

Semi-Supervised 3D Abdominal Multi-Organ Segmentation via Deep Multi-Planar Co-Training

Yuyin Zhou¹ Yan Wang¹ Peng Tang² Song Bai³ Wei Shen¹ Elliot K. Fishman⁴ Alan Yuille¹

¹The Johns Hopkins University ²Huazhong University of Science and Technology

³University of Oxford ⁴The Johns Hopkins University School of Medicine

{zhouyuyiner, wyanny.9, tangpeng723, songbai.site, shenwei1231}@gmail.com

alan.l.yuille@gmail.com, efishman@jhmi.edu

Abstract

In multi-organ segmentation of abdominal CT scans, most existing fully supervised deep learning algorithms require lots of voxel-wise annotations, which are usually difficult, expensive, and slow to obtain. In comparison, massive unlabeled 3D CT volumes are usually easily accessible. Current mainstream works to address semi-supervised biomedical image segmentation problem are mostly graph-based. By contrast, deep network based semi-supervised learning methods have not drawn much attention in this field. In this work, we propose Deep Multi-Planar Co-Training (DMPCT), whose contributions can be divided into two folds: 1) The deep model is learned in a co-training style which can mine consensus information from multiple planes like the sagittal, coronal, and axial planes; 2) Multi-planar fusion is applied to generate more reliable pseudo-labels, which alleviates the errors occurring in the pseudo-labels and thus can help to train better segmentation networks. Experiments are done on our newly collected large dataset with 100 unlabeled cases as well as 210 labeled cases where 16 anatomical structures are manually annotated by four radiologists and confirmed by a senior expert. The results suggest that DMPCT significantly outperforms the fully supervised method by more than 4% especially when only a small set of annotations is used.

1. Introduction

Multi-organ segmentation of radiology images is a critical task which is essential to many clinical applications such as computer-aided diagnosis, computer-aided surgery, and radiation therapy. Compared with other internal human structures like brain or heart, segmenting abdominal organs appears to be much more challenging due to the low contrast and high variability of shape in CT images. In this paper, we focus on the problem of multi-organ segmentation

in abdominal regions, *e.g.*, liver, pancreas, kidney, *etc.*

Fully supervised approaches can usually achieve high accuracy with a large labeled training set which consists of pairs of radiology images as well as their corresponding pixel-wise label maps. However, it is quite time-consuming and costly to obtain such a large training set especially in the medical imaging domain due to the following reasons: 1) precise annotations of radiology images must be hand annotated by experienced radiologists and carefully checked by additional experts and 2) contouring organs or tissues in 3D volumes requires tedious manual input. By contrast, large unannotated datasets of CT images are much easier to obtain. Thereby our study mainly focuses on multi-organ segmentation in a semi-supervised fashion, *i.e.*, how to fully leverage unlabeled data to boost performance, so as to alleviate the need for such a large annotated training set.

In the biomedical imaging domain, traditional methods for semi-supervised learning usually adopt graph-based methods [14, 18] with a clustering assumption to segment pixels (voxels) into meaningful regions, *e.g.*, superpixels. These methods were studied for tissue or anatomical structures segmentation in 3D brain MR images, ultrasound images, *etc.* Other machine learning methods such as kernel-based large margin algorithms [27] have been suggested for white matter hyperintensities segmentation. Although widely applied to biomedical imaging segmentation tasks in the past decade, the traditional methods cannot always produce a satisfactory result due to the lack of advanced techniques.

With the recent advance of deep learning and its applications [23, 38, 37, 39], fully convolutional networks (FCNs) [24] have been successfully applied to many biomedical segmentation tasks such as neuronal structures segmentation [9, 13, 30, 35], single organ segmentation [32, 48, 47], and multi-organ segmentation [33, 42] in a fully supervised manner. Their impressive performances have shown that we are now equipped with much more powerful tech-

niques than traditional methods. Nevertheless, network-based semi-supervised learning for biomedical image segmentation has not drawn enough attention. The current usage of deep learning for semi-supervised multi-organ segmentation in the biomedical imaging domain is to train an FCN on both labeled and unlabeled data, and alternately update automated segmentations (pseudo-labels) for unlabeled data and the network parameters [5]. However, if an error occurs in the initial pseudo-label of the unlabeled data, the error will be reinforced by the network during the following iterations. How to improve the quality of pseudo-labels for unlabeled data hence becomes a promising direction to alleviate this negative effect.

In this paper, we exploit the fact that CT scans are high-resolution three-dimensional volumes which can be represented by multiple planes, *i.e.*, the axial, coronal, and sagittal planes. Taking advantages of this multi-view property, we propose Deep Multi-Planar Co-Training (DMPCT), a systematic EM-like semi-supervised learning framework. DMPCT consists of a teacher model, a multi-planar fusion module, and a student model. While the teacher model is trained from multiple planes separately in a slice-by-slice manner with a few annotations, the key advantage of DMPCT is that it enjoys the additional benefit of continuously generating more reliable pseudo-labels by the multi-planar fusion module, which can afterward help train the student model by making full usage of massive unlabeled data. As there are multiple segmentation networks corresponding to different planes in the teacher model and the student model, co-training [7, 26] is introduced so that these networks can be trained simultaneously in our unified framework and benefit from each other. We evaluate our algorithm on our newly collected large dataset and observe a significant improvement of 4.23% compared with the fully supervised method. At last, as DMPCT is a generic and flexible framework, it can be envisioned that better backbone models and fusion strategies can be easily plugged into our framework. Our unified system can be also practically useful for current clinical environments due to the efficiency in leveraging massive unlabeled data to boost segmentation performance.

2. Related Work

Fully-supervised multi-organ segmentation. Early studies of abdominal organ segmentation focused on atlas-based methods [22, 11, 43]. The frameworks are usually problematic because 1) they are not able to capture the large inter-subject variations of abdominal regions and 2) computational time is tightly dependent on the number of atlases. Recently, learning-based approaches with relatively large dataset have been introduced for multi-organ segmentation [17, 34, 8]. Especially, deep Convolutional Neural Networks (CNNs) based methods have achieved a great success

in the medical image segmentation [33, 10, 16, 41, 42, 46, 21] in the last few years. Compared with multi-atlas-based approaches, CNNs based methods are generally more efficient and accurate. CNNs based methods for multi-organ segmentation can be divided into two major categories: 3D CNNs [33, 10, 16] based and 2D CNNs [41, 42, 46, 21] based. 3D CNNs usually adopt the sliding-window strategy to avoid the *out of memory* problem, leading to high time complexity. Compared with 3D CNNs, 2D CNNs based algorithms can be directly end-to-end trained using 2D deep networks, which is less time-consuming.

Semi-supervised learning. The most commonly used techniques for semi-supervised learning include self-training [31, 28], co-training [7], multi-view learning [44] and graph-based methods [6, 40].

In self-training, the classifier is iteratively re-trained using the training set augmented by adding the unlabeled data with their own predictions. The procedure repeated until some convergence criteria are satisfied. In such case, one can imagine that a classification mistake can reinforce itself. Self-training has achieved great performances in many computer vision problems [31, 28] and recently has been applied to deep learning based semi-supervised learning in the biomedical imaging domain [5].

Co-training [7] assumes that (1) features can be split into two independent sets and (2) each sub-feature set is sufficient to train a good classifier. During the learning process, each classifier is retrained with the additional training examples given by the other classifier. Co-training utilizes multiple sets of independent features which describe the same data, and therefore tends to yield more accurate and robust results than self-training [36]. Multi-view learning [44], in general, defines learning paradigms that utilize the agreement among different learners. Co-training is one of the earliest schemes for multi-view learning.

Graph-based semi-supervised methods define a graph where the nodes are labeled and unlabeled examples in the dataset, and edges reflect the similarity of examples. These methods have been widely adopted in non-deep-learning based semi-supervised learning algorithms in the biomedical imaging domain [14, 18, 27].

Different from other methods, our work tactfully embeds the multi-view property of 3D medical data into the co-training framework, which is simple and effective.

