |
| DeepLabv3+ [20] | 54.7M/4.0ms | × | 92.21 ±3.38 | 0.48 ±0.58 |
| CA-Net(Ours) | 2.8M/2.1ms | | 92.08±2.67 | 0.58±0.56 |

Table VI

Quantitative Evaluation of Different Spatial Attention Methods for Placenta and Fetal Brain Segmentation. s-AG Means Single-Pathway AG, T-AG Means Dual-Pathway AG, n-Local Means non-Local Networks. Js-A Is Our Proposed Multi-Scale Spatial Attention That Combines Non-Local Block and Dual-Pathway AG

| Network | Placenta | | Fetal Brain | |
|-------------|-------------------|------------------|-------------------|------------------|
| | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) |
| Baseline | 84.79±8.45 | 0.77±0.95 | 93.20±5.96 | 0.38±0.92 |
| s-AG [9] | 84.71±6.62 | 0.72±0.61 | 93.97±3.19 | 0.47±0.73 |
| t-AG | 85.26±6.81 | 0.71±0.70 | 94.70±3.63 | 0.30±0.40 |
| n-Local | 85.43±6.80 | 0.66±0.55 | 94.57±3.48 | 0.37±0.53 |
| Js-A | 85.65±6.19 | 0.58±0.43 | 95.47±2.43 | 0.30±0.49 |

Table VII

Comparison Experiment on Channel Attention-Based Networks For Fetal Mri Segmentation. ENC, DEC, AND Enc&Dec Means the Channel Attention Blocks Are Located in the Encoder, Decoder and Both Encoder and Decoder, Respectively

| Network | P_{max} | Para | Placenta | | Fetal Brain | |
|-----------------|-----------|------|-------------------|------------------|-------------------|------------------|
| | | | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) |
| Baseline | - | 1.9M | 84.79±8.45 | 0.77±0.95 | 93.20±5.96 | 0.38±0.92 |
| C-A(Enc) | × | 2.7M | 85.42±6.46 | 0.51±0.32 | 95.42±2.24 | 0.36±0.24 |
| C-A(Enc) | | 2.7M | 86.12±7.00 | 0.50±0.44 | 95.60±3.30 | 0.31±0.36 |
| C-A(Dec) | × | 2.7M | 86.17±5.70 | 0.44±0.29 | 95.61±3.69 | 0.33±0.42 |
| C-A(Dec) | | 2.7M | 86.65±5.99 | 0.52±0.40 | 95.69±2.66 | 0.28±0.39 |
| C-A(Enc&Dec) | × | 3.4M | 85.83±7.02 | 0.52±0.40 | 95.60±2.29 | 0.26±0.46 |
| C-A(Enc&Dec) | | 3.4M | 86.26±6.68 | 0.51±0.43 | 94.39±4.14 | 0.47±0.64 |

Table VIII

Comparison Between Different Variants of Scale Attention-Based Networks. L-A (1-K) Represents the Features From Scale 1 To K Were Concatenated For Scale Attention

| Network | Placenta | | Fetal Brain | |
|-----------------|-------------------|------------------|-------------------|------------------|
| | Dice(%) | ASSD(pix) | Dice(%) | ASSD(pix) |
| Baseline | 84.79±8.45 | 0.77±0.95 | 93.20±5.96 | 0.38±0.92 |
| L-A(1-2) | 86.17±6.02 | 0.59±0.41 | 94.19±3.29 | 0.27±0.36 |
| L-A(1-3) | 86.17±6.53 | 0.50±0.37 | 94.61±3.13 | 0.54±0.63 |
| L-A(1-4) | 86.21±5.96 | 0.52±0.58 | 95.18±3.22 | 0.27±0.59 |
| L-A(1-5) | 86.09±6.10 | 0.61±0.47 | 95.05±2.51 | 0.24±0.47 |

Table IX

Quantitative Comparison of Partial and Comprehensive Attention Methods for Fetal MRI Segmentation. SA, CA and LA Are Our Proposed Spatial, Channel And Scale Attention Modules Respectively

| Network | Placenta | | Fetal Brain | |
|---------------------|-------------------|------------------|-------------------|------------------|
| | Dice(%) | ASSD(pix) | Dice(%) | ASSD(pix) |
| Baseline | 84.79±8.45 | 0.77±0.95 | 93.20±5.96 | 0.38±0.92 |
| SA | 85.65±6.19 | 0.58±0.43 | 95.47±2.43 | 0.30±0.49 |
| LA | 86.21±5.96 | 0.52±0.58 | 95.18±3.22 | 0.27±0.59 |
| CA | 86.65±5.99 | 0.52±0.40 | 95.39±2.66 | 0.28±0.39 |
| SA+LA | 86.20±6.26 | 0.54±0.42 | 94.59±3.14 | 0.35±0.53 |
| CA+LA | 86.50±6.89 | 0.47±0.29 | 95.29±3.10 | 0.32±0.59 |
| SA+CA | 86.68±4.82 | 0.48±0.42 | 95.42±2.44 | 0.25±0.45 |
| CA-Net(Ours) | 87.08±6.07 | 0.52±0.39 | 95.88±2.07 | 0.16±0.26 |

