

A 3D Coarse-to-Fine Framework for Volumetric Medical Image Segmentation

Zhuotun Zhu¹, Yingda Xia¹, Wei Shen^{1,2}, Elliot K. Fishman³, Alan L. Yuille¹

¹Johns Hopkins University, ²Shanghai University, ³Johns Hopkins University School of Medicine

{ztzhu, yxia25, ayuille1}@jhu.edu wei.shen@t.shu.edu.cn efishman@jhmi.edu

Abstract

In this paper, we adopt 3D Convolutional Neural Networks to segment volumetric medical images. Although deep neural networks have been proven to be very effective on many 2D vision tasks, it is still challenging to apply them to 3D tasks due to the limited amount of annotated 3D data and limited computational resources. We propose a novel 3D-based coarse-to-fine framework to **effectively** and **efficiently** tackle these challenges. The proposed 3D-based framework outperforms the 2D counterpart to a large margin since it can leverage the rich spatial information along all three axes. We conduct experiments on two datasets which include healthy and pathological pancreases respectively, and achieve the current state-of-the-art in terms of Dice-Sørensen Coefficient (DSC). On the NIH pancreas segmentation dataset, we outperform the previous best by an average of over 2%, and the worst case is improved by 7% to reach almost 70%, which indicates the reliability of our framework in clinical applications.

1. Introduction

Driven by the huge demands for computer-aided diagnosis systems, automatic organ segmentation from medical images, such as computed tomography (CT) and magnetic resonance imaging (MRI), has become an active research topic in both the medical image processing and computer vision communities. It is a prerequisite step for many clinical applications, such as diabetes inspection, organic cancer diagnosis, and surgical planning. Therefore, it is well worth exploring automatic segmentation systems to accelerate the computer-aided diagnosis in medical image analysis.

In this paper, we focus on pancreas segmentation from CT scans, one of the most challenging organ segmentation problems [34][25]. As shown in Fig. 1, the main difficulties stem from three aspects: 1) the small size of the pancreas in the whole abdominal CT volume; 2) the large variations in texture, location, shape and size of the pancreas; 3) the abnormalities, like pancreatic cysts, can alter the appearance of pancreases a lot.




Figure 1. An illustration of normal pancreases on NIH dataset [25] and abnormal cystic pancreases on JHMI dataset [33] shown in the first and second row respectively. Normal pancreas regions are masked as red and abnormal pancreas regions are marked as blue. The pancreas usually occupies a small region in a whole CT scan. Best viewed in color.

Following the rapid development of deep neural networks [15][29] and their successes in many computer vision tasks, such as semantic segmentation [18][4], edge detection [28][30][27] and 3D shape retrieval [35][7], many deep learning based methods have been proposed for pancreas segmentation and achieved considerable progress [34][25][26]. However, these methods are based on 2D fully convolutional networks (FCNs) [18], which perform segmentation slice by slice while CT volumes are indeed 3D data. Although these 2D methods use strategies to fuse the output from different 2D views to obtain 3D segmentation results, they inevitably lose some 3D context, which is important for capturing the discriminative features of the pancreas with respect to background regions.

Using 3D deep networks for organ segmentation is a recent trend but not yet applied to the pancreas. An obstacle is that training 3D deep networks suffers from the “out of memory” problem. 2D FCNs can accept a whole 2D slice as input, but 3D FCNs cannot be fed a whole 3D volume due to the limited GPU memory size. A common solution is to train 3D FCNs from small sub-volumes and test them in a sliding-window manner [20][1][5][3][31], *i.e.*, performing 3D segmentation on densely and uniformly sampled sub-volumes one by one. Usually, these neighboring sampled sub-volumes overlap with each other to improve the robustness of the final 3D results. It is worth noting that the overlap size is a trade-off between the segmentation accu-

racy and the time cost. Setting a larger/smaller overlap size generally leads to a better/worse segmentation accuracy but takes more/less time during testing.

To address these issues, we propose a concise and effective framework based on 3D deep networks for pancreas segmentation, which can simultaneously achieve high segmentation accuracy and low time cost. Our framework is formulated in a coarse-to-fine manner. In the training stage, we first train a 3D FCN from the sub-volumes sampled from an entire CT volume. We call this *ResDSN Coarse* model, which aims to obtain the rough location of the target pancreas from the whole CT volume by making full use of the overall 3D context. Then, we train another 3D FCN from the sub-volumes sampled only from the ground truth bounding boxes of the target pancreas. We call this the *ResDSN Fine* model, which can refine the segmentation based on the coarse result. In the testing stage, we first apply the coarse model in the sliding-window manner to a whole CT volume to extract the most probable location of the pancreas. Since we only need a rough location for the target pancreas in this step, the overlap size is set to a small value. Afterwards, we apply the fine model in the sliding-window manner to the coarse pancreas region, but by setting a larger overlap size. Thus, we can efficiently obtain a fine segmentation result and we call the coarse-to-fine framework by *ResDSN C2F*.

Note that, the meaning of “coarse-to-fine” in our framework is twofold. First, it means the input region of interests (RoIs) for *ResDSN Coarse* model and *ResDSN Fine* model are different, i.e., a whole CT volume for the former one and a rough bounding box of the target pancreas for the latter one. We refer to this as coarse-to-fine RoIs, which is designed to achieve better segmentation performance. The coarse step removes a large amount of the unrelated background region, then with a relatively smaller region to be sampled as input, the fine step can much more easily learn cues which distinguish the pancreas from the local background, i.e., exploit local context which makes it easier to obtain a more accurate segmentation result. Second, it means the overlap sizes used for *ResDSN Coarse* model and *ResDSN Fine* model during inference are different, i.e., small and large overlap sizes for them, respectively. We refer to this as coarse-to-fine overlap sizes, which is designed for efficient 3D inference.