3. Deep Multi-Planar Co-Training

We propose Deep Multi-Planar Co-Training (DMPCT), a semi-supervised multi-organ segmentation method which exploits multi-planar information to generate pseudo-labels for unlabeled 3D CT volumes. Assume that we are given a 3D CT volume dataset \mathcal{S} containing K organs. This includes labeled volumes $\mathcal{S}_L = \{(\mathbf{I}_m, \mathbf{Y}_m)\}_{m=1}^L$ and unlabeled




Figure 1. Illustration of the Deep Multi-Planar Co-Training (DMPCT) framework. (a) We first train a teacher model on the labeled dataset. (b) The trained model is used to assign pseudo-labels to the unlabeled data using our multi-planar fusion module as demonstrated in Figure 2. (c) Finally, we train a student model over the union of both the labeled and the unlabeled data. Step (b) and (c) are performed in an iterative manner.

labeled volumes $\mathcal{S}_L = \{\mathbf{I}_m\}_{m=l+1}^M$, where \mathbf{I}_m and \mathbf{Y}_m denote a 3D input volume and its corresponding ground-truth segmentation mask. l and $M-l$ are the numbers of labeled and unlabeled volumes, respectively. Typically $l \ll M$. As shown in Figure 1, DMPCT involves the following steps:

- **Step 1:** train a *teacher model* on the manually labeled data \mathcal{S}_L in the fully supervised setting (see Sec. 3.1).
- **Step 2:** the trained model is then used to assign pseudo-labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ to the unlabeled data \mathcal{S}_U by fusing the estimations from all planes (see Sec. 3.2).
- **Step 3:** train a *student model* on the union of the manually labeled data and automatically labeled data $\mathcal{S}_L \cup \{(\mathbf{I}_m, \hat{\mathbf{Y}}_m)\}_{m=l+1}^M$ (see Sec. 3.3).
- **Step 4:** perform step 2 & 3 in an iterative manner.

3.1. Teacher Model

We train the teacher model on the labeled dataset \mathcal{S}_L . By splitting each volume and its corresponding label mask from the sagittal (S), coronal (C), and axial (A) planes, we can get three sets of 2D slices, *i.e.*, $\mathcal{S}_L^V = \{(\mathbf{I}_n^V, \mathbf{Y}_n^V)\}_{n=1}^{N_V}$, $V \in \{S, C, A\}$, where N_V is the number of 2D slices obtained from plane V . We train a 2D-FCN model (we use [24] as our reference CNN model throughout this paper) to perform segmentation from each plane individually.

Without loss of generality, let $\mathbf{I}^V \in \mathbb{R}^{W \times H}$ and $\mathbf{Y}^V = \{y_i^V\}_{i=1}^{W \times H}$ denote a 2D slice and its corresponding label

mask in \mathcal{S}_L^V , where $y_i^V \in \{0, 1, \dots, K\}$ is the organ label (0 means background) of the i -th pixel in \mathbf{I}^V . Consider a segmentation model $\mathbb{M}^V : \hat{\mathbf{Y}} = \mathbf{f}(\mathbf{I}^V; \theta)$, where θ denotes the model parameters and $\hat{\mathbf{Y}}$ denotes the prediction for \mathbf{I}^V . Our objective function is

$$\mathcal{L}(\mathbf{I}^V, \mathbf{Y}^V; \theta) = -\frac{1}{W \times H} \left[\sum_{i=1}^{W \times H} \sum_{k=0}^K \mathbb{1}(y_i^V = k) \log p_{i,k}^V \right], \quad (1)$$

where $p_{i,k}^V$ denotes the probability of the i -th pixel been classified as label k on 2D slice \mathbf{I}^V and $\mathbb{1}(\cdot)$ is the indicator function. We train the teacher model by optimizing \mathcal{L} w.r.t. θ by stochastic gradient descent.

3.2. Multi-Planar Fusion Module

Given a well-trained teacher model $\{\mathbb{M}^V | V \in \{S, C, A\}\}$, our goal of the multi-planar fusion module is to generate the pseudo-labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ for the unlabeled data \mathcal{S}_U . We first make predictions on the 2D slices from each plane and then reconstruct the 3D volume by stacking all slices back together. Several previous studies [20, 29, 3, 4] suggest that combining predictions from multiple views can often improve the accuracy and the robustness of the final decision since complementary information can be exploited from multiple views simultaneously. Thereby, the fused prediction from multiple planes is superior to any estimation of a single plane. The overall module is shown in Figure 2.

More specifically, majority voting is applied to fuse the

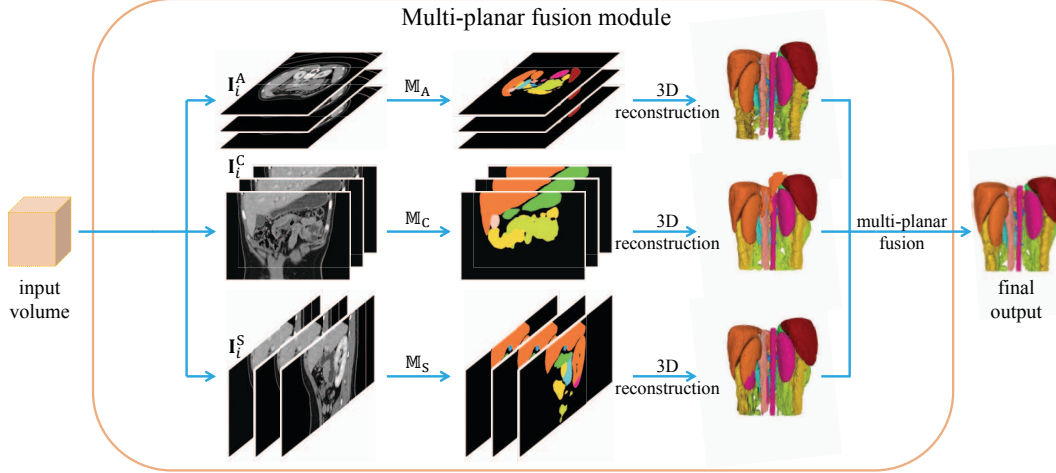


Figure 2. Illustration of the multi-planar fusion module, where the input 3D volume is first parsed into 3 sets of slices along the sagittal, coronal, and axial planes to be evaluated respectively. Then the final 3D estimation is obtained by fusing predictions from each individual plane.

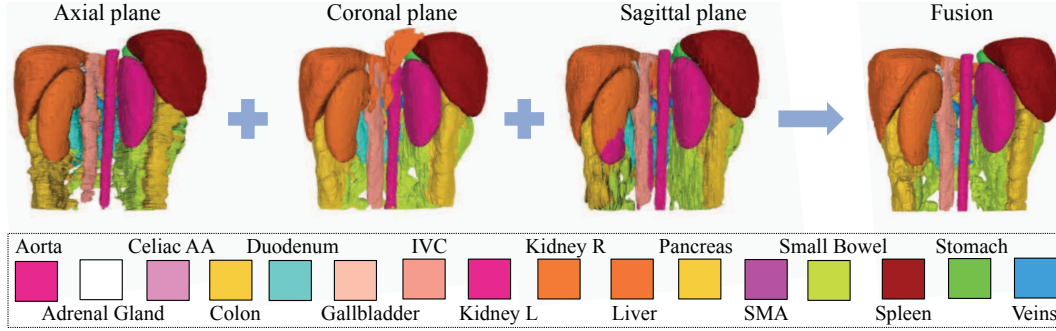


Figure 3. An example of 3D predictions reconstructed from the sagittal, coronal, and axial planes as well as their fusion output. Estimations from single planes are already reasonably well, whereas the single fusion outcome is superior to estimation from any single plane.

hard estimations by seeking an agreement among different planes. If the predictions from all planes do not agree on a voxel, then we select the prediction for that voxel with the maximum confidence. As simple as this strategy might sound, this method has been shown to result in highly robust and efficient outcome in various previous studies [1, 20, 29, 48]. The final decision for the i -th voxel y_i^* of $\hat{\mathbf{Y}}_m$ is:

$$y_i^* = \begin{cases} y_i^V, & \text{if } \exists V, V' \in \{S, C, A\}, V \neq V' \mid y_i^V = y_i^{V'} \\ y_i^{V^*}, & \text{otherwise} \end{cases}, \quad (2)$$

where $V^* = \arg \max_{V \in \{S, C, A\}} \max_j p_{i,j}^V$, $p_{i,j}^S$, $p_{i,j}^C$, and $p_{i,j}^A$ denote the probabilities of the i -th pixel classified as label j from the sagittal, coronal, and axial planes, respectively. y_i^V denotes the hard estimation for the i -th pixel on plane V , i.e., $y_i^V = \arg \max_j p_{i,j}^V$.

As shown in Figure 3, our multi-planar fusion module improves both over- and under-estimation by fusing aspects from different planes and therefore yields a much better

outcome. Note that other rules [2, 41] can also be easily adapted to this module. We do not focus on discussing the influence of the fusion module in this paper, although intuitively better fusion module should lead to higher performance.

3.3. Student Model

After generating the pseudo-labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ for the unlabeled dataset \mathcal{S}_U , the training set can be then enlarged by taking the union of both the labeled and the unlabeled dataset, i.e., $\mathcal{S} = \mathcal{S}_L \cup \{(\mathbf{I}_m, \hat{\mathbf{Y}}_m)\}_{m=l+1}^M$. The student model is trained on this augmented dataset \mathcal{S} the same way we train the teacher model as described in Sec. 3.1. The overall training procedure is summarized in Algorithm 1. In the training stage, we first train a teacher model in a supervised manner and then use it to generate the pseudo-labels for the unlabeled dataset. Then we alternate the training of the student model and the pseudo-label generation procedures in an iterative manner to optimize the student model T times. In the testing stage, we follow the method in Sec.

