# Thickened 2D Networks for 3D Medical Image Segmentation

Qihang Yu<sup>1</sup> Yingda Xia<sup>1</sup> Lingxi Xie<sup>2</sup> Elliot K. Fishman<sup>3</sup> Alan L. Yuille<sup>1</sup> <sup>1</sup> The Johns Hopkins University <sup>2</sup> Noah's Ark Lab, Huawei Inc. <sup>3</sup> The Johns Hopkins Medical Institute

{yucornetto, 198808xc, alan.l.yuille}@gmail.com yxia25@jhu.edu efishman@jhmi.edu

### Abstract

There has been a debate in medical image segmentation on whether to use 2D or 3D networks, where both pipelines have advantages and disadvantages. This paper presents a novel approach which thickens the input of a 2D network, so that the model is expected to enjoy both the stability and efficiency of 2D networks as well as the ability of 3D networks in modeling volumetric contexts. A major information loss happens when a large number of 2D slices are fused at the first convolutional layer, resulting in a relatively weak ability of the network in distinguishing the difference among slices. To alleviate this drawback, we propose an effective framework which (i) postpones slice fusion and (ii) adds highway connections from the pre-fusion layer so that the prediction layer receives slice-sensitive auxiliary cues. Experiments on segmenting a few abdominal targets in particular blood vessels which require strong 3D contexts demonstrate the effectiveness of our approach.

# 1. Introduction

Medical image segmentation is an important prerequisite of computer-assisted diagnosis (CAD) which implies a wide range of clinical applications. Recent years, with the development of deep learning, convolutional neural networks have been widely applied to this area, boosting the performance of previously designed handcrafted approaches by a large margin. On some easy organs and soft tissues, deep networks have almost achieved or even surpassed humanlevel accuracy. However, for some difficult targets in particular small blood vessels, segmentation accuracy is still below satisfactory.

The key reason that leads to this deficit lies in the major difference from medical images and natural images, *i.e.*, medical data are often in a 3D form. To deal with volumetric input, two main flowcharts exist. The first method borrows the idea from natural image segmentation, cutting the 3D volume into 2D slices, and train a 2D network which



Figure 1: An example for superior mesenteric artery (SMA) in the CT scan. SMA is small, and its shape is like a specially contiguous pipe, so it is difficult to segment when looked into from a single viewpoint. Upperleft: the axial view with the SMA area highlighted in red. Upperright: the label of SMA in a 3D view. Lowerleft and lowerright: segmentation results of the baseline with 3-slice inputs and the proposed model with 15-slice inputs (with DSC provided). The thicker inputs lead to much better segmentation.

deals with each slice individually or sequentially. Another way is to train a 3D network to deal with volumetric data directly. Both of these two methods have their own advantages and disadvantages. A 2D network makes use of information in the entire 2D slice, but suffers from the lack of information of relationship among slices (see Figure 1). A 3D approach, in opposite, has the perception of 3D contexts, but also suffers from two weaknesses: (i) the lack of pre-trained models which makes the training process unstable; and (ii) in the testing stage, fusing patch-wise output into the final prediction remains tricky and time-consuming. This paper presents a novel framework which **uses a 2D network to mimic 3D segmentation**. Our idea is motivated by a few prior work for pseudo-3D segmentation [30][27], in which three neighboring 2D slices are stacked during training and testing, so that the 2D network sees a small range of 3D contexts. However, these approaches still produced unsatisfying results for the blood vessels, because recognizing these targets require even stronger 3D contexts – see Figure 1 for an example. To deal with this problem, a natural choice is to continue thickening the input of 2D networks, allowing richer 3D contexts to be seen. However, this strategy reaches a plateau quickly at 5 or 6 slices, and adding more slices to 2D networks causes accuracy drop.

The essence lies in information loss. Technically, for a 2D network, after the first convolutional layer, information from all input slices are mixed together, and, without a specifically designed scheme, it is difficult to discriminate them from each other. For a regular organ like pancreas, this does not bring in too much trouble (using 3 slices and 12 slices as input produce comparable accuracy), but for a blood vessel like superior mesenteric artery (SMA), this can cause significant accuracy drop (from 73.32% to 71.07%).

We argue that it is the early fusion (information from all slices are fused in the first convolution) that leads to the information loss. Therefore, our solution is to postpone the stage that these information are put together. Our algorithm can be built upon an arbitrary segmentation baseline. It is composed of two parts. First, the thickened input is cut into several groups, each of which has a smaller number, say 3, of slices. We partition the baseline network into two parts, with the first part being applied to each group individually, and the second part processing the re-grouped output of the first part. With this design, we postpone the stage that information along the third dimension is fused, and thus alleviate the extent of information loss. In addition, to improve the discriminative power of the fused features, we insert a highway connection between the pre-fusion stage and the decision stage, *i.e.*, before the up-sampling operations. The options of fusing these two sources of information are also studied in an empirical manner.

Experiments are performed on several abdominal organs individually, including two regular organs and three blood vessels in our own dataset, and the hepatic vessels in the Medcial Segmentation Decathlon (MSD) dataset. Our approach achieves consistent accuracy gain in particular in the blood vessels, demonstrating that 3D contexts are indeed incorporated into 2D networks.