Table X

Quantitative Evaluations of the State-of-the-art Methods and Our Proposed FA-Net for Fetal MRI Segmentation on Three Views (Axial, Coronal and Sagittal). Inf-T Means the Inference Time on Whole Test Dataset. E-Able Means the Method Is Explainable

| Network | Para/Inf-T | E-able | Placenta | | | | | | | | | | |
|---------------------|-------------|--------|-------------------|------------------|-------------------|------------------|-------------------|------------------|-------------------|------------------|---------|-------------|--|
| | | | Axial | | | Coronal | | | Sagittal | | | Whole | |
| | | | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | |
| Baseline(U-Net [6]) | 1.9M/0.9ms | × | 85.37±7.12 | 0.38±0.18 | 83.32±5.41 | 0.90±0.88 | 85.28±10.48 | 1.06±1.33 | 84.79±8.45 | 0.77±0.95 | | | |
| Attention U-Net [9] | 2.1M/1.1ms | | 86.00±5.94 | 0.44±0.29 | 84.56±7.47 | 0.71±0.45 | 85.74±10.20 | 1.00±1.27 | 85.52±7.89 | 0.72±0.83 | | | |
| DenseASPP [35] | 8.3M/3.1ms | × | 84.29±5.84 | 0.67±0.70 | 82.11±7.28 | 0.60±0.21 | 84.92±10.02 | 1.06±1.17 | 83.93±7.84 | 0.80±0.84 | | | |
| DeepLabv3+(DRN) | 40.7M/1.6ms | × | 85.79±6.00 | 0.60±0.56 | 82.34±6.69 | 0.73±0.46 | 85.98±7.11 | 0.90±0.58 | 84.91±6.59 | 0.75±0.54 | | | |
| RefineNet [37] | 46.3M/2.2ms | × | 86.30±5.72 | 0.55±0.48 | 83.25±5.55 | 0.64±0.40 | 86.67±7.04 | 0.88±0.79 | 85.60±6.18 | 0.70±0.60 | | | |
| DeepLabv3+ [20] | 54.7M/3.4ms | × | 86.34±6.30 | 0.46±0.50 | 83.38±7.23 | 0.54±0.37 | 86.90±7.84 | 0.58±0.49 | 85.76±7.17 | 0.57±0.45 | | | |
| CA-Net(Ours) | 2.8M/1.5ms | | 87.72±4.43 | 0.53±0.42 | 85.10±5.83 | 0.52±0.35 | 87.88±7.67 | 0.50±0.45 | 87.08±6.07 | 0.52±0.39 | | | |

| Network | Para/Inf-T | E-able | Brain | | | | | | | | | | |
|---------------------|-------------|--------|-------------------|------------------|-------------------|------------------|-------------------|------------------|-------------------|------------------|---------|-------------|--|
| | | | Axial | | | Coronal | | | Sagittal | | | Whole | |
| | | | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | Dice(%) | ASSD(pixel) | |
| Baseline(U-Net [6]) | 1.9M/0.9ms | × | 91.76±7.44 | 0.22±0.25 | 95.67±2.28 | 0.15±0.15 | 92.84±6.04 | 0.71±1.48 | 93.20±5.96 | 0.38±0.92 | | | |
| Attention U-Net [9] | 2.1M/1.1ms | | 93.27±3.21 | 0.20±0.16 | 94.79±5.49 | 0.57±1.36 | 94.85±3.73 | 0.19±0.21 | 94.25±4.03 | 0.30±0.70 | | | |
| DenseASPP [35] | 8.3M/3.1ms | × | 91.08±5.52 | 0.24±0.21 | 93.42±4.82 | 0.25±0.23 | 92.02±6.04 | 0.68±1.76 | 92.08±5.43 | 0.41±1.09 | | | |
| DeepLabv3+(DRN) | 40.7M/1.6ms | × | 92.63±3.00 | 0.33±0.42 | 94.94±2.49 | 0.14±0.09 | 92.36±0.09 | 0.94±2.75 | 93.10±5.90 | 0.51±1.70 | | | |
| RefineNet [37] | 46.3M/2.2ms | × | 94.04±1.85 | 0.25±0.45 | 91.90±11.10 | 0.25±0.41 | 90.36±10.91 | 0.81±2.09 | 92.05±8.77 | 0.46±1.32 | | | |
| DeepLabv3+ [20] | 54.7M/3.4ms | × | 93.70±4.17 | 0.19±0.31 | 95.50±2.60 | 0.09±0.03 | 94.97±4.27 | 0.58±1.24 | 94.65±3.87 | 0.31±0.78 | | | |
| CA-Net(Ours) | 2.8M/1.5ms | | 95.51±2.66 | 0.21±0.23 | 96.52±1.61 | 0.09±0.05 | 95.68±3.49 | 0.09±0.09 | 95.88±2.07 | 0.16±0.26 | | | |