Algorithm 1 Deep Multi-Planar Co-Training for Multi-Organ Segmentation

Require: A set of labeled data $\mathcal{S}_L = \{(\mathbf{I}_m, \mathbf{Y}_m)\}_{m=1}^l$ and unlabeled volumes $\mathcal{S}_U = \{\mathbf{I}_m\}_{m=l+1}^M$.

Ensure: A trained multi-organ segmentation model $\{\mathbb{M}^S, \mathbb{M}^C, \mathbb{M}^A\}$.

$\mathcal{S} \leftarrow \mathcal{S}_L$

for $t = 1$ to T **do**

 Parse \mathcal{S} into $\mathcal{S}^S, \mathcal{S}^C, \mathcal{S}^A$.

 Train $\mathbb{M}^S, \mathbb{M}^C$, and \mathbb{M}^A on $\mathcal{S}^S, \mathcal{S}^C$, and \mathcal{S}^A respectively.

 Generate pseudo-class labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ for the unlabeled dataset \mathcal{S}_U by Eq. 2

 Augment the training set \mathcal{S} by adding the self-labeled examples to \mathcal{S}_L , *i.e.*, $\mathcal{S} = \mathcal{S}_L \cup \{(\mathbf{I}_m, \hat{\mathbf{Y}}_m)\}_{m=l+1}^M$.

end for

 Parse \mathcal{S} into $\mathcal{S}^S, \mathcal{S}^C, \mathcal{S}^A$.

 Train $\mathbb{M}^S, \mathbb{M}^C$, and \mathbb{M}^A on $\mathcal{S}^S, \mathcal{S}^C$, and \mathcal{S}^A respectively.

3.2 to generate the final estimation using the T -th student model.

4. Experiments

4.1. Dataset and Evaluation

Our fully-labeled dataset includes 210 contrast-enhanced abdominal clinical CT images in the portal venous phase, in which we randomly choose 50/30/80 patients for training, validation, and testing, unless otherwise specified. A total of 16 structures (Aorta, Adrenal gland, Celiac AA, Colon, Duodenum, Gallbladder, Interior Vena Cava (IVC), Kidney (left, right), Liver, Pancreas, Superior Mesenteric Artery (SMA), Small bowel, Spleen, Stomach, Veins) for each case were segmented by four experienced radiologists, and confirmed by an independent senior expert. Our unlabeled dataset consists of 100 unlabeled cases acquired from a local hospital. To the best of our knowledge, this is the largest abdominal CT dataset with the most number of organs segmented. Each CT volume consists of $319 \sim 1051$ slices of 512×512 pixels, and have voxel spatial resolution of $([0.523 \sim 0.977] \times [0.523 \sim 0.977] \times 0.5)\text{mm}^3$. The metric we use is the Dice-Sørensen Coefficient (DSC), which measures the similarity between the prediction voxel set \mathcal{Z} and the ground-truth set \mathcal{Y} , with the mathematical form of $\text{DSC}(\mathcal{Z}, \mathcal{Y}) = \frac{2 \times |\mathcal{Z} \cap \mathcal{Y}|}{|\mathcal{Z}| + |\mathcal{Y}|}$. For each organ, we report an average DSC together with the standard deviation over all the testing cases.

4.2. Implementation Details

We set the learning rate to be 10^{-9} . The teacher model and the student model are trained for 80,000 and 160,000 iterations respectively. The validation set is used for tuning

the hyper-parameters. Similar to [19], we use three windows of $[-125, 275]$, $[-160, 240]$, and $[-1000, 1000]$ Housefield Units as the three input channels respectively. The intensities of each slice are rescaled to $[0.0, 1.0]$. Similar to [48, 45, 41], we initialize the network parameters θ by using the FCN-8s model [24] pre-trained on the PascalVOC image segmentation dataset. The iteration number T in Algorithm 1 is set to 2, *i.e.*, $T = 2$, as the performance of the validation set gets saturated.

4.3. Comparison with the Baseline

We show that our proposed DMPCT works better than other methods: 1) fully supervised learning method [24] (denoted as FCN), and 2) single planar based semi-supervised learning approach [5] (denoted as SPSL). Both 1) and 2) are applied on each individual plane separately, and then the final result is obtained via multi-planar fusion (see Sec 3.2). As shown in Table 1, with 50 labeled data, by varying the number of unlabeled data from 0 to 100, the average DSC of DMPCT increases from 73.71% to 77.94% and the standard deviation decreases from 9.97% to 8.51%. Compared with SPSL, our proposed DMPCT can boost the performance in both settings (*i.e.*, 50 labeled data + 50 unlabeled data and 50 labeled data + 100 unlabeled data). Besides, the p -values for testing significant difference between our DMPCT (50 labeled data + 100 unlabeled data) and FCN (50 labeled data + 0 unlabeled data) for 16 organs are shown in the last column of Table 1, which suggests significant statistical improvements among almost all organs. Figure 4 shows comparison results of our DMPCT and the fully supervised method by box plots.

It is noteworthy that greater improvements are observed especially for those difficult organs, *i.e.*, organs either small in sizes or with complex geometric characteristics. Table 1 indicates that our DMPCT approach boosts the segmentation performance of these small hard organs by 5.54% (Pancreas), 8.72% (Colon), 9.89% (Duodenum), 8.12% (Small bowels) and 5.48% (Veins), 5.76% (IVC). This promising result indicates that our method distills a reasonable amount of knowledge from the unlabeled data. An example is shown in Figure 5. In this particular case, the DSCs for Celiac AA, Colon, Duodenum, IVC, Pancreas and Veins are boosted from 60.13%, 46.79%, 71.08%, 69.23%, 63.48% to 79.45%, 83.81%, 77.59%, 74.75%, 75.31% respectively.

4.4. Discussion

4.4.1 Amount of labeled data

For ablation analysis, we enlarge the labeled training set to 100 cases and keep the rest of the settings the same. As shown in Figure 6, with more labeled data, the semi-supervised methods (DMPCT, SPSL) still obtain better performance than the supervised method (FCN), while the performance gain becomes less prominent. This is probably

Table 1. The comparison of segmentation accuracy (DSC, %) by using 50 labeled data and varying the number of unlabeled data (e.g., 50-0 indicates 50 labeled data and 0 unlabeled data). We report the mean and standard deviation over 80 cases. The p -values for testing significant difference between DMPCT (50-100) and FCN (50-0) are shown. Significant statistical improvement is shown in *italic* with $p < 0.05$. See Section 4.3 for definitions of FCN, SPSL, and DMPCT (Ours).

Organ Type	FCN	SPSL		DMPCT (Ours)		p -value
	50 - 0	50 - 50	50 - 100	50 - 50	50 - 100	
Aorta	89.14 \pm 7.95	91.10 \pm 5.52	90.76 \pm 5.90	91.43 \pm 4.88	91.54 \pm 4.65	3.32×10^{-5}
Adrenal gland	26.45 \pm 12.1	29.92 \pm 14.7	26.93 \pm 15.6	30.58 \pm 12.7	35.48 \pm 11.8	1.98×10^{-15}
Celiac AA	35.01 \pm 19.7	37.27 \pm 19.0	39.78 \pm 18.4	36.25 \pm 20.5	40.50 \pm 18.9	1.00×10^{-5}
Colon	71.81 \pm 14.9	78.28 \pm 13.0	79.58 \pm 12.9	79.61 \pm 12.3	80.53 \pm 11.6	7.69×10^{-12}
Duodenum	54.89 \pm 15.5	57.77 \pm 17.3	62.22 \pm 14.8	66.95 \pm 12.6	64.78 \pm 13.8	1.95×10^{-19}
Gallbladder	86.53 \pm 6.21	87.87 \pm 5.45	88.02 \pm 5.83	88.45 \pm 5.07	87.77 \pm 6.29	0.002
IVC	77.67 \pm 9.49	81.28 \pm 8.87	82.63 \pm 7.31	83.49 \pm 6.94	83.43 \pm 7.02	9.30×10^{-14}
Kidney (L)	95.12 \pm 5.01	95.59 \pm 4.97	95.88 \pm 3.68	95.82 \pm 3.60	96.09 \pm 3.42	3.69×10^{-6}
Kidney (R)	95.69 \pm 2.36	95.77 \pm 4.93	96.14 \pm 2.94	96.17 \pm 2.75	96.26 \pm 2.29	1.74×10^{-7}
Liver	95.45 \pm 2.41	96.06 \pm 0.99	96.07 \pm 1.03	96.11 \pm 0.97	96.15 \pm 0.92	0.005
Pancreas	76.49 \pm 11.6	80.12 \pm 7.52	80.93 \pm 6.84	81.46 \pm 6.32	82.03 \pm 6.16	2.97×10^{-8}
SMA	52.26 \pm 17.1	51.81 \pm 18.2	51.94 \pm 17.1	49.40 \pm 19.2	52.70 \pm 17.7	0.667
Small bowel	71.13 \pm 13.1	78.93 \pm 12.6	79.97 \pm 12.8	79.49 \pm 12.1	79.25 \pm 12.6	2.53×10^{-22}
Spleen	94.81 \pm 2.64	95.46 \pm 2.09	95.58 \pm 1.90	95.73 \pm 2.03	95.98 \pm 1.59	1.83×10^{-10}
Stomach	91.38 \pm 3.94	92.62 \pm 3.71	92.92 \pm 3.65	93.33 \pm 3.47	93.42 \pm 3.21	3.30×10^{-23}
Veins	64.75 \pm 15.4	70.43 \pm 14.3	69.66 \pm 14.6	69.82 \pm 14.5	70.23 \pm 14.4	4.16×10^{-15}
Mean	73.71 \pm 9.97	76.32 \pm 9.58	76.87 \pm 9.08	77.20 \pm 8.75	77.94 \pm 8.51	4.74×10^{-90}