The remainder of this paper is organized as follows. Section 2 briefly reviews related work, and Section 3 presents our approach. After experiments are shown in Section 4, we conclude this work in Section 5.

### 2. Related Work

**Semantic Segmentation** Semantic segmentation is a critical problem in computer vision. Conventional methods based on graph [1] or handcrafted local features [24] have been gradually replaced by techniques from deep learning, which could produce higher segmentation accuracy [15, 4]. Various deep network architectures have been proposed and achieved great performance on large natural image datasets [11]. As the segmentation networks could be extended to more and more tasks, researchers also attempt to apply segmentation network to medical imaging analysis, which is usually based on volumetric data [27, 30, 19].

**Computer aided diagnosis** Computer aided diagnosis (CAD) is a research area aiming at helping human doctors in clinics. Recently a lot of CAD approaches are based on medical imaging analysis to get accurate descriptions of the scanned organs, soft tissues, *etc.*. One topic with great importance in this area is object segmentation, *i.e.*, determining which voxels belong to the target in 3D data. The progress of deep learning brought more powerful and efficient solutions. It has been proved successfully not only in natural image area [15, 5] but also medical image area [18, 17], outperforming conventional approaches, *e.g.*, when segmenting the liver [13], the lung [10], or the pancreas [6, 20]. These deep learning methods in medical image segmentation can be classified into two main types according to their way to deal with 3D volume data:

**2D Networks** Researchers taking this method usually cut each 3D volume into 2D slices, and train a 2D network to process each of them individually [18]. Such method often suffers from missing 3D contextual information, for which various techniques were adopted, such as using 2.5D data (stacking a few 2D images as different input channels) [19, 20], training deep networks from different viewpoints and fusing multi-view information at the final stage [30, 26, 27], and applying a recurrent network to process sequential data [2, 3].

**3D Networks** In this method, researchers directly train a 3D network to deal with the volumetric data [7, 17]. These approaches, while being able to see more information, often require much larger memory consumption, so some existing methods worked on small patches [8, 31] and fused the outputs of all patches. In addition, unlike 2D networks that can borrow pre-trained models from natural image datasets [11], 3D networks need to be trained from scratch, which means it could suffer from unstable convergence properties [21]. One way is to implement a pseudo 3D convolution by a 2D-followed-by-1D convolution [14]. A discussion on 2D vs. 3D models for medical imaging segmentation is available in [12].



Figure 2: Network Architecture. Left one is the flowchart of our model, the network (in this paper it refers to a ResNet50) is divided into two parts. By multiplexing the first part of network, little extra parameters are added. The right one is the design of a channel-wise non-local module. It would compute the relationship between channels of final feature map and channels of single slice group feature map, thus leading to a slice-sensitive result. The right part illustrates the detail of Channel Non-local in the left part.

### 3. Our Approach

#### 3.1. Problem Statement

In our task, a CT scan can be regarded as a 3D volume **X** of size  $W \times H \times L$  which is annotated with a binary ground-truth segmentation **Y** where  $y_i = 1$  means a foreground voxel. We need to obtain a binary output volume **Z** of the same dimension. Denote  $\mathcal{Y}$  a and  $\mathcal{Z}$  as the set of foreground voxels in the ground truth and prediction. *i.e.*,  $\mathcal{Y} = \{ i \mid y_i = 1 \}$  and  $\mathcal{Z} = \{ i \mid z_i = 1 \}$ .

# 3.2. Common 2D Segmentation Network

Suppose our input is a 3D volume **X** of size  $W \times H \times L$ . When training 2D deep networks for 3D segmentation, each 3D volume is sliced along three axes, the *coronal*, *sagittal*, and *axial*. We denote these 2D slices with  $\mathbf{X}_{C,w}$  (w = 1, 2, ..., W),  $\mathbf{X}_{S,h}$  (h = 1, 2, ..., H),  $\mathbf{X}_{A,l}$ (l = 1, 2, ..., L), where the subscripts C, S and A stand for *coronal*, *sagittal*, and *axial*, respectively. On each axis, an individual 2D deep network like UNet [18] or Deeplab [4] is trained, which are denoted by  $\mathbb{M}_C$ ,  $\mathbb{M}_S$ , and  $\mathbb{M}_A$ , respectively. Without loss of generality, we consider a 2D slice along the axial view, denoted by  $\mathbf{X}_{A,l}$ . Our goal is to infer a binary segmentation mask  $\mathbf{Z}_{A,l}$  of the same dimensionality. Usually, this is achieved by first computing a probability map  $\mathbf{P}_{A,l} = \mathbf{f}[\mathbf{X}_{A,l};\theta]$ , where  $\mathbf{f}[\cdot;\theta]$  is a deep segmentation network with  $\theta$  being network parameters, and afterwards,  $\mathbf{P}_{A,l}$  is binarized into  $\mathbf{Z}_{A,l}$  using a fixed threshold of 0.5, *i.e.*,  $\mathbf{Z}_{A,l} = \mathbb{I}[\mathbf{P}_{A,l} \ge 0.5]$ , and final axial prediction needs to be grouped  $\mathbf{Z}_A = [\mathbf{Z}_{A,1}, \dots, \mathbf{Z}_{A,L}]$ . The final prediction are averaged probability on three axes  $\mathbf{P} = \frac{(\mathbf{P}_C + \mathbf{P}_S + \mathbf{P}_A)}{3}$ , and  $\mathbf{Z} = \mathbb{I}[\mathbf{P} \ge 0.5]$ .