To our best knowledge, we are one of the first studies to segment the challenging *normal* and *abnormal* pancreases using 3D networks which leverage the rich spatial information. The effectiveness and efficiency of the proposed 3D coarse-to-fine framework are demonstrated on two pancreas segmentation datasets where we achieve the state-of-the-art with relative low time cost. It is worth mentioning that, although our focus is pancreas segmentation, our framework is generic and can be directly applied to segmenting other medical organs.

2. Related Work

The medical image analysis community is facing a revolution brought by the fast development of deep networks [15][29]. Deep Convolutional Neural Networks (CNNs) based methods have dominated the research area of volumetric medical image segmentation in the last few years. Generally speaking, CNN-based methods for volumetric medical image segmentation can be divided into two major categories: 2D CNNs based and 3D CNNs based.

2.1. 2D CNNs for Volumetric Segmentation

2D CNNs based methods [25][26][11][21][23] performed volumetric segmentation slice by slice from different views, and then fused the 2D segmentation results to obtain a 3D Volumetric Segmentation result. In the early stage, the 2D segmentation based models were trained from image patches and tested in a patch by patch manner [25], which is time consuming. Since the introduction of fully convolution networks (FCNs) [18], almost all the 2D segmentation methods are built upon 2D FCNs to perform holistic slice segmentation during both training and testing. Havaei *et al* [11] proposed a two-pathway FCN architecture, which exploited both local features as well as more global contextual features simultaneously by the two pathways. Roth *et al* [26] performed Pancreas segmentation by a holistic learning approach, which first segment pancreas regions by holistically-nested networks [30] and then refine them by the boundary maps obtained by robust spatial aggregation using random forest. The U-Net [23] is one of the most popular FCN architectures for medical image segmentation, which is a encoder-decoder network, but with additional short connection between encoder and decoder paths. Based on the fact that a pancreas only takes up a small fraction of the whole scan, Zhou *et al.* [34] proposed to find the rough pancreas region and then learn a FCN based fixed-point model to refine the pancreas region iteratively. Their method is also based on a coarse-to-fine framework, but it only considered coarse-to-fine RoIs. Besides coarse-to-fine RoIs, our coarse-to-fine method also takes coarse-to-fine overlap sizes into account, which is designed specifically for efficient 3D inference.

2.2. 3D CNNs for Volumetric Segmentation

Although 2D CNNs based methods achieved considerable progresses, they are not optimal for medical image segmentation, as they cannot make full use of the 3D context encoded in volumetric data. Several 3D CNNs based segmentation methods have been proposed. The 3D U-Net [5] extended the previous 2D U-Net architecture [23] by replacing all 2D operations with their 3D counterparts. Based on the architecture of the 3D U-Net, the V-Net [20] introduced residual structures [12] (short term skip connection) into each stage of the network. Chen *et al* [3] proposed a




Figure 2. Flowchart of the proposed 3D coarse-to-fine segmentation system in the testing phase. We first apply “ResDSN Coarse” with a small overlapped sliding window to obtain a rough pancreas region and then use the “ResDSN Fine” model to refine the results with a large overlapped sliding window. Best viewed in color.

deep voxel-wise residual network for 3D brain segmentation. Both I2I-3D [19] and 3D-DSN [6] included auxiliary supervision via side outputs into their 3D deep networks. Despite the success of 3D CNNs as a technique for segmenting the target organs, such as prostate [20] and kidney [5], very few techniques have been developed for leveraging 3D spatial information on the challenging pancreas segmentation. Gibson *et al.* [8] proposed the DenseVNet which is however constrained to have shallow encoders due to the computationally demanding dense connections. Roth *et al.* [24] extended 3D U-Net to segment pancreas, which has the following shortcomings, 1) the input of their networks is fixed to $120 \times 120 \times 120$, which is very computationally demanding due to this large volume size, 2) the rough pancreas bounding box is resampled to a fixed size as their networks input, which loses information and flexibility, and cannot deal with the intrinsic large variations of pancreas in shape and size. Therefore, we propose our 3D coarse-to-fine framework that works on both normal and abnormal to ensure both low computation cost and high pancreas segmentation accuracy.

3. Method

In this section, we elaborate our proposed 3D coarse-to-fine framework which includes a *coarse* stage and a *fine* stage afterwards. We first formulate a segmentation model that can be generalized to both *coarse* stage and *fine* stage. Later in Sec. 3.1 and Sec. 3.2, we will customize the segmentation model to these two stages, separately.

We denote a 3D CT-scan volume by \mathbf{X} . This is associated with a human-labeled per-voxel annotation \mathbf{Y} , where both \mathbf{X} and \mathbf{Y} have size $W \times H \times D$, which corresponds to axial, sagittal and coronal views, separately. The ground-truth segmentation mask \mathbf{Y} has a binary value y_i , $i = 1, \dots, WHD$, at each spatial location i where $y_i = 1$ indicates that x_i is a pancreas voxel. Denote a segmentation model by $\mathbb{M} : \mathbf{P} = \mathbf{f}(\mathbf{X}; \Theta)$, where Θ indicates model parameters and \mathbf{P} means the binary prediction volume. Specifically in a neural network with L layers and parameters $\Theta = \{\mathcal{W}, \mathcal{B}\}$, \mathcal{W} is a set of weights and \mathcal{B} is a set of biases, where $\mathcal{W} = \{\mathbf{W}^1, \mathbf{W}^2, \dots, \mathbf{W}^L\}$ and

$\mathcal{B} = \{\mathbf{B}^1, \mathbf{B}^2, \dots, \mathbf{B}^L\}$. Given that $p(y_i|x_i; \Theta)$ represents the predicted probability of a voxel x_i being what is the labeled class at the final layer of the output, the negative log-likelihood loss can be formulated as:

$$\mathcal{L} = \mathcal{L}(\mathbf{X}; \Theta) = - \sum_{x_i \in \mathbf{X}} \log(p(y_i|x_i; \Theta)). \quad (1)$$

It is also known as the cross entropy loss in our binary segmentation setting. By thresholding $p(y_i|x_i; \Theta)$, we can obtain the binary segmentation mask \mathbf{P} .