Figure 4. Performance comparison (DSC, %) in box plots of 16 organs by using 50 labeled data and varying the number of unlabeled data (e.g., 50-0 indicates 50 labeled data and 0 unlabeled data). See Section 4.3 for definitions of FCN and DMPCT (Ours).

because the network is already trained well when large training set is available. We believe that if much more unlabeled data can be provided the performance should go up considerably. In addition, we find that DMPCT outperforms SPSL in every setting, which further demonstrates the usefulness of multi-planar fusion in our co-training framework.

4.4.2 Comparison with 3D network-based self-training

Various previous studies [25, 41] demonstrate that 2D multi-planar fusion outperforms directly 3D learning in the fully supervised setting. 3D CNNs come with an increased number of parameters, significant memory and computational requirements. Due to GPU memory restrictions,

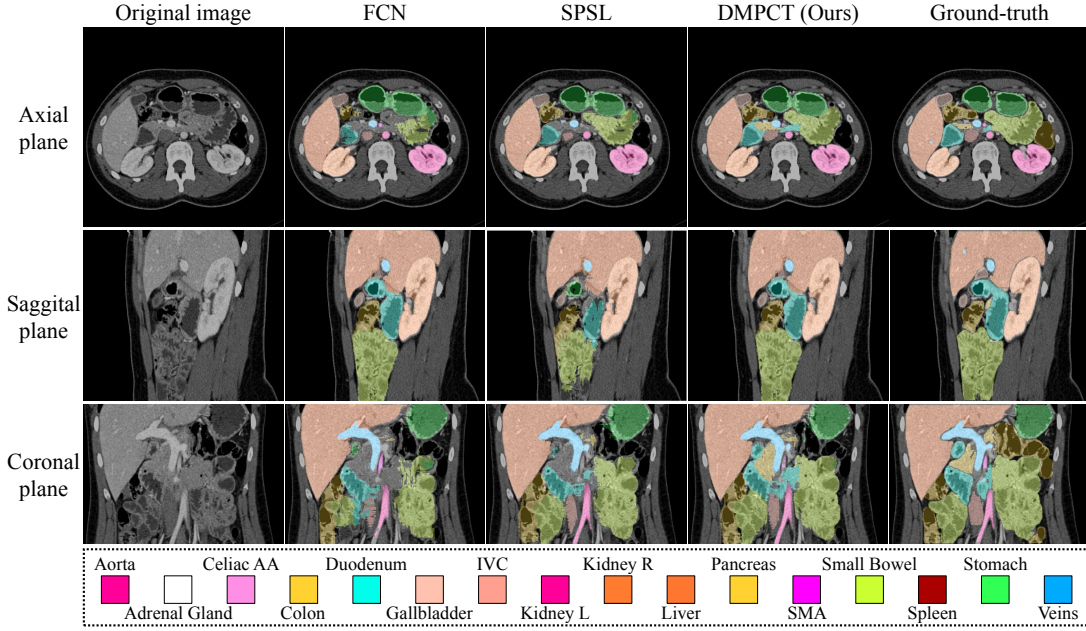


Figure 5. Comparisons among FCN, SPSL, and DMPCT (Ours) viewed from multiple planes. 50 labeled cases are used for all methods. 100 unlabeled cases are used for the SPSL and DMPCT. For this particular case, FCN obtains an average DSC of 72.75%, SPSL gets 78.87%, and DMPCT (Ours) gets 80.75%. See Section 4.3 for definitions of FCN, SPSL, and DMPCT (Ours). Best viewed in color.

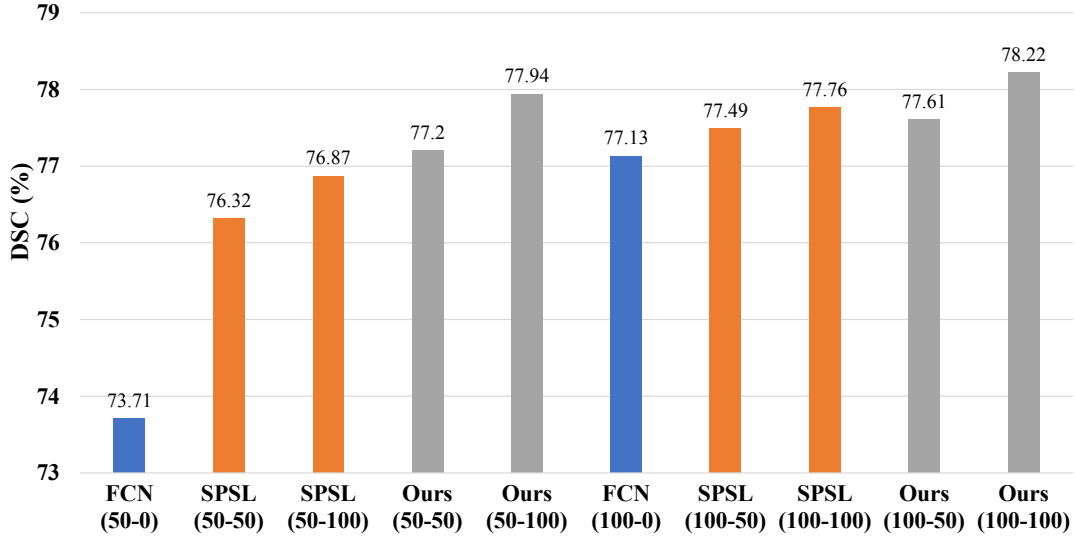


Figure 6. Ablation study on numbers of labeled data and unlabeled data. Mean DSC of all testing cases under all settings (e.g., 50-0 indicates 50 labeled data and 0 unlabeled data). See Section 4.3 for definitions of FCN, SPSL, and DMPCT (Ours).

these 3D CNN approaches which adopt the sliding-window strategy do not act on the entire 3D CT volume, but instead on local 3D patches [12, 15, 17]. This results in the lack of holistic information and low efficiency. In order to prove that DMPCT outperforms direct 3D learning in the semi-supervised setting, we also implement a patch-based 3D UNet [12]. 3D UNet gets 69.66% in terms of mean DSC using 50 labeled data. When adding 100 unlabeled data the

performance even drops to 65.21%. This clearly shows that in 3D learning the teacher model is not trained well, thus the errors of the pseudo-labels are reinforced during student model training.

Table 2. Cross-dataset generalization results.

Organ	Spleen	Kidney (R)	Kidney (L)	Gall Bladder	Liver
FCN	71.85 \pm 26.13	54.44 \pm 20.04	54.98 \pm 26.63	48.13 \pm 26.07	85.46 \pm 16.81
DMPCT (Ours)	83.68 \pm 16.53	71.36 \pm 20.85	69.95 \pm 20.50	60.05 \pm 26.91	92.11 \pm 6.46
Organ	Stomach	Aorta	IVC	Veins	Pancreas
FCN	38.89 \pm 23.86	70.43 \pm 19.70	53.67 \pm 18.40	35.54 \pm 18.94	39.40 \pm 25.34
DMPCT (Ours)	54.78 \pm 26.57	76.05 \pm 15.99	68.18 \pm 14.58	37.52 \pm 15.86	60.05 \pm 16.61

4.4.3 Comparison with traditional co-training

In order to show that our DMPCT outperforms traditional co-training algorithm [7], we also select only the most confident samples during each iteration. Here the confidence score is measured by the entropy of probability distribution for each voxel in one slice. Under the setting of 50 labeled cases and 50 unlabeled cases, we select top 5000 samples with the highest confidence in each iteration. The whole training process takes about 6-7 iterations for each plane. The complete training requires more than 50 hours. Compared with our approach, this method requires much more time to converge. It obtains a mean DSC of 76.52%, slightly better than SPSL but worse than our DMPCT, which shows that selecting the most confident samples during training may not be a wise choice for deep network based semi-supervised learning due to its low efficiency.