# 3.3. Thickened 2D Segmentation Network

#### 3.3.1 Thickened 2D Inputs and Information Loss

Thickened 2D Inputs for Pseudo-3D segmentation A typical method to improve 2D network performance on 3D data is to use three neighboring slices as input instead of replicating one slice for three times. This brings slight 3D information and enables the 2D network to perform a pseudo-3D segmentation, which means 2D segmentation that incorporates 3D information. Segmentation on a 3D volume data mainly relies on two kinds of information, one is called intra-slice information, which is what we can know about a slice from itself. Another is called inter-slice information, which is what we can know about a slice from its neighboring slices. Motivated by this, we propose to thicken 2D inputs. By adding more inter-slice information, we believe that it could help improve the pseudo-3D segmentation one step forward, especially for blood vessels like SMA which features in its contiguous shape along some axis 1. With thickened 2D inputs, the segmentation is not based on a single slice  $\mathbf{X}_{A,l}$  but on a k-slice group

 $\mathbf{X}_{A,l}^{k} = [\mathbf{X}_{A,l}, \mathbf{X}_{A,l+1}, \dots, \mathbf{X}_{A,l+k-1}].$  The model takes a k-slice group as input and also outputs a corresponding k-slice prediction. And the formulation becomes  $\mathbf{P}_{A,l}^{k} =$  $\mathbf{f}[\mathbf{X}_{A,l}^{k}; \theta]$ , where  $\mathbf{P}_{A,l}^{k} = [\mathbf{P}_{A,l}, \mathbf{P}_{A,l+1}, \dots, \mathbf{P}_{A,l+k-1}],$ and  $\mathbf{Z}_{A,l}^{k} = \mathbb{I}[\mathbf{P}_{A,l}^{k} \geq 0.5].$ 

**Regroup Thickened 2D Outputs** When grouping thickened 2D outputs into a 3D volume, there can exist overlapped slices, so we need a function G to combine all predictions into one volume, thus we have  $\mathbf{Z}_{A} = G[\mathbf{Z}_{A,1}^{k}, \mathbf{Z}_{A,2}^{k}, \dots, \mathbf{Z}_{A,L}^{k}]$ , in our paper, function G is averaging the probability of overlapped slices and concatenate each slice prediction into one volume.

**Information Loss** Despite that increasing slice thickness from 1 to 3 can boost performance, we observe that increasing thickness to a larger number can result in performance drop. For instance, in our experiments 3, using 12-slice input instead of 3-slice on SMA leads to a drop from 73.32% to 71.07%. We notice that the reason is the information loss brought by 2D convolution operation which fuses different slices. As we know, a 2D convolution can be regarded as weighted-sum of all input channels for each output channel, so there is no special connection between specific output channel and corresponding input channel, leading to a confusion of intra-slice information and inter-slice information. Say  $\mathbf{R}^{C}$  is a feature map with C channels, in a typical ResNet network, the feature map mapping is  $\mathbf{R}^k \to \mathbf{R}^{64} \to \cdots \to \mathbf{R}^{2048} \to \mathbf{R}^{256} \to \mathbf{R}^k$ . For the input  $\mathbf{R}^k$  and output  $\mathbf{R}^k$ , their channels should have one-toone relationship in a k-slice medical image segmentation. When k is small, say 1 or 3, the network could figure out the mapping relations. But when K is large, say 12, the network could be confused after so many 2D convolutions which result in a loss of slice-sensitive information, and this information loss leads to a worse result.

To address this problem, we propose two methods to reduce the information loss when fusing slices.

#### 3.3.2 Solution One: Postpone Fusion Stage

To reduce the information loss, we use multiple smallthickness groups instead of a whole large-thickness group, and we also postpone the fusion stage. If every slice is encoded into same feature space before fusion, then the fusion stage will result in fewer loss. If we fuse each mini-group at the very last layer, there will be no confusion at all, since all slices individually go through the whole network except last layer. However, if the fusion stage is too late, we cannot make full use of inter-slice information with limited computation. So we need to find a balance point where the loss is little and the model could benefit from fused feature most. To achieve this, we divide the original 2D network  $\mathbf{f}[\cdot; \theta]$ into two parts,  $\mathbf{f}_1[\cdot; \theta_1]$  and  $\mathbf{f}_2[\cdot; \theta_2]$ . In the first part of network, each mini-group will forward propagate individually, and get fused as one unity in second part. Thus a segmentation procedure on a k-slice group is:

$$\begin{cases} \mathbf{X}_{A,l}^{k} = [\mathbf{X}_{A,l}, \mathbf{X}_{A,l+1}, \dots, \mathbf{X}_{A,l+k-1}] \\ \mathbf{O}_{A,l} = \mathbf{f}_{1}[\mathbf{X}_{A,l}; \theta_{1}] \\ \mathbf{O}_{A,l}^{k} = \mathbf{f}_{c}[\mathbf{O}_{A,l}, \mathbf{O}_{A,l+1}, \dots, \mathbf{O}_{A,l+k-1}] \\ \mathbf{Z}_{A,l}^{k} = \mathbf{f}_{2}[\mathbf{O}_{A,l}^{k}; \theta_{2}] \\ \mathbf{Z}_{A} = G[\mathbf{Z}_{A,1}^{k}, \mathbf{Z}_{A,2}^{k}, \dots, \mathbf{Z}_{A,L}^{k}] \end{cases}$$
(1)

where  $\mathbf{O}_{A,l}$  is an intermediate feature map for slice l at prefusion layer, and  $\mathbf{O}_{A,l}^k = [\mathbf{O}_{A,l}, \mathbf{O}_{A,l+1}, \dots, \mathbf{O}_{A,l+k-1}]$ .  $\mathbf{f}_c$  is a convolution to fuse features from different slices.  $\mathbf{f}_1$ ,  $\mathbf{f}_2$  are two parts of a complete network  $\mathbf{f}$ . G is a function grouping all outputs into one volume, in this paper, G concatenates all output slices and average overlapped slices.

As we show in ablation study Table 3, this modification alone can help the network better figure our the corresponding relationship between input slices and output slices, thus leading to a better result.

#### 3.3.3 Solution Two: Feature Separable Attention

We further address the problem of information loss. Though postpone fusion stage can help reduce information loss and make it easier for network to distinguish the output for each slice, there still exists some problem. The final feature map is a mixture of inter-slice information and intra-slice information, and the network needs to predict for different slice based on this same mixed feature map, which would be hard if no specific-slice information is given. So we introduce slice-sensitive auxiliary cues from pre-fusion layer to the final feature map. We add a highway connection from prefusion layer to prediction layer.

Channel-wise Non-local Module Non-local module could serve as a good way for attention [22] or different kinds of feature fusion [25]. Inspired by non-local neural networks [23] and its following work [28], we also apply non-local module as a way to introduce slice-sensitive information. Yet in our situation, the main problem lies in confusion in channel level (see Section 3.3.1). So we design a channel-wise non-local module to perform the interaction between the slice-specific feature and final mixed feature. In our work, we propose channel-wise non-local module(CNLM) to further address the problem of information loss. Suppose the input feature map has a shape of (C, H, W). Instead of computing the relationship between each location pair, we compute the correspondence between each pair of feature dimension (channel), which can be formulated as:

$$\mathbf{y}_p = \frac{1}{Z(x)} \sum_{\forall q} f(x_p, x_q) g(x_q) \tag{2}$$

where  $p, q \in \{1, ..., C\}$ ,  $x_p, x_q$  represent  $H \times W$  dimensional vectors for the  $p^{th}$  and  $q^{th}$  channel. A pairwise function f computes a scalar between channel p and all channel q. The unary function g computes a representation of the input x at the channel q. This module intuitively tells where to look at for each dimension of high level feature, and thus separate information for each slice from the mixed-up information during computation.

After equipped with channel-wise non-local module, the segmentation procedure becomes:

$$\begin{cases} \mathbf{X}_{A,l}^{k} = [\mathbf{X}_{A,l}, \mathbf{X}_{A,l+1}, \dots, \mathbf{X}_{A,l+k-1}] \\ \mathbf{O}_{A,l} = \mathbf{f}_{1}[\mathbf{X}_{A,l}; \theta_{1}] \\ \mathbf{O}_{A,l}^{k} = \mathbf{f}_{c}[\mathbf{O}_{A,l}, \mathbf{O}_{A,l+1}, \dots, \mathbf{O}_{A,l+k-1}] \\ \mathbf{U}_{A,l}^{k} = \mathbf{f}_{2}[\mathbf{O}_{A,l}^{k}; \theta_{2}] \\ \mathbf{Z}_{A,l} = \text{CNLM}[\mathbf{U}_{A,l}^{k}, \mathbf{O}_{A,l}; \eta] \\ \mathbf{Z}_{A} = G[\mathbf{Z}_{A,1}^{k}, \mathbf{Z}_{A,2}^{k}, \dots, \mathbf{Z}_{A,L}^{k}] \end{cases}$$
(3)

where CNLM is channel-wise non-local module with parameters  $\eta$ .  $\mathbf{U}_{A,l}^k$  is the feature for all k slices, which is the output of second part of network.

#### **3.4. Implementation Details**

**Backbone** We follow the design of Deeplab [4] to modify a ResNet50 [9] into segmentation network with decoder. We divide the input slices into 3-slice groups, which means there are 3 slices in each group. Given that a 2D network like ResNet50 can handle 3 slices well, thus by 3-slicemini-grouping we can use a much larger slice thickness under same GPU memory limitation. To get meaningful intra-slice information, each 3-slice group will go through the first convolution and the following layer1 part of the network. Afterwards, we concatenate the feature from different groups on channel axis and use two convolutions to compress it, one will compress the channel number to its half and another to 256. After going through the remaining part of the network as one unity, we get the mixture of intra-slice and inter-slice information we want. To extract the targeted 3-slice features, we use channel-wise non-local module to deal with specific-slice information and the allslice information. Thus we get a prediction focusing on the targeted slices. The architecture is shown in left part of Figure 2.