We also add some auxiliary layers to such a neural network (will be called mainstream network in the rest of the paper), which produces side outputs under deep supervision [17]. These auxiliary layers form a branch network and facilitate the feature learning at lower layer of the mainstream network. Each branch network shares the weights of the first d layers from the mainstream network, which is denoted by $\Theta_d = \{\mathcal{W}_d, \mathcal{B}_d\}$. Apart from the shared weights, it owns weights $\hat{\Theta}_d$ to output the per-voxel prediction. Similarly, the loss of an auxiliary network can be formulated as:

$$\mathcal{L}_d(\mathbf{X}; \Theta_d, \hat{\Theta}_d) = \sum_{x_i \in \mathbf{X}} -\log(p(y_i|x_i; \Theta_d, \hat{\Theta}_d)), \quad (2)$$

which is abbreviated as \mathcal{L}_d . Finally, stochastic gradient descent is applied to minimize the negative log-likelihood, which is given by following regularized objective function:

$$\mathcal{L}_{overall} = \mathcal{L} + \sum_{d \in \mathcal{D}} \xi_d \mathcal{L}_d + \lambda (\|\Theta\|^2 + \sum_{d \in \mathcal{D}} \|\hat{\Theta}_d\|^2), \quad (3)$$

where \mathcal{D} is a set of branch networks for auxiliary supervisions, ξ_d balances the importance of each auxiliary network and l_2 regularization is added to the objective to prevent the networks from overfitting. For conciseness concerns in the following sections, we keep a segmentation model that is obtained from the overall function described in Eq. 3 denoted by $\mathbb{M} : \mathbf{P} = \mathbf{f}(\mathbf{X}; \Theta)$, where Θ includes parameters of the mainstream network and auxiliary networks.

3.1. Coarse Stage

In the *coarse* stage, the input of “ResDSN Coarse” is sampled from the whole CT-scan volume denoted by \mathbf{X}^C ,

on which the *coarse* segmentation model $\mathbb{M}^C : \mathbf{P}^C = \mathbf{f}^C(\mathbf{X}^C; \Theta^C)$ is trained on. All the C superscripts depict the *coarse* stage. The goal of this stage is to efficiently produce a rough binary segmentation \mathbf{P}^C from the complex background, which can get rid of regions that are segmented as non-pancreas with a high confidence to obtain an approximate pancreas volume. Based on this approximate pancreas volume, we can crop from the original input \mathbf{X}^C with a rectangular cube derived from \mathbf{P}^C to obtain a smaller 3D image space \mathbf{X}^F , which is surrounded by simplified and less variable context compared with \mathbf{X}^C . The mathematic definition of \mathbf{X}^F is formulated as:

$$\mathbf{X}^F = \text{Crop}[\mathbf{X}^C \otimes \mathbf{P}^C; \mathbf{P}^C, m], \quad (4)$$

where \otimes means an element-wise product. The function $\text{Crop}[\mathbf{X}; \mathbf{P}, m]$ denotes cropping \mathbf{X} via a rectangular cube that covers all the 1's voxels of a binary volume \mathbf{P} added by a padding margin m along three axes. Given \mathbf{P} , the functional constraint imposed on \mathbf{X} is that they have exactly the same dimensionality in 3D space. The padding parameter m is empirically determined in experiments, where it is used to better segment the boundary voxels of pancreas during the *fine* stage. The Crop operation acts as a dimensionality reduction to facilitate the fine segmentation, which is crucial to cut down the consuming time of segmentation. It is well-worth noting that the 3D locations of the rectangular cube which specifies where to crop \mathbf{X}^F from \mathbf{X}^C is recorded to map the *fine* segmentation results back their positions in the full CT scan.

3.2. Fine Stage

In the *fine* stage, the input of the ConvNet is sampled from the cropped volume \mathbf{X}^F , on which we train the *fine* segmentation model $\mathbb{M}^F : \mathbf{P}^F = \mathbf{f}^F(\mathbf{X}^F; \Theta^F)$, where the F superscripts indicate the *fine* stage. The goal of this stage is to refine the coarse segmentation results from previous stage. In practice, \mathbf{P}^F has the same volumetric size of \mathbf{X}^F , which is smaller than the original size of \mathbf{X}^C .

3.3. Coarse-to-Fine Segmentation

Our segmentation task is to give a volumetric prediction on every voxel of \mathbf{X}^C , so we need to map the \mathbf{P}^F back to exactly the same size of \mathbf{X}^C given by:

$$\mathbf{P}^{C2F} = \text{DeCrop}[\mathbf{P}^F \odot \mathbf{P}^C; \mathbf{X}^F, \mathbf{X}^C], \quad (5)$$

where \mathbf{P}^{C2F} denotes the final volumetric segmentation, and \odot means an element-wise replacement, and DeCrop operation defined on $\mathbf{P}^F, \mathbf{P}^C, \mathbf{X}^F$ and \mathbf{X}^C is to replace a pre-defined rectangular cube inside \mathbf{P}^C by \mathbf{P}^F , where the replacement locations are given by the definition of cropping \mathbf{X}^F from \mathbf{X}^C given in Eq. 4.