4.4.4 Cross dataset generalization

We apply our trained DMPCT model (50 labeled data + 100 unlabeled data) and baseline FCN model (50 labeled data + 0 unlabeled data) on a public available abdominal CT datasets¹ with 13 anatomical structures labeled *without any further re-training* on new data cases. 10 out of the 13 structures are evaluated which are also manually annotated in our own dataset and we find that our proposed method improves the overall mean DSC and also reduces the standard deviation significantly, as shown in Table 2. The overall mean DSC as well as the standard deviation for the 10 organs is improved from 59.23 \pm 22.20% to 67.38 \pm 19.64%. We also directly test our models on the NIH pancreas segmentation dataset of 82 cases² and observe that our DMPCT model achieves an average DSC of 66.16%, outperforming the fully supervised method, with an average DSC of 58.73%, by more than 7%. This may demonstrate that our approach, which leverages more unlabeled data from multiple planes, turns out to be much more generalizable than the baseline model.

¹30 training data sets at <https://www.synapse.org/#!/Synapse:syn3193805/wiki/217789>

²<https://wiki.cancerimagingarchive.net/display/Public/Pancreas-CT>

4.4.5 Computation time

In our experiments, the teacher model training process takes about 4.94 hours on an NVIDIA TITAN Xp GPU card for 80,000 iterations over all the training cases. The average computation time for generating pseudo-label as well as testing per volume depends on the volume of the target structure, and the average computation time for 16 organs is approximately 4.5 minutes, which is comparable to other recent methods [48, 32] even for single structure inference. The student model training process takes about 9.88 hours for 160,000 iterations.

5. Conclusion

In this paper, we designed a systematic framework DMPCT for multi-organ segmentation in abdominal CT scans, which is motivated by the traditional co-training strategy to incorporate multi-planar information for the unlabeled data during training. The pseudo-labels are iteratively updated by inferencing comprehensively on multiple configurations of unlabeled data with a multi-planar fusion module. We evaluate our approach on our own large newly collected high-quality dataset. The results show that 1) our method outperforms the fully supervised learning approach by a large margin; 2) it outperforms the single planar method, which further demonstrates the benefit of multi-planar fusion; 3) it can learn better if more unlabeled data provided especially when the scale of labeled data is small. Our framework can be practical in assisting radiologists for clinical applications since the annotation of multiple organs in 3D volumes requires massive labor from radiologists. Our framework is not specific to a certain structure, but shows robust results in multiple complex anatomical structures within efficient computational time. We believe that our algorithm may achieve even higher accuracy if a more powerful backbone network or an advanced fusion algorithm is employed, which we leave as the future work.

Acknowledgement. This work was supported by the Lustgarten Foundation for Pancreatic Cancer Research and also supported by NSFC No. 61672336. We thank Prof. Seyoun Park, Dr. Lingxi Xie, Cihang Xie, Zhishuai Zhang, Fengze Liu, Zhuotun Zhu and Yingda Xia for instructive discussions.

References

- [1] P. Aljabar, R. A. Heckemann, A. Hammers, J. V. Hajnal, and D. Rueckert. Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. *NeuroImage*, 46(3):726–738, 2009. 4
- [2] A. J. Asman and B. A. Landman. Non-local statistical label fusion for multi-atlas segmentation. *MIA*, 17(2):194–208, 2013. 4
- [3] S. Bai, X. Bai, Z. Zhou, Z. Zhang, and L. J. Latecki. Gift: A real-time and scalable 3d shape search engine. In *CVPR*, pages 5023–5032, 2016. 3
- [4] S. Bai, X. Bai, Z. Zhou, Z. Zhang, Q. Tian, and L. J. Latecki. Gift: Towards scalable 3d shape retrieval. *TMM*, 19(6):1257–1271, 2017. 3
- [5] W. Bai, O. Oktay, M. Sinclair, H. Suzuki, M. Rajchl, G. Tarroni, B. Glocker, A. King, P. M. Matthews, and D. Rueckert. Semi-supervised learning for network-based cardiac mr image segmentation. In *MICCAI*, pages 253–260, 2017. 2, 5
- [6] A. Blum and S. Chawla. Learning from labeled and unlabeled data using graph mincuts. 2001. 2
- [7] A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, pages 92–100, 1998. 2, 8
- [8] T. Brosch and A. Saalbach. Foveal fully convolutional nets for multi-organ segmentation. In *Medical Imaging 2018: Image Processing*, volume 10574, page 105740U, 2018. 2
- [9] H. Chen, X. Qi, J.-Z. Cheng, P.-A. Heng, et al. Deep contextual networks for neuronal structure segmentation. In *AAAI*, pages 1167–1173, 2016. 1
- [10] S. Chen, H. Roth, S. Dorn, M. May, A. Cavallaro, M. M. Lell, M. Kachelrieß, H. Oda, K. Mori, and A. Maier. Towards automatic abdominal multi-organ segmentation in dual energy ct using cascaded 3d fully convolutional network. *arXiv preprint arXiv:1710.05379*, 2017. 2
- [11] C. Chu, M. Oda, T. Kitasaka, K. Misawa, M. Fujiwara, Y. Hayashi, Y. Nimura, D. Rueckert, and K. Mori. Multi-organ Segmentation based on Spatially-Divided Probabilistic Atlas from 3D Abdominal CT Images. In *MICCAI*, pages 165–172, 2013. 2
- [12] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432, 2016. 7
- [13] D. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. In *NIPS*, pages 2843–2851, 2012. 1
- [14] A. Ciurte, X. Bresson, O. Cuisenaire, N. Houhou, S. Nedevschi, J.-P. Thiran, and M. B. Cuadra. Semi-supervised segmentation of ultrasound images based on patch representation and continuous min cut. *PloS one*, 9(7):e100972, 2014. 1, 2
- [15] Q. Dou, H. Chen, Y. Jin, L. Yu, J. Qin, and P.-A. Heng. 3d deeply supervised network for automatic liver segmentation from ct volumes. In *MICCAI*, pages 149–157, 2016. 7
- [16] 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. 2
- [17] E. Gibson, F. Giganti, Y. Hu, E. Bonmati, S. Bandula, K. Gurusamy, B. R. Davidson, S. P. Pereira, M. J. Clarkson, and D. C. Barratt. Towards image-guided pancreas and biliary endoscopy: automatic multi-organ segmentation on abdominal ct with dense dilated networks. In *MICCAI*, pages 728–736, 2017. 2, 7
- [18] L. Gu, Y. Zheng, R. Bise, I. Sato, N. Imanishi, and S. Aiso. Semi-supervised learning for biomedical image segmentation via forest oriented super pixels (voxels). In *MICCAI*, pages 702–710, 2017. 1, 2
- [19] A. P. Harrison, Z. Xu, K. George, L. Lu, R. M. Summers, and D. J. Mollura. Progressive and multi-path holistically nested neural networks for pathological lung segmentation from ct images. In *MICCAI*, pages 621–629, 2017. 5
- [20] R. A. Heckemann, J. V. Hajnal, P. Aljabar, D. Rueckert, and A. Hammers. Automatic anatomical brain mri segmentation combining label propagation and decision fusion. *NeuroImage*, 33(1):115–126, 2006. 3, 4
- [21] P. Hu, F. Wu, J. Peng, Y. Bao, F. Chen, and D. Kong. Automatic abdominal multi-organ segmentation using deep convolutional neural network and time-implicit level sets. *International journal of computer assisted radiology and surgery*, 12(3):399–411, 2017. 2
- [22] J. E. Iglesias and M. R. Sabuncu. Multi-atlas segmentation of biomedical images: A survey. *MIA*, 24(1):205–219, 2015. 2
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 1
- [24] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 1, 3, 5
- [25] A. Prason, K. Petersen, C. Igel, F. Lauze, E. Dam, and M. Nielsen. Deep Feature Learning for Knee Cartilage Segmentation Using a Triplanar Convolutional Neural Network. *MICCAI*, 2013. 6
- [26] S. Qiao, W. Shen, Z. Zhang, B. Wang, and A. Yuille. Deep co-training for semi-supervised image recognition. In *ECCV*, 2018. 2
- [27] C. Qin, R. G. Moreno, C. Bowles, C. Ledig, P. Scheltens, F. Barkhof, H. Rhodius-Meester, B. M. Tijms, A. W. Lemstra, W. M. van der Flier, B. Glocker, and D. Rueckert. A semi-supervised large margin algorithm for white matter hyperintensity segmentation. In *MICCAI*, pages 104–112, 2016. 1, 2
- [28] I. Radosavovic, P. Dollr, R. Girshick, G. Gkioxari, and K. He. Data distillation: Towards omni-supervised learning. In *CVPR*, 2018. 2
- [29] T. Rohlfing, R. Brandt, R. Menzel, and C. R. Maurer Jr. Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. *NeuroImage*, 21(4):1428–1442, 2004. 3, 4
- [30] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 1