**Postpone Fusion Time** As we mentioned in Section 3.3.2, the later we fuse all this intra-slice feature, the more distinguishable features we would get. This serves as an initial encoding of different slices so that fusing them will be easier and better. We found that fusion after layer1 is good enough and could save more memory and computation resources compared with later stages. Yet we also be-

Algorithm 1 The Testing Phase of Thickened 2D Networks Input:

input volume **X**, viewpoint  $\mathcal{V} = \{C, S, A\}$ ; network parameters  $\theta^v$  which consists of  $\theta_1^v$  and  $\theta_2^v$ , and channel non-local module parameters  $\eta^v, v \in \mathcal{V}$ ; slice thickness k, threshold thr; **Output:** 

### segmentation volume **Z**;

1:  $\mathbf{I}^{v} \leftarrow \mathbf{X}, v \in \mathcal{V};$ 2:  $\mathbf{I}_{1}^{v}, \mathbf{I}_{2}^{v}, \dots, \mathbf{I}_{k}^{v} \leftarrow \mathbf{I}^{v}$ 3: for i in  $1, 2, \dots, k$ : 4:  $\mathbf{O}_{i}^{v} \leftarrow \mathbf{f}_{1}[\mathbf{I}_{i}^{v}; \theta_{1}^{v}]$ 5:  $\mathbf{O}^{v} \leftarrow \mathbf{f}_{c}[\mathbf{O}_{1}^{v}, \dots, \mathbf{O}_{k}^{v}]$ 6:  $\mathbf{U}^{v} \leftarrow \mathbf{f}_{2}[\mathbf{O}^{v}; \theta_{2}^{v}]$ 7: for i in  $1, 2, \dots, k$ : 8:  $\mathbf{P}_{i}^{v} \leftarrow \mathbf{CNLM}[\mathbf{U}^{v}, \mathbf{O}_{i}^{v}; \eta^{v}]$ 9:  $\mathbf{P}^{v} \leftarrow [\mathbf{P}_{1}^{v}, \dots, \mathbf{P}_{k}^{v}]$ 10:  $\mathbf{P} = \frac{\mathbf{P}^{C} + \mathbf{P}^{S} + \mathbf{P}^{A}}{3}, \mathbf{Z} = \mathbb{I}[\mathbf{P} \ge 0.5]$ 11: return  $\mathbf{Z}$ 

lieve that we need to postpone the fusion stage furthermore if we want to thicken the input slices a step forward.

**Channel-wise Non-local Module** Different from the original non-local network [23], here the confusion mainly comes from channel-level instead of pixel-level, so we change the transpose and matrix multiplication operations to fit the channel-level situation. Besides, we also replace two of the three linear part ( $1 \times 1$  convolution followed by a max pooling) with an adaptive average pooling to  $32 \times 32$  size followed by a  $1 \times 1$  convolution as scaling part. The average pooling and convolution are meant to reduce the variance of different receptive field. We also add a switchable normalization [16] to reinforce the module. Channel-wise non-local module here serves as an attention module and aims at introducing the targeted individual slice information in the fused feature map.

**Training** We train our models using SGD with a minibatch size of 8 on 4 GPUs (*i.e.*, 2 samples per GPU). The model is trained for 100k iterations, with a learning rate of 0.005, which is decreased by a factor of 10 at iteration 70k and 90k. The SGD momentum is 0.9 and weight decay 0.0005. Besides, we use the ImageNet pretrained ResNet50 weight as our model's initialization. To align different slice size so that the model could be trained parallel, images are cropped or zero padded to make sure that each slice size should be  $512 \times 512$ . The model on each axis is trained separately. The loss term we use is the DSC loss proposed in VNet paper [17]. We represent each probability map with  $P_{A,l}^K$ , and denote the loss by  $L\{\mathbf{Y}_{A,l}^K, \mathbf{P}_{A,l}^K\}$ . Here  $\mathbf{Y}_{A,l}^K$  is the ground-truth segmenta-

SMA	3-slice baseline	VNet	VNet (ROI)	15-slice ours	
$DSC_C$	71.20%	-	-	73.06%	
DSC <sub>S</sub>	70.71%	-	-	72.97%	
$DSC_A$	69.62%	-	-	72.64%	
$DSC_F$	73.32%	70.36%	73.49%	74.27%	
Celiac AA					
$DSC_C$	58.78%	-	-	60.61%	
$DSC_S$	52.61%	-	-	53.09%	
DSC <sub>A</sub>	58.06%	-	-	59.98%	
$DSC_F$	59.97%	54.38%	63.38%	60.38%	
Duodenum					
$DSC_C$	67.78%	-	-	70.16%	
$DSC_S$	64.43%	-	-	68.31%	
$DSC_A$	68.87%	-	-	70.26%	
$DSC_F$	72.55%	60.62%	71.43%	73.67%	
Pancreas					
$DSC_C$	85.43%	-	-	86.53%	
$DSC_S$	85.30%	-	-	86.34%	
DSC <sub>A</sub>	84.71%	-	-	86.29%	
$DSC_F$	87.15%	84.12%	86.17%	87.36%	
Vein					
$DSC_C$	74.01%	-	-	75.45%	
DSC <sub>S</sub>	73.11%	-	-	74.64%	
DSC <sub>A</sub>	73.35%	-	-	75.39%	
$DSC_F$	75.77%	75.06%	76.90%	76.27%	