All in all, our entire 3D-based coarse-to-fine segmentation framework during testing is illustrated in Fig 2.

3.4. Network Architecture

As shown in Fig. 3, we provide an illustration of our convolutional network architecture. Inspired by V-Net [20], 3D U-Net [5], and VoxResNet [3], we have an encoder path followed by a decoder path each with four resolution steps. The left part of network acts as a feature extractor to learn higher and higher level of representations while the right part of network decompresses compact features into finer and finer resolution to predict the per-voxel segmentation. The padding and stride of each layer (Conv, Pooling, DeConv) are carefully designed to make sure the densely predicted output is the same size as the input.

The encoder sub-network on the left is divided into different steps that work on different resolutions. Each step consists of one to two convolutions, where each convolution is composed of $3 \times 3 \times 3$ convolution followed by a batch normalization (BN [13]) and a rectified linear unit (ReLU [22]) to reach better convergence, and then a max pooling layer with a kernel size of $2 \times 2 \times 2$ and strides of two to reduce resolutions and learn more compact features. The downsampling operation implemented by max-pooling can reduce the size of the intermediate feature maps while increasing the size of the receptive fields. Having fewer size of activations makes it possible to double the number of channels during feature aggregation given the limited computational resource.

The decoder sub-network on the right is composed of several steps that operate on different resolutions as well. Each step has two convolutions with each one followed by a BatchNorm and a ReLU, and afterwards a Deconvolution with a kernel size of $4 \times 4 \times 4$ and strides of two is connected to expand the feature maps and finally predict the segmentation mask at the last layer. The upsampling operation that is carried out by deconvolution enlarges the resolution between each step, which increases the size of the intermediate activations so that we need to halve the number of channels due to the limited memory of the GPU card.

Apart from the left and right sub-networks, we impose a residual connection [12] to bridge short-cut connections of features between low-level layers and high-level layers. During the forward phase, the low-level cues extracted by networks are directly added to the high-level cues, which can help elaborate the fine-scaled segmentation, *e.g.*, small parts close to the boundary which may be ignored during the feature aggregation due to the large size of receptive field at high-level layers. As for the backward phase, the supervision cues at high-level layers can be back-propagated through the short-cut way via the residual connections. This type of mechanism can prevent networks from gradient vanishing and exploding [9], which hampers network convergence during training.

We have one mainstream loss layer connected from “Res/Conv1b” and another two auxiliary loss layers con-




Figure 3. Illustration of our 3D convolutional neural network for volumetric segmentation. The encoder path is composed from “Conv1a” to “Conv4b” while the decoder path is from “DeConv3a” to “Res/Conv1b”. Each convolution or deconvolution layer consists of one convolution followed by a BatchNorm and a ReLU. To clarify, “Conv1a, 32, $3 \times 3 \times 3$ ” means the convolution operation with 32 channels and a kernel size of $3 \times 3 \times 3$. “Pooling 1, max, 2” means the max pooling operation with kernel size of $2 \times 2 \times 2$ and a stride of two. Long residual connections are illustrated by the blue concrete lines. Blocks with same color mean the same operations. Best viewed in color.

nected from “Conv2b” and “Conv3b” to the ground truth label, respectively. For the mainstream loss in “Res/Conv1b” as the last layer which has the same size of data flow as one of the input, a $1 \times 1 \times 1$ convolution is followed to reduce the number of channels to the number of label classes which is 2 in our case. As for the two auxiliary loss layers, deconvolution layers are connected to upsample feature maps to be the same as input.

The deep supervision imposed by auxiliary losses provides robustness to hyper-parameters choice, in that the low-level layers are guided by the direct segmentation loss, leading to faster convergence rate. Throughout this work, we have two auxiliary branches where the default parameters are $\xi_1 = 0.2$ and $\xi_2 = 0.4$ in Eq. 3 to control the importance of deep supervisions compared with the major supervision from the mainstream loss for all segmentation networks.

As shown in Table 1, we give the detailed comparisons of network configurations in terms of four aspects: long residual connection, short residual connection, deep supervision and loss function. Our backbone network architecture, named as “ResDSN”, is proposed with different strategies in terms of combinations of long residual connection and short residual connection compared with VoxResNet [3], 3D HED [19], 3D DSN [6] and MixedResNet [32]. In this table, we also depict “FResDSN” and “SResDSN”, where “FResDSN” and “SResDSN” are similar to MixedResNet [32] and VoxResNet [3], respectively. As confirmed by our quantitative experiments in Sec. 5.1, instead of adding short residual connections to the network, e.g., “FResDSN” and “SResDSN”, we only choose the long residual element-wise sum, which can be more computationally efficient while even performing better than the

Method	Long Res	Short Res	Deep Super	Loss
ResDSN (Ours)	Sum	No	Yes	CE
FResDSN	Sum	Sum	Yes	CE
SResDSN	No	Sum	Yes	CE
3D U-Net [5]	Concat	No	No	CE
V-Net [20]	Concat	Sum	No	DSC
VoxResNet [3]	No	Sum	Yes	CE
MixedResNet [32]	Sum	Sum	Yes	CE
3D DSN [6]	No	No	Yes	CE
3D HED [19]	Concat	No	Yes	CE

Table 1. Configurations comparison of different 3D segmentation networks on medical image analysis. For all the abbreviated phrases, “Long Res” means long residual connection, “Short Res” means short residual connection, “Deep Super” means deep supervision implemented by auxiliary loss layers, “Concat” means concatenation, “DSC” means Dice-Sørensen Coefficient and “CE” means cross-entropy. For residual connection, it has two types: concatenation (“Concat”) or element-wise sum (“Sum”).