- [31] C. Rosenberg, M. Hebert, and H. Schneiderman. Semi-supervised self-training of object detection models. In *WACV/MOTION*, pages 29–36, 2005. 2
- [32] H. R. Roth, L. Lu, N. Lay, A. P. Harrison, A. Farag, A. Sohn, and R. M. Summers. Spatial aggregation of holistically-nested convolutional neural networks for automated pancreas localization and segmentation. *MIA*, 45:94 – 107, 2018. 1, 8
- [33] H. R. Roth, H. Oda, Y. Hayashi, M. Oda, N. Shimizu, M. Fujiwara, K. Misawa, and K. Mori. Hierarchical 3d fully convolutional networks for multi-organ segmentation. *arXiv preprint arXiv:1704.06382*, 2017. 1, 2
- [34] H. R. Roth, C. Shen, H. Oda, T. Sugino, M. Oda, Y. Hayashi, K. Misawa, and K. Mori. A multi-scale pyramid of 3d fully convolutional networks for abdominal multi-organ segmentation. *arXiv preprint arXiv:1806.02237*, 2018. 2
- [35] W. Shen, B. Wang, Y. Jiang, Y. Wang, and A. Yuille. Multi-stage multi-recursive-input fully convolutional networks for neuronal boundary detection. In *ICCV*, pages 2391–2400, 2017. 1
- [36] R. T. Sousa and J. Gama. Comparison between co-training and self-training for single-target regression in data streams using amrules. 2017. 2
- [37] P. Tang, X. Wang, S. Bai, W. Shen, X. Bai, W. Liu, and A. L. Yuille. Pcl: Proposal cluster learning for weakly supervised object detection. *TPAMI*, 2018. 1
- [38] P. Tang, X. Wang, X. Bai, and W. Liu. Multiple instance detection network with online instance classifier refinement. In *CVPR*, pages 2843–2851, 2017. 1
- [39] P. Tang, X. Wang, A. Wang, Y. Yan, W. Liu, J. Huang, and A. Yuille. Weakly supervised region proposal network and object detection. In *ECCV*, pages 370–386, 2018. 1
- [40] J. Wang, T. Jebara, and S.-F. Chang. Semi-supervised learning using greedy max-cut. *JMLR*, 14(Mar):771–800, 2013. 2
- [41] Y. Wang, Y. Zhou, W. Shen, S. Park, E. K. Fishman, and A. L. Yuille. Abdominal multi-organ segmentation with organ-attention networks and statistical fusion. *arXiv preprint arXiv:1804.08414*, 2018. 2, 4, 5, 6
- [42] Y. Wang, Y. Zhou, P. Tang, W. Shen, E. K. Fishman, and A. L. Yuille. Training multi-organ segmentation networks with sample selection by relaxed upper confident bound. In *MICCAI*, 2018. 1, 2
- [43] R. Wolz, C. Chu, K. Misawa, M. Fujiwara, K. Mori, and D. Rueckert. Automated abdominal multi-organ segmentation with subject-specific atlas generation. *TMI*, 32(9):1723–1730, 2013. 2
- [44] C. Xu, D. Tao, and C. Xu. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*, 2013. 2
- [45] 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 *CVPR*, 2018. 5
- [46] X. Zhou, T. Ito, R. Takayama, S. Wang, T. Hara, and H. Fujita. Three-dimensional ct image segmentation by combining 2d fully convolutional network with 3d majority voting. In *Deep Learning and Data Labeling for Medical Applications*, pages 111–120. 2016. 2
- [47] 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
- [48] 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*, pages 693–701, 2017. 1, 4, 5, 8

Semi-Supervised 3D Abdominal Multi-Organ Segmentation via Deep Multi-Planar Co-Training

Yuyin Zhou¹ Yan Wang¹ Peng Tang² Song Bai³ Wei Shen¹ Elliot K. Fishman⁴ Alan Yuille¹

¹The Johns Hopkins University ²Huazhong University of Science and Technology

³University of Oxford ⁴The Johns Hopkins University School of Medicine

{zhouyuyiner, wyanny.9, tangpeng723, songbai.site, shenwei1231}@gmail.com

alan.l.yuille@gmail.com, efishman@jhmi.edu

Abstract

In multi-organ segmentation of abdominal CT scans, most existing fully supervised deep learning algorithms require lots of voxel-wise annotations, which are usually difficult, expensive, and slow to obtain. In comparison, massive unlabeled 3D CT volumes are usually easily accessible. Current mainstream works to address semi-supervised biomedical image segmentation problem are mostly graph-based. By contrast, deep network based semi-supervised learning methods have not drawn much attention in this field. In this work, we propose Deep Multi-Planar Co-Training (DMPCT), whose contributions can be divided into two folds: 1) The deep model is learned in a co-training style which can mine consensus information from multiple planes like the sagittal, coronal, and axial planes; 2) Multi-planar fusion is applied to generate more reliable pseudo-labels, which alleviates the errors occurring in the pseudo-labels and thus can help to train better segmentation networks. Experiments are done on our newly collected large dataset with 100 unlabeled cases as well as 210 labeled cases where 16 anatomical structures are manually annotated by four radiologists and confirmed by a senior expert. The results suggest that DMPCT significantly outperforms the fully supervised method by more than 4% especially when only a small set of annotations is used.

1. Introduction

Multi-organ segmentation of radiology images is a critical task which is essential to many clinical applications such as computer-aided diagnosis, computer-aided surgery, and radiation therapy. Compared with other internal human structures like brain or heart, segmenting abdominal organs appears to be much more challenging due to the low contrast and high variability of shape in CT images. In this paper, we focus on the problem of multi-organ segmentation

in abdominal regions, *e.g.*, liver, pancreas, kidney, *etc.*

Fully supervised approaches can usually achieve high accuracy with a large labeled training set which consists of pairs of radiology images as well as their corresponding pixel-wise label maps. However, it is quite time-consuming and costly to obtain such a large training set especially in the medical imaging domain due to the following reasons: 1) precise annotations of radiology images must be hand annotated by experienced radiologists and carefully checked by additional experts and 2) contouring organs or tissues in 3D volumes requires tedious manual input. By contrast, large unannotated datasets of CT images are much easier to obtain. Thereby our study mainly focuses on multi-organ segmentation in a semi-supervised fashion, *i.e.*, how to fully leverage unlabeled data to boost performance, so as to alleviate the need for such a large annotated training set.

In the biomedical imaging domain, traditional methods for semi-supervised learning usually adopt graph-based methods [14, 18] with a clustering assumption to segment pixels (voxels) into meaningful regions, *e.g.*, superpixels. These methods were studied for tissue or anatomical structures segmentation in 3D brain MR images, ultrasound images, *etc.* Other machine learning methods such as kernel-based large margin algorithms [27] have been suggested for white matter hyperintensities segmentation. Although widely applied to biomedical imaging segmentation tasks in the past decade, the traditional methods cannot always produce a satisfactory result due to the lack of advanced techniques.

With the recent advance of deep learning and its applications [23, 38, 37, 39], fully convolutional networks (FCNs) [24] have been successfully applied to many biomedical segmentation tasks such as neuronal structures segmentation [9, 13, 30, 35], single organ segmentation [32, 48, 47], and multi-organ segmentation [33, 42] in a fully supervised manner. Their impressive performances have shown that we are now equipped with much more powerful tech-

niques than traditional methods. Nevertheless, network-based semi-supervised learning for biomedical image segmentation has not drawn enough attention. The current usage of deep learning for semi-supervised multi-organ segmentation in the biomedical imaging domain is to train an FCN on both labeled and unlabeled data, and alternately update automated segmentations (pseudo-labels) for unlabeled data and the network parameters [5]. However, if an error occurs in the initial pseudo-label of the unlabeled data, the error will be reinforced by the network during the following iterations. How to improve the quality of pseudo-labels for unlabeled data hence becomes a promising direction to alleviate this negative effect.

In this paper, we exploit the fact that CT scans are high-resolution three-dimensional volumes which can be represented by multiple planes, *i.e.*, the axial, coronal, and sagittal planes. Taking advantages of this multi-view property, we propose Deep Multi-Planar Co-Training (DMPCT), a systematic EM-like semi-supervised learning framework. DMPCT consists of a teacher model, a multi-planar fusion module, and a student model. While the teacher model is trained from multiple planes separately in a slice-by-slice manner with a few annotations, the key advantage of DMPCT is that it enjoys the additional benefit of continuously generating more reliable pseudo-labels by the multi-planar fusion module, which can afterward help train the student model by making full usage of massive unlabeled data. As there are multiple segmentation networks corresponding to different planes in the teacher model and the student model, co-training [7, 26] is introduced so that these networks can be trained simultaneously in our unified framework and benefit from each other. We evaluate our algorithm on our newly collected large dataset and observe a significant improvement of 4.23% compared with the fully supervised method. At last, as DMPCT is a generic and flexible framework, it can be envisioned that better backbone models and fusion strategies can be easily plugged into our framework. Our unified system can be also practically useful for current clinical environments due to the efficiency in leveraging massive unlabeled data to boost segmentation performance.