Table 1: Experiments on SMA, celiac aa, pancreas, duodenum, and vein.  $DSC_C$  means dice score based on coronal plane,  $DSC_S$  is on sagittal plane,  $DSC_A$  refers to axial plane, and  $DSC_F$  is three-axes fusion result. We evaluate the final result with DSC and compare our approach with 2D baseline and 3D network. ROI means the 3D network only deal with data inside ground-truth bounding-box even when testing.

tion mask, and  $L\{\mathbf{Y}, \mathbf{P}\} = 1 - \frac{2 \times \sum_i \mathbf{Y}_i \mathbf{P}_i}{\sum_i \mathbf{Y}_i + \mathbf{P}_i}$ . The whole training procedure takes 14 hours on 4 NVIDIA TITAN Xp GPUs.

**Testing** We slice the volume into 15-slice groups while each group has 14 slices overlapped by each other. The final prediction is averaged on each slice and each viewpoint. The output ranges in [0,1] and we use a threshold = 0.5 to get final prediction.

### 4. Experiments

#### 4.1. Dataset and Evaluation

To verify that our approach can be applied to various blood vessels and organs, we collect a large dataset which contains 200 CT scans. This corpus took 4 full-time radiologists around 3 months to annotate. To the best of our knowledge, this dataset is large and contains more organs than any public datasets. We choose several blood vessels which require more inter-slice information and also other

Hepathic Vessel	3-slice baseline	VNet	VNet(ROI)	15-slice ours
$DSC_C$	56.19%	-	-	56.23%
$DSC_S$	55.19%	-	-	56.80%
DSC <sub>A</sub>	54.63%	-	-	58.33%
$DSC_F$	58.09%	51.68%	51.68%	58.86%

Table 2: Experiments on the MSD dataset for hepatic vessel segmentation.

challenging organs like pancreas and duodenum. We randomly partition the dataset into two folds, one contains 150 cases for training, while another consisting of 50 cases for testing. Each organ is trained and tested individually. When a pixel is predicted as more than one organ, we choose the one with the largest confidence score.

We also conduct experiments on a public dataset from Medical Segmentation Decalthon which contains 303 cases of Hepatic Vessels. We use 228 cases for training and 75 cases for testing. The accuracy of segmentation is evaluated by the Dice-Sørensen coefficient (DSC): DSC  $(\mathcal{Y}, \mathcal{Z}) = \frac{2 \times |\mathcal{Y} \cap \mathcal{Z}|}{|\mathcal{Y}| + |\mathcal{Z}|}$ . This metric falls in the range of [0,1], with 1 implying perfect segmentation.

#### 4.2. Experiment results on Multi-organ Dataset

Results are summarized in Table 1. We have conducted experiments on several challenging blood vessels and organs, including SMA, celiac aa, vein and pancreas, duodenum. On blood vessels, single slice based network usually cannot produce satisfying results, and inter-slice information plays an important role here. We also try to apply our algorithm to other small and challenging organs, proving that our method also works for other organs besides blood vessels. With our method, the inter-slice information and intra-slice information could collaborate in a better way. We notice that there is a significant improvement in terms of single axis performance, even comparable with a fusion result. We also compare the result with a patch-based 3D network VNet [17] to verify our model can make a good use of 3D information. As we show, due to lack of intra-slice information, a 3D network suffers from serious false positive especially on small blood vessels which only take up thousands of voxels. Our results are also comparable with a VNet tested with ROI (which means the model only deal with volumes inside the ground truth bounding box, so there is no false positive). Our model can benefit from both 2D network advantage and 3D network advantage.

#### 4.3. Experiment results on MSD Dataset

We further apply our approach on a public dataset – Hepatic Vessel segmentation in Medical Segmentation Decathlon. Hepatic vessels are challenging and hard if addressed by common 2D networks. The dataset contains 303 cases. We split the dataset into 228 cases for training and 75

	3-slice baseline	12-slice baseline	12-slice None	12-slice Concat	12-slice Dot	12-slice channel non-local
$DSC_C$	71.20%	70.15%	72.73%	72.27%	71.16%	73.23%
DSC <sub>S</sub>	70.71%	67.03%	71.91%	71.83%	71.75%	71.67%
$DSC_A$	69.62%	70.35%	72.18%	72.04%	71.21%	72.54%
DSC <sub>F</sub>	73.32%	71.07%	73.63%	73.07%	73.74%	74.17%

Table 3: Compare different settings. Here baseline means directly input multi-slice at the beginning. None means only postpone the fusion stage. Concat, dot, and channel-wise non-local mean different way to introduce slice-sensitive intra-information to the final feature map.