“FResDSN” architecture which is equipped with both long and short residual connections. Moreover, ResDSN has noticeable differences with respect to the V-Net [20] and 3D U-Net [5]. On the one hand, compared with 3D U-Net and V-Net which concatenate the lower-level local features to higher-level global features, we adopt the element-wise sum between these features, which outputs less number of channels for efficient computation. On the other hand, we introduce deep supervision via auxiliary losses into the network to yield better convergence.

4. Experiments

In this section, we first describe in detail how we conduct training and testing in the *coarse* and *fine* stages, respectively. Then we are going to compare our proposed method

with previous state-of-the-art on two pancreas datasets: NIH pancreas dataset [25] and JHMI pathological pancreas dataset [33].

4.1. Network Training and Testing

All our experiments were run on a desktop equipped with the NVIDIA TITAN X (Pascal) GPU and deep neural networks were implemented based on the CAFFE [14] platform customized to support 3D operations for all necessary layers, *e.g.*, “convolution”, “deconvolution” and “pooling”, *etc.* For the data pre-processing, we simply truncated the raw intensity values to be in $[-100, 240]$ and then normalized each raw CT case to have zero mean and unit variance to decrease the data variance caused by the physical processes [10] of medical images. As for the data augmentation in the training phase, unlike sophisticated processing used by others, *e.g.*, elastic deformation [20][23], we utilized simple but effective augmentations on all training patches, *i.e.*, rotation (90° , 180° , and 270°) and flip in all three axes (axial, sagittal and coronal), to increase the number of 3D training samples which can alleviate the scarce of CT scans with expensive human annotations. Note that different CT cases have different physical resolutions, but we keep their resolutions unchanged. The input size of all our networks is denoted by $W_I \times H_I \times D_I$, where $W_I = H_I = D_I = 64$.

For the *coarse* stage, we randomly sampled $64 \times 64 \times 64$ sub-volumes from the whole CT scan in the training phase. In this case, a sub-volume can either cover a portion of pancreas voxels or be cropped from regions with non-pancreas voxels at all, which acts as a hard negative mining to reduce the false positive. In the testing phase, a sliding window was carried out to the whole CT volume with a *coarse* stepsize that has small overlaps within each neighboring sub-volume. Specifically, for a testing volume with a size of $W \times H \times D$, we have a total number of $(\lfloor \frac{W}{W_I} \rfloor + n) \times (\lfloor \frac{H}{H_I} \rfloor + n) \times (\lfloor \frac{D}{D_I} \rfloor + n)$ sub-volumes to be fed into the network and then combined to obtain the final prediction, where n is a parameter to control the sliding overlaps that a larger n results in a larger overlap and vice versa. In the *coarse* stage for the low time cost concern, we set $n = 6$ to efficiently locate the rough region of pancreas \mathbf{X}^F defined in Eq. 4 from the whole CT scan \mathbf{X}^C .

For the *fine* stage, we randomly cropped $64 \times 64 \times 64$ sub-volumes constrained to be from the pancreas regions defined by ground-truth labels during training. In this case, a training sub-volume was assured to cover pancreatic voxels, which was specifically designed to be capable of segmentation refinement. In the testing phase, we only applied the sliding window on \mathbf{X}^F with a size of $W_F \times H_F \times D_F$. The total number of sub-volumes to be tested is $(\lfloor \frac{W_F}{W_I} \rfloor + n) \times (\lfloor \frac{H_F}{H_I} \rfloor + n) \times (\lfloor \frac{D_F}{D_I} \rfloor + n)$. In the *fine* stage for the high accuracy performance concern, we set $n = 12$ to accurately estimate the pancreatic mask \mathbf{P}^F from the rough segmen-

tation volume \mathbf{X}^F . In the end, we mapped the \mathbf{P}^F back to \mathbf{P}^C to obtain \mathbf{P}^{C2F} for the final pancreas segmentation as given in Eq. 5, where the mapping location is given by the cropped location of \mathbf{X}^F from \mathbf{X}^C .

After we get the final binary segmentation mask, we denote \mathcal{P} and \mathcal{Y} to be the set of pancreas voxels in the prediction and ground truth, separately, *i.e.*, $\mathcal{P} = \{i | p_i = 1\}$ and $\mathcal{Y} = \{i | y_i = 1\}$. The evaluation metric is defined by the Dice-Sørensen Coefficient (DSC) formulated as $\text{DSC}(\mathcal{P}, \mathcal{Y}) = \frac{2 \times |\mathcal{P} \cap \mathcal{Y}|}{|\mathcal{P}| + |\mathcal{Y}|}$. This evaluation measurement ranges in $[0, 1]$ where 1 means a perfect prediction.