2. Related Work

Fully-supervised multi-organ segmentation. Early studies of abdominal organ segmentation focused on atlas-based methods [22, 11, 43]. The frameworks are usually problematic because 1) they are not able to capture the large inter-subject variations of abdominal regions and 2) computational time is tightly dependent on the number of atlases. Recently, learning-based approaches with relatively large dataset have been introduced for multi-organ segmentation [17, 34, 8]. Especially, deep Convolutional Neural Networks (CNNs) based methods have achieved a great success

in the medical image segmentation [33, 10, 16, 41, 42, 46, 21] in the last few years. Compared with multi-atlas-based approaches, CNNs based methods are generally more efficient and accurate. CNNs based methods for multi-organ segmentation can be divided into two major categories: 3D CNNs [33, 10, 16] based and 2D CNNs [41, 42, 46, 21] based. 3D CNNs usually adopt the sliding-window strategy to avoid the *out of memory* problem, leading to high time complexity. Compared with 3D CNNs, 2D CNNs based algorithms can be directly end-to-end trained using 2D deep networks, which is less time-consuming.

Semi-supervised learning. The most commonly used techniques for semi-supervised learning include self-training [31, 28], co-training [7], multi-view learning [44] and graph-based methods [6, 40].

In self-training, the classifier is iteratively re-trained using the training set augmented by adding the unlabeled data with their own predictions. The procedure repeated until some convergence criteria are satisfied. In such case, one can imagine that a classification mistake can reinforce itself. Self-training has achieved great performances in many computer vision problems [31, 28] and recently has been applied to deep learning based semi-supervised learning in the biomedical imaging domain [5].

Co-training [7] assumes that (1) features can be split into two independent sets and (2) each sub-feature set is sufficient to train a good classifier. During the learning process, each classifier is retrained with the additional training examples given by the other classifier. Co-training utilizes multiple sets of independent features which describe the same data, and therefore tends to yield more accurate and robust results than self-training [36]. Multi-view learning [44], in general, defines learning paradigms that utilize the agreement among different learners. Co-training is one of the earliest schemes for multi-view learning.

Graph-based semi-supervised methods define a graph where the nodes are labeled and unlabeled examples in the dataset, and edges reflect the similarity of examples. These methods have been widely adopted in non-deep-learning based semi-supervised learning algorithms in the biomedical imaging domain [14, 18, 27].

Different from other methods, our work tactfully embeds the multi-view property of 3D medical data into the co-training framework, which is simple and effective.

3. Deep Multi-Planar Co-Training

We propose Deep Multi-Planar Co-Training (DMPCT), a semi-supervised multi-organ segmentation method which exploits multi-planar information to generate pseudo-labels for unlabeled 3D CT volumes. Assume that we are given a 3D CT volume dataset \mathcal{S} containing K organs. This includes labeled volumes $\mathcal{S}_L = \{(\mathbf{I}_m, \mathbf{Y}_m)\}_{m=1}^L$ and unlabeled

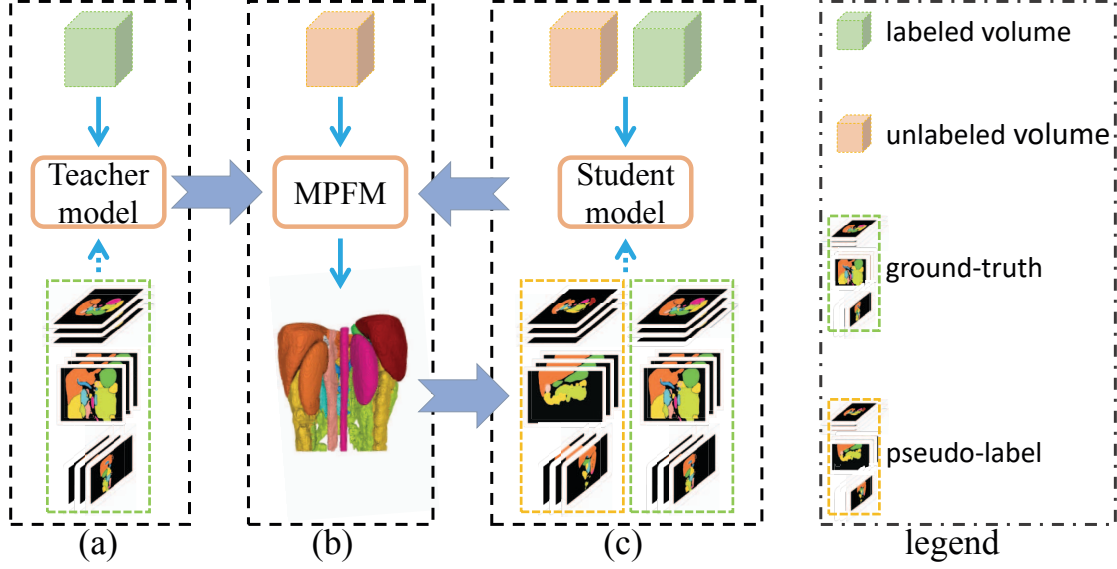


Figure 1. Illustration of the Deep Multi-Planar Co-Training (DMPCT) framework. (a) We first train a teacher model on the labeled dataset. (b) The trained model is used to assign pseudo-labels to the unlabeled data using our multi-planar fusion module as demonstrated in Figure 2. (c) Finally, we train a student model over the union of both the labeled and the unlabeled data. Step (b) and (c) are performed in an iterative manner.

labeled volumes $\mathcal{S}_L = \{\mathbf{I}_m\}_{m=l+1}^M$, where \mathbf{I}_m and \mathbf{Y}_m denote a 3D input volume and its corresponding ground-truth segmentation mask. l and $M-l$ are the numbers of labeled and unlabeled volumes, respectively. Typically $l \ll M$. As shown in Figure 1, DMPCT involves the following steps:

- **Step 1:** train a *teacher model* on the manually labeled data \mathcal{S}_L in the fully supervised setting (see Sec. 3.1).
- **Step 2:** the trained model is then used to assign pseudo-labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ to the unlabeled data \mathcal{S}_U by fusing the estimations from all planes (see Sec. 3.2).
- **Step 3:** train a *student model* on the union of the manually labeled data and automatically labeled data $\mathcal{S}_L \cup \{(\mathbf{I}_m, \hat{\mathbf{Y}}_m)\}_{m=l+1}^M$ (see Sec. 3.3).
- **Step 4:** perform step 2 & 3 in an iterative manner.

3.1. Teacher Model

We train the teacher model on the labeled dataset \mathcal{S}_L . By splitting each volume and its corresponding label mask from the sagittal (S), coronal (C), and axial (A) planes, we can get three sets of 2D slices, *i.e.*, $\mathcal{S}_L^V = \{(\mathbf{I}_n^V, \mathbf{Y}_n^V)\}_{n=1}^{N_V}$, $V \in \{S, C, A\}$, where N_V is the number of 2D slices obtained from plane V . We train a 2D-FCN model (we use [24] as our reference CNN model throughout this paper) to perform segmentation from each plane individually.

Without loss of generality, let $\mathbf{I}^V \in \mathbb{R}^{W \times H}$ and $\mathbf{Y}^V = \{y_i^V\}_{i=1}^{W \times H}$ denote a 2D slice and its corresponding label

mask in \mathcal{S}_L^V , where $y_i^V \in \{0, 1, \dots, K\}$ is the organ label (0 means background) of the i -th pixel in \mathbf{I}^V . Consider a segmentation model $\mathbb{M}^V : \hat{\mathbf{Y}} = \mathbf{f}(\mathbf{I}^V; \theta)$, where θ denotes the model parameters and $\hat{\mathbf{Y}}$ denotes the prediction for \mathbf{I}^V . Our objective function is

$$\mathcal{L}(\mathbf{I}^V, \mathbf{Y}^V; \theta) = -\frac{1}{W \times H} \left[\sum_{i=1}^{W \times H} \sum_{k=0}^K \mathbb{1}(y_i^V = k) \log p_{i,k}^V \right], \quad (1)$$

where $p_{i,k}^V$ denotes the probability of the i -th pixel been classified as label k on 2D slice \mathbf{I}^V and $\mathbb{1}(\cdot)$ is the indicator function. We train the teacher model by optimizing \mathcal{L} w.r.t. θ by stochastic gradient descent.

3.2. Multi-Planar Fusion Module

Given a well-trained teacher model $\{\mathbb{M}^V | V \in \{S, C, A\}\}$, our goal of the multi-planar fusion module is to generate the pseudo-labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ for the unlabeled data \mathcal{S}_U . We first make predictions on the 2D slices from each plane and then reconstruct the 3D volume by stacking all slices back together. Several previous studies [20, 29, 3, 4] suggest that combining predictions from multiple views can often improve the accuracy and the robustness of the final decision since complementary information can be exploited from multiple views simultaneously. Thereby, the fused prediction from multiple planes is superior to any estimation of a single plane. The overall module is shown in Figure 2.