	3-slice ours	6-slice ours	9-slice ours	12-slice ours	15-slice ours	18-slice ours	21-slice ours	24-slice ours
$DSC_C$	70.37%	72.49%	73.07%	73.23%	73.06%	73.04%	72.73%	72.72%
DSC <sub>S</sub>	70.49%	71.34%	71.65%	71.67%	72.97%	72.17%	73.17%	72.46%
$DSC_A$	70.70%	70.59%	72.76%	72.54%	72.64%	-	-	-
$DSC_F$	72.61%	73.21%	74.05%	74.17%	74.27%	73.61%	74.17%	73.78%

Table 4: Different slice thickness for our method on SMA. We tried different slice thickness from 3 to 24. It can be observed that our model performs well even for a much larger slice thickness. '-' means the model does not converge, if some plane result is missing, the fusion will be the average on the remaining planes.

	SMA	Celiac AA	Duodenum
3-slice baseline	18.60	12.79	21.88
12-slice baseline	20.70	12.99	20.68
15-slice Ours	11.42	11.37	19.05

Table 5: Inter-slice similarity among different methods compared with GT. The results are based on axial plane. Smaller number is better.

cases for validation. Results are shown in Table 2. Our results outperforms both 2D baseline and 3D baseline, which illustrates the ability of our approach in incorporating more inter-slice information.

### 4.4. Ablation Experiments

We have done two ablation experiments, one is to know each modification influence and another to know compare our method performance under different slice thickness.

**Different type of introducing slice-sensitive information** We have tried different ways to combine the specific-slice feature and all-slice feature based on 12-slice models. The results are shown in Table 3. As we can see, directly fusing the different types of feature can already bring benefits to the model. But the improvement will be less if we introduce slice-sensitive information in a naive way, like concatenate or element-wise dot. With a channel-wise non-local module, the feature maps will be updated based on channelwise relationship before sum, thus the performance can be boosted higher.

**More slices input** We tried different slice thickness to test our model capacity, from at least 3-slice to at most 24-slice under 12G GPU memory. The results are summarized in Table 4. We found that the results keep increas-

ing until slice thickness reaches 15, while the axial model over 18-slice model cannot converge. The trend of performance increase has a different relationship when slice thickness increases. For plane X, a major improvement happens when slice thickness increase from 3 to 6 (+2.12%). And slice thickness over 9 produce similar results. For plane Y, there happen two major improvements, one is from 3 to 6 (+0.85%) and another is from 12 to 15 (+1.30%), while 6slice, 9-slice, and 12-slice have similar accuracy. For plane Z, when increasing slice numbers from 6 to 9, the result increases by 2.17%. It seems that there exist some bottlenecks when increasing slice thickness, and breaking through the bottleneck can bring most significant improvement.

# 4.5. Visualization and Diagnoses

**Qualitative Results** We visualize a case of Hepathic Vessel and a case of celiac AA segmentation (see Fig 3) with ITK-SNAP [29]. First two rows are for Hepathic Vessel, we could see that it is contiguous with multi-branch, just like a tree. Compared with baseline, our method can do better in capturing this vessels' contiguity. As for Celiac AA, the third row is our baseline result and the fourth is results of our 15-slice thickened 2D network. In the baseline, single-view prediction suffers from lack of contextual information, leading to an unsatisfying result. In our method, even single view could capture the contiguous shape of celiac aa and achieve a good result, even better than baseline fusion. We observe no obvious failure (a typical fail in baseline single axis result) and a better segmentation consistency in our approach.

**Statistical Diagnose** To prove that our method achieves a better performance because it can distinguish each slice more clearly, we design a statistics called *inter-slice simi*-



Figure 3: 3D visualization for Hepathic Vessel (the first two rows) and Celiac AA (the last two rows with first baseline and second our method) segmentation result. This figure is best viewed in color.

*larity.* We compute the dice score between each two neighboring slices along a certain axis, forming a vector. We use this vector to represent a volume prediction neighboring slice similarity. By comparing this vector from prediction with that from ground-truth, we can evaluate each method on how well they can distinguish neighboring slices. We use L2-distance to evaluate the similarity between two vectors. By comparing our algorithm with 3-slice baseline and 12-slice baseline on all testing cases, we show that our method has a better distinguish ability, which makes it possible for our model to predict more accurate on each slice even the there is a thickened input. The result is summarized in Table 5.

#### **5.** Conclusions

This paper is motivated by the need of capturing 3D contexts, and aims at designing a network structure which works on thickened 2D inputs. The major obstacle is the in-

formation loss brought by fusion along the third dimension. We propose two keys to deal with this issue, *i.e.*, postponing the stage of information fusion, and creating a shortcut connection between the pre-fusion stage and the final decision stage. Evaluated on a few medical image segmentation datasets, our approach reports higher segmentation accuracy with thickened input data, demonstrating the effectiveness of 3D contexts.

Our research sheds light on designing efficient 3D networks for segmenting volumetric data. The success of our approach provides a side evidence that both 2D and 3D networks are not the optimal solution, as 2D methods benefit from natural image pre-training but inevitably lack contexts, meanwhile 3D methods usually waste time and memory for unnecessary computations. Absorbing advantages from both of them remains an open problem and deserves more efforts in future research.