4.2. NIH Pancreas Dataset

We conduct experiments on the NIH pancreas segmentation dataset [25], which contains 82 contrast-enhanced abdominal CT volumes provided by an experienced radiologist. The size of CT volumes is $512 \times 512 \times D$, where $D \in [181, 466]$ and their spatial resolutions are $w \times h \times d$, where $d = 1.0\text{mm}$ and $w = h$ that ranges from 0.5mm to 1.0mm . Data pre-processing and data augmentation were described in Sec. 4.1. Note that we did not normalize the spatial resolution into the same one since we wanted to impose the networks to learn to deal with the variations between different volumetric cases. Following the training protocol [25], we perform 4-fold cross-validation in a random split from 82 patients for training and testing folds, where each testing fold has 21, 21, 20 and 20 cases, respectively. We trained networks illustrated in Fig. 3 by SGD optimizer with a 16 mini-batch, a 0.9 momentum, a base learning rate to be 0.01 via polynomial decay (the power is 0.9) in a total of 80,000 iterations, and the weight decay 0.0005. Both training networks in the *coarse* and *fine* stages shared the same training parameter settings except that they took a $64 \times 64 \times 64$ input sampled from different underlying distributions described in Sec. 4.1, which included the details of testing settings as well. We average the score map of overlapped regions from the sliding window and throw away small isolated predictions whose portions are smaller than 0.2 of the total prediction, which can remove small false positives. For DSC evaluation, we report the average with standard deviation, max and min statistics over all 82 testing cases as shown in Table 2.

First of all, our overall coarse-to-fine framework outperforms previous state-of-the-art by nearly 2.2% (Cai *et al.* [2] and Zhou *et al.* [34]) in terms of average DSC, which is a large improvement. The lower standard deviation of DSC shows that our method is the most stable and robust across all different CT cases. Although the enhancement of max DSC of our framework is small due to the saturation, the improvement of min DSC over the second best (Dou *et*

¹The results are reported by our runs using the same cross-validation splits where code is available from their GitHub: <https://github.com/yulequan/HeartSeg>.

Method	Mean DSC	Max DSC	Min DSC
ResDSN C2F (Ours)	84.59 ± 4.86%	91.45%	69.62%
ResDSN Coarse (Ours)	83.18 ± 6.02%	91.33%	58.44%
Cai <i>et al.</i> [2]	82.4 ± 6.7%	90.1%	60.0%
Zhou <i>et al.</i> [34]	82.37 ± 5.68%	90.85%	62.43%
Dou ¹ <i>et al.</i> [6]	82.25 ± 5.91%	90.32%	62.53%
Roth <i>et al.</i> [26]	78.01 ± 8.20%	88.65%	34.11%
Yu ¹ <i>et al.</i> [31]	71.96 ± 15.34%	89.27%	0%

Table 2. Evaluation of different methods on the NIH dataset. Our proposed framework achieves state-of-the-art by a large margin compared with previous state-of-the-arts.

al. [6]) is from 62.53% to 69.62%, which is a more than 7% advancement. The worst case almost reaches 70%, which is a reasonable and acceptable segmentation result. After coarse-to-fine, the segmentation result of the worst case is improved by more than 11% after the 3D-based refinement from the 3D-based coarse result. The overall average DSC was also improved by 1.41%, which proves the effectiveness of our framework.

As shown in Fig 4, we report the segmentation results by “ResDSN Coarse” and “ResDSN C2F” on the same slice for comparison. Note that yellow regions are the correctly predicted pancreas. For the NIH case #33, which is the min DSC case reported by both “ResDSN Coarse” and “ResDSN C2F”, the “ResDSN C2F” successfully predict more correct pancreas regions at the bottom, which is obviously missed by “ResDSN Coarse”. If the coarse segmentation is bad, *e.g.*, case #33 and #63, our 3D coarse-to-fine can significantly improve the segmentation results by as much as 10% in DSC. However, if the coarse segmentation is already very good, *e.g.*, case #74, our proposed method cannot improve too much. We conclude that our proposed “ResDSN C2F” shows its advancement over 2D methods by aggregating rich spatial information and is more powerful than other 3D methods on the challenging pancreas segmentation task.

4.3. JHMI Pathological Pancreas Dataset

We verified our proposed idea on the JHMI pathological cyst dataset [33] of abdominal CT scans as well. Different from the NIH healthy pancreas dataset, this dataset includes pathological cysts where some can be or can become cancerous. The pancreatic cancer stage largely influences the morphology of the pancreas [16] that makes this dataset extremely challenging for considering the large variants.

This dataset has a total number of 131 contrast-enhanced abdominal CT volumes with human-labeled pancreas annotations. The size of CT volumes is $512 \times 512 \times D$, where $D \in [358, 1121]$ that spans a wider variety of thickness than one of the NIH dataset. Following the training protocol [33], we conducted 4-fold cross validation on this dataset where each testing fold has 33, 33, 32 and 33 cases, respectively. We trained networks illustrated in Fig. 3 in

both the *coarse* and *fine* stage with the same training settings as on the NIH except that we trained a total of 300,000 iterations on this pathological dataset since a pancreas with cysts is more difficult to segment than a normal case. In the testing phase, we vote the prediction map of overlapped regions from the sliding window and ignore small isolated pancreas predictions whose portions are smaller than 0.05 of the total prediction. As shown in Table. 3, we compare our framework with only one available published results on this dataset. “ResDSN C2F” achieves an average 80.56% DSC that consistently outperforms the 2D based coarse-to-fine method [33], which confirms the advantage of leveraging the rich spatial information along three axes. What’s more, the “ResDSN C2F” improves the “ResDSN Coarse” by 2.60% in terms of the mean DSC, which is a remarkable improvement that proves the effectiveness of the proposed 3D coarse-to-fine framework. Both [33] and our method have multiple failure cases whose testing DSC are 0, which indicates the segmentation of pathological organs is a more tough task. Due to these failure cases, we observe a large deviation on this pathological pancreas dataset compared with results on the NIH healthy pancreas dataset.