More specifically, majority voting is applied to fuse the

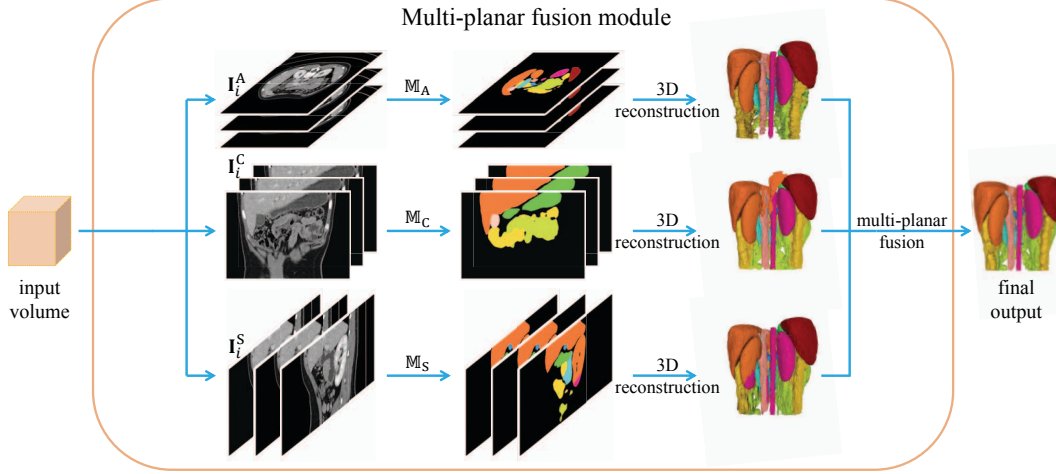


Figure 2. Illustration of the multi-planar fusion module, where the input 3D volume is first parsed into 3 sets of slices along the sagittal, coronal, and axial planes to be evaluated respectively. Then the final 3D estimation is obtained by fusing predictions from each individual plane.




Figure 3. An example of 3D predictions reconstructed from the sagittal, coronal, and axial planes as well as their fusion output. Estimations from single planes are already reasonably well, whereas the single fusion outcome is superior to estimation from any single plane.

hard estimations by seeking an agreement among different planes. If the predictions from all planes do not agree on a voxel, then we select the prediction for that voxel with the maximum confidence. As simple as this strategy might sound, this method has been shown to result in highly robust and efficient outcome in various previous studies [1, 20, 29, 48]. The final decision for the i -th voxel y_i^* of $\hat{\mathbf{Y}}_m$ is:

$$y_i^* = \begin{cases} y_i^V, & \text{if } \exists V, V' \in \{S, C, A\}, V \neq V' \mid y_i^V = y_i^{V'} \\ y_i^{V^*}, & \text{otherwise} \end{cases}, \quad (2)$$

where $V^* = \arg \max_{V \in \{S, C, A\}} \max_j p_{i,j}^V$, $p_{i,j}^S$, $p_{i,j}^C$, and $p_{i,j}^A$ denote the probabilities of the i -th pixel classified as label j from the sagittal, coronal, and axial planes, respectively. y_i^V denotes the hard estimation for the i -th pixel on plane V , i.e., $y_i^V = \arg \max_j p_{i,j}^V$.

As shown in Figure 3, our multi-planar fusion module improves both over- and under-estimation by fusing aspects from different planes and therefore yields a much better

outcome. Note that other rules [2, 41] can also be easily adapted to this module. We do not focus on discussing the influence of the fusion module in this paper, although intuitively better fusion module should lead to higher performance.

3.3. Student Model

After generating the pseudo-labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ for the unlabeled dataset \mathcal{S}_U , the training set can be then enlarged by taking the union of both the labeled and the unlabeled dataset, i.e., $\mathcal{S} = \mathcal{S}_L \cup \{(\mathbf{I}_m, \hat{\mathbf{Y}}_m)\}_{m=l+1}^M$. The student model is trained on this augmented dataset \mathcal{S} the same way we train the teacher model as described in Sec. 3.1. The overall training procedure is summarized in Algorithm 1. In the training stage, we first train a teacher model in a supervised manner and then use it to generate the pseudo-labels for the unlabeled dataset. Then we alternate the training of the student model and the pseudo-label generation procedures in an iterative manner to optimize the student model T times. In the testing stage, we follow the method in Sec.

Algorithm 1 Deep Multi-Planar Co-Training for Multi-Organ Segmentation

Require: A set of labeled data $\mathcal{S}_L = \{(\mathbf{I}_m, \mathbf{Y}_m)\}_{m=1}^l$ and unlabeled volumes $\mathcal{S}_U = \{\mathbf{I}_m\}_{m=l+1}^M$.

Ensure: A trained multi-organ segmentation model $\{\mathbb{M}^S, \mathbb{M}^C, \mathbb{M}^A\}$.

$\mathcal{S} \leftarrow \mathcal{S}_L$

for $t = 1$ to T **do**

 Parse \mathcal{S} into $\mathcal{S}^S, \mathcal{S}^C, \mathcal{S}^A$.

 Train $\mathbb{M}^S, \mathbb{M}^C$, and \mathbb{M}^A on $\mathcal{S}^S, \mathcal{S}^C$, and \mathcal{S}^A respectively.

 Generate pseudo-class labels $\{\hat{\mathbf{Y}}_m\}_{m=l+1}^M$ for the unlabeled dataset \mathcal{S}_U by Eq. 2

 Augment the training set \mathcal{S} by adding the self-labeled examples to \mathcal{S}_L , *i.e.*, $\mathcal{S} = \mathcal{S}_L \cup \{(\mathbf{I}_m, \hat{\mathbf{Y}}_m)\}_{m=l+1}^M$.

end for

 Parse \mathcal{S} into $\mathcal{S}^S, \mathcal{S}^C, \mathcal{S}^A$.

 Train $\mathbb{M}^S, \mathbb{M}^C$, and \mathbb{M}^A on $\mathcal{S}^S, \mathcal{S}^C$, and \mathcal{S}^A respectively.

3.2 to generate the final estimation using the T -th student model.

4. Experiments

4.1. Dataset and Evaluation

Our fully-labeled dataset includes 210 contrast-enhanced abdominal clinical CT images in the portal venous phase, in which we randomly choose 50/30/80 patients for training, validation, and testing, unless otherwise specified. A total of 16 structures (Aorta, Adrenal gland, Celiac AA, Colon, Duodenum, Gallbladder, Interior Vena Cava (IVC), Kidney (left, right), Liver, Pancreas, Superior Mesenteric Artery (SMA), Small bowel, Spleen, Stomach, Veins) for each case were segmented by four experienced radiologists, and confirmed by an independent senior expert. Our unlabeled dataset consists of 100 unlabeled cases acquired from a local hospital. To the best of our knowledge, this is the largest abdominal CT dataset with the most number of organs segmented. Each CT volume consists of $319 \sim 1051$ slices of 512×512 pixels, and have voxel spatial resolution of $([0.523 \sim 0.977] \times [0.523 \sim 0.977] \times 0.5)\text{mm}^3$. The metric we use is the Dice-Sørensen Coefficient (DSC), which measures the similarity between the prediction voxel set \mathcal{Z} and the ground-truth set \mathcal{Y} , with the mathematical form of $\text{DSC}(\mathcal{Z}, \mathcal{Y}) = \frac{2 \times |\mathcal{Z} \cap \mathcal{Y}|}{|\mathcal{Z}| + |\mathcal{Y}|}$. For each organ, we report an average DSC together with the standard deviation over all the testing cases.

4.2. Implementation Details

We set the learning rate to be 10^{-9} . The teacher model and the student model are trained for 80,000 and 160,000 iterations respectively. The validation set is used for tuning

the hyper-parameters. Similar to [19], we use three windows of $[-125, 275]$, $[-160, 240]$, and $[-1000, 1000]$ Housefield Units as the three input channels respectively. The intensities of each slice are rescaled to $[0.0, 1.0]$. Similar to [48, 45, 41], we initialize the network parameters θ by using the FCN-8s model [24] pre-trained on the PascalVOC image segmentation dataset. The iteration number T in Algorithm 1 is set to 2, *i.e.*, $T = 2$, as the performance of the validation set gets saturated.

4.3. Comparison with the Baseline

We show that our proposed DMPCT works better than other methods: 1) fully supervised learning method [24] (denoted as FCN), and 2) single planar based semi-supervised learning approach [5] (denoted as SPSL). Both 1) and 2) are applied on each individual plane separately, and then the final result is obtained via multi-planar fusion (see Sec 3.2). As shown in Table 1, with 50 labeled data, by varying the number of unlabeled data from 0 to 100, the average DSC of DMPCT increases from 73.71% to 77.94% and the standard deviation decreases from 9.97% to 8.51%. Compared with SPSL, our proposed DMPCT can boost the performance in both settings (*i.e.*, 50 labeled data + 50 unlabeled data and 50 labeled data + 100 unlabeled data). Besides, the p -values for testing significant difference between our DMPCT (50 labeled data + 100 unlabeled data) and FCN (50