# References

- A. M. Ali, A. A. Farag, and A. S. El-Baz. Graph cuts framework for kidney segmentation with prior shape constraints. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 384– 392. Springer, 2007. 2
- [2] J. Cai, L. Lu, Z. Zhang, F. Xing, L. Yang, and Q. Yin. Pancreas segmentation in mri using graph-based decision fusion on convolutional neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 442–450. Springer, 2016. 2
- [3] J. Chen, L. Yang, Y. Zhang, M. Alber, and D. Z. Chen. Combining fully convolutional and recurrent neural networks for 3d biomedical image segmentation. In *Advances in neural information processing systems*, pages 3036–3044, 2016. 2
- [4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018. 2, 3, 5
- [5] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3640–3649, 2016. 2
- [6] C. Chu, M. Oda, T. Kitasaka, K. Misawa, M. Fujiwara, Y. Hayashi, Y. Nimura, D. Rueckert, and K. Mori. Multiorgan segmentation based on spatially-divided probabilistic atlas from 3d abdominal ct images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 165–172. Springer, 2013. 2
- [7] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3D u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, 2016. 2
- [8] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.-M. Jodoin, and H. Larochelle. Brain tumor segmentation with deep neural networks. *Medical image analysis*, 35:18–31, 2017. 2
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [10] S. Hu, E. A. Hoffman, and J. M. Reinhardt. Automatic lung segmentation for accurate quantitation of volumetric x-ray ct images. *IEEE transactions on medical imaging*, 20(6):490– 498, 2001. 2
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 2
- [12] M. Lai. Deep learning for medical image segmentation. arXiv preprint arXiv:1505.02000, 2015. 2
- [13] X. Li, H. Chen, X. Qi, Q. Dou, C.-W. Fu, and P. A. Heng. Hdenseunet: Hybrid densely connected unet for liver and liver tumor segmentation from ct volumes. *IEEE Transactions on Medical Imaging*, 2017. 2
- [14] S. Liu, D. Xu, S. K. Zhou, O. Pauly, S. Grbic, T. Mertelmeier, J. Wicklein, A. Jerebko, W. Cai, and D. Comaniciu. 3d

anisotropic hybrid network: Transferring convolutional features from 2d images to 3d anisotropic volumes. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 851–858. Springer, 2018. 2

- [15] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [16] P. Luo, J. Ren, and Z. Peng. Differentiable learning-tonormalize via switchable normalization. arXiv preprint arXiv:1806.10779, 2018. 5
- [17] F. Milletari, N. Navab, and S.-A. Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In 2016 Fourth International Conference on 3D Vision (3DV), pages 565–571. IEEE, 2016. 2, 5, 6
- [18] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2, 3
- [19] H. R. Roth, L. Lu, A. Farag, H.-C. Shin, J. Liu, E. B. Turkbey, and R. M. Summers. Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In *MICCAI*, 2015. 2
- [20] H. R. Roth, L. Lu, A. Farag, A. Sohn, and R. M. Summers. Spatial aggregation of holistically-nested networks for automated pancreas segmentation. In *MICCAI*, 2016. 2
- [21] N. Tajbakhsh, J. Y. Shin, S. R. Gurudu, R. T. Hurst, C. B. Kendall, M. B. Gotway, and J. Liang. Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE transactions on medical imaging*, 35(5):1299– 1312, 2016. 2
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. 4
- [23] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 4, 5
- [24] Z. Wang, K. K. Bhatia, B. Glocker, A. Marvao, T. Dawes, K. Misawa, K. Mori, and D. Rueckert. Geodesic patch-based segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 666–673. Springer, 2014. 2
- [25] C.-Y. Wu, C. Feichtenhofer, H. Fan, K. He, P. Krähenbühl, and R. Girshick. Long-term feature banks for detailed video understanding. *arXiv preprint arXiv:1812.05038*, 2018. 4
- [26] Y. Xia, L. Xie, F. Liu, Z. Zhu, E. K. Fishman, and A. L. Yuille. Bridging the gap between 2d and 3d organ segmentation with volumetric fusion net. In *International Conference* on Medical Image Computing and Computer-Assisted Intervention, pages 445–453. Springer, 2018. 2
- [27] Q. Yu, L. Xie, Y. Wang, Y. Zhou, E. K. Fishman, and A. L. Yuille. Recurrent saliency transformation network: Incorporating multi-stage visual cues for small organ segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2

- [28] K. Yue, M. Sun, Y. Yuan, F. Zhou, E. Ding, and F. Xu. Compact generalized non-local network. In Advances in Neural Information Processing Systems, pages 6511–6520, 2018. 4
- [29] P. A. Yushkevich, J. Piven, H. Cody Hazlett, R. Gimpel Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006. 7
- [30] Y. Zhou, L. Xie, W. Shen, Y. Wang, E. K. Fishman, and A. L. Yuille. A fixed-point model for pancreas segmentation in abdominal ct scans. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 693–701. Springer, 2017. 2
- [31] Z. Zhu, Y. Xia, W. Shen, E. K. Fishman, and A. L. Yuille. A 3d coarse-to-fine framework for automatic pancreas segmentation. *arXiv preprint arXiv:1712.00201*, 2017. 2