Method	Mean DSC
ResDSN C2F (Ours)	80.56 ± 13.36%
ResDSN Coarse (Ours)	77.96 ± 13.36%
Zhou <i>et al.</i> [33]	79.23 ± 9.72%

Table 3. Evaluations on the JHMI pathological pancreas.

5. Discussion

In this section, we conduct the ablation studies about residual connection, time efficiency and deep supervision to further investigate the effectiveness and efficiency of our proposed framework for pancreas segmentation.

5.1. Residual Connection

We discuss how different combinations of residual connections contribute to the pancreas segmentation task on the NIH dataset. All the residual connections are implemented in the element-wise sum and they shared exactly the same deep supervision connections, cross-validation splits, data input, training and testing settings except that the residual structure is different from each other. As given in Table. 4, we compare four configurations of residual connections of 3D based networks only in the *coarse* stage. The major differences between our backbone network “ResDSN” with respect to “FResDSN”, “SResDSN” and “DSN” are depicted in Table. 1. “ResDSN” outperforms other network architectures in terms of average DSC and a small standard deviation even through the network is not as sophisticated as “FResDSN”, which is the reason we adopt “ResDSN” for efficiency concerns in the *coarse* stage.




Figure 4. Examples of segmentation results reported by “ResDSN Coarse” and “ResDSN C2F” on a same slice in the axial view from NIH case #33, #63 and #74, respectively. Numbers after “Coarse” or “C2F” mean testing DSC. Red, green and yellow indicate the ground truth, prediction and overlapped regions, respectively. Best viewed in color.

Method	Mean DSC	Max DSC	Min DSC
ResDSN Coarse (Ours)	83.18 ± 6.02%	91.33%	58.44%
FResDSN Coarse	83.11 ± 6.53%	91.34%	61.97%
SResDSN Coarse	82.82 ± 5.97%	90.33%	62.43%
DSN [6] Coarse	82.25 ± 5.91%	90.32%	62.53%

Table 4. Evaluation of different residual connections on NIH.

Method	Mean DSC	n	Testing Time (s)
ResDSN C2F (Ours)	84.59 ± 4.86%	6&12	245
ResDSN Coarse (Ours)	83.18 ± 6.02%	6	111
ResDSN Fine (Ours)	83.96 ± 5.65%	12	382

Table 5. Average time cost in the testing phase, where n controls the overlap size of sliding windows during the inference.

5.2. Time Efficiency

We discuss the time efficiency of the proposed coarse-to-fine framework with a smaller overlap in the *coarse* stage for the low consuming time concern while a larger one in the *fine* stage for the high prediction accuracy concern. The overlap size depends on how large we choose n defined in Sec 4.1. We choose $n = 6$ during the coarse stage while $n = 12$ during the fine stage. Experimental results are shown in Table 5. “ResDSN Coarse” is the most efficient while the accuracy is the worst among three methods, which makes sense that we care more of the efficiency to obtain a rough pancreas segmentation. “ResDSN Fine” is to use a large overlap on an entire CT scan to do the segmentation which is the most time-consuming. In our coarse-to-fine framework, we combine the two advantages together to propose “ResDSN C2F” which can achieve the best segmentation results while the average testing time cost for each case is reduced by 36% from 382s to 245s compared with “ResDSN Fine”. In comparison, it takes an experienced board certified Abdominal Radiologist 20 mins for one case, which verifies the clinical use of our framework.

5.3. Deep Supervision

We discuss how effective of the auxiliary losses to demonstrate the impact of the deep supervision on our 3D coarse-to-fine framework. Basically, we train our mainstream networks without any auxiliary losses for both

Method	Mean DSC	Max DSC	Min DSC
ResDSN C2F (Ours)	84.59 ± 4.86%	91.45%	69.62%
Res C2F	84.06 ± 6.51%	91.72%	51.83%

Table 6. Discussions of the deep supervision on NIH.

coarse and fine stages, denoted as “Res C2F”, while keeping all other settings as the same, *e.g.*, cross-validation splits, data pre-processing and post-processing. As shown in Table 6, “ResDSN C2F” outperforms “Res C2F” by 17.79% to a large extent on min DSC and 0.53% better on average DSC though it’s a little bit worse on max DSC. We conclude that 3D coarse-to-fine with deep supervisions performs better and especially more stable on the pancreas segmentation.

6. Conclusion

In this work, we propose a novel 3D network called “ResDSN” integrated with a coarse-to-fine framework to simultaneously achieve high segmentation accuracy and low time cost. The backbone network “ResDSN” is carefully designed to only have long residual connections for efficient inference. To our best knowledge, we are one of the first works to segment the challenging pancreas using 3D networks which leverage the rich spatial information to achieve the state-of-the-art. On widely-used datasets, the worst segmentation case is experimentally improved a lot by our coarse-to-fine framework. What’s more, our coarse-to-fine framework can work on both normal and abnormal pancreases to achieve good segmentation accuracy.

Though this work mainly focuses on segmentation for the pancreas, we can naturally apply the proposed idea to other small organs, *e.g.*, spleen, duodenum and gallbladder, *etc.* In the future, we will target on error causes that lead to inaccurate segmentation to make our framework more stable, and extend our 3D coarse-to-fine framework to cyst segmentation which can cause cancerous tumors, and the very important tumor segmentation [36] task.

Acknowledgements This work was supported by the Lustgarten Foundation for Pancreatic Cancer Research, and National Natural Science Foundation of China No. 61672336. We appreciate enormous help from Seyoun Park, Lingxi Xie, Yuyin Zhou, Yan Wang, Fengze Liu, and valuable discussions from Qing Liu, Yan Zheng, Chenxi Liu, Zhe Ren.

References

- [1] T. D. Bui, J. Shin, and T. Moon. 3D densely convolution networks for volumetric segmentation. *arXiv:1709.03199*, 2017. [1](#)
- [2] J. Cai, L. Lu, Y. Xie, F. Xing, and L. Yang. Improving deep pancreas segmentation in CT and MRI images via recurrent neural contextual learning and direct loss function. 2017. [6](#), [7](#)
- [3] H. Chen, Q. Dou, L. Yu, J. Qin, and P.-A. Heng. Voxresnet: Deep voxelwise residual networks for brain segmentation from 3D MR images. *NeuroImage*, 2017. [1](#), [2](#), [4](#), [5](#)
- [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. *arXiv:1606.00915*, 2016. [1](#)
- [5] 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. [1](#), [2](#), [3](#), [4](#), [5](#)
- [6] Q. Dou, L. Yu, H. Chen, Y. Jin, X. Yang, J. Qin, and P.-A. Heng. 3D deeply supervised network for automated segmentation of volumetric medical images. *MIA*, 41:40–54, 2017. [3](#), [5](#), [7](#), [8](#)
- [7] Y. Fang, J. Xie, G. Dai, M. Wang, F. Zhu, T. Xu, and E. Wong. 3D deep shape descriptor. In *CVPR*, 2015. [1](#)
- [8] E. Gibson, F. Giganti, Y. Hu, E. Bonmati, S. Bandula, K. Gurusamy, B. Davidson, S. P. Pereira, M. J. Clarkson, and D. C. Barratt. Automatic multi-organ segmentation on abdominal ct with dense v-networks. *TMI*, 2018. [3](#)
- [9] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. [4](#)
- [10] P. Gravel, G. Beaudoin, and J. A. De Guise. A method for modeling noise in medical images. *TMI*, 23(10):1221–1232, 2004. [6](#)
- [11] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. C. Courville, Y. Bengio, C. Pal, P. Jodoin, and H. Larochelle. Brain tumor segmentation with deep neural networks. *MIA*, 35:18–31, 2017. [2](#)
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [2](#), [4](#)
- [13] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. [4](#)
- [14] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional Architecture for Fast Feature Embedding. *MM*, 2014. [6](#)
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. [1](#), [2](#)
- [16] A. A. Lasboo, P. Rezai, and V. Yaghmai. Morphological analysis of pancreatic cystic masses. *Academic radiology*, 17(3):348–351, 2010. [7](#)
- [17] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu. Deeply-supervised nets. In *AISTATS*, 2015. [3](#)
- [18] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. [1](#), [2](#)
- [19] J. Merkow, A. Marsden, D. Kriegman, and Z. Tu. Dense volume-to-volume vascular boundary detection. In *MICCAI*, 2016. [3](#), [5](#)
- [20] F. Milletari, N. Navab, and S.-A. Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 2016. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [21] P. Moeskops, J. M. Wolterink, B. H. M. van der Velden, K. G. A. Gilhuijs, T. Leiner, M. A. Viergever, and I. Išgum. Deep learning for multi-task medical image segmentation in multiple modalities. *CoRR*, abs/1704.03379, 2017. [2](#)
- [22] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. [4](#)
- [23] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. [2](#), [6](#)
- [24] H. Roth, M. Oda, N. Shimizu, H. Oda, Y. Hayashi, T. Kitasaka, M. Fujiwara, K. Misawa, and K. Mori. Towards dense volumetric pancreas segmentation in CT using 3D fully convolutional networks. In *SPIE*, 2018. [3](#)
- [25] 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. [1](#), [2](#), [6](#)
- [26] 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. [1](#), [2](#), [7](#)
- [27] W. Shen, B. Wang, Y. Jiang, Y. Wang, and A. L. Yuille. Multi-stage multi-recursive-input fully convolutional networks for neuronal boundary detection. In *ICCV*, pages 2410–2419, 2017. [1](#)
- [28] W. Shen, X. Wang, Y. Wang, X. Bai, and Z. Zhang. Deep-contour: A deep convolutional feature learned by positive-sharing loss for contour detection. In *CVPR*, 2015. [1](#)
- [29] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. [1](#), [2](#)
- [30] S. Xie and Z. Tu. Holistically-nested edge detection. In *ICCV*, 2015. [1](#), [2](#)
- [31] L. Yu, J.-Z. Cheng, Q. Dou, X. Yang, H. Chen, J. Qin, and P.-A. Heng. Automatic 3D cardiovascular MR segmentation with densely-connected volumetric convnets. In *MICCAI*, 2017. [1](#), [7](#)
- [32] L. Yu, X. Yang, H. Chen, J. Qin, and P. Heng. Volumetric convnets with mixed residual connections for automated prostate segmentation from 3D MR images. In *AAAI*, 2017. [5](#)
- [33] Y. Zhou, L. Xie, E. K. Fishman, and A. L. Yuille. Deep supervision for pancreatic cyst segmentation in abdominal CT scans. In *MICCAI*, 2017. [1](#), [6](#), [7](#)
- [34] 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 *MICCAI*, 2017. [1](#), [2](#), [6](#), [7](#)
- [35] Z. Zhu, X. Wang, S. Bai, C. Yao, and X. Bai. Deep learning representation using autoencoder for 3D shape retrieval. *Neurocomputing*, 204:41–50, 2016. [1](#)
- [36] Z. Zhu, Y. Xia, L. Xie, E. K. Fishman, and A. L. Yuille. Multi-scale coarse-to-fine segmentation for screening pancreatic ductal adenocarcinoma. *arXiv*, 2018. [8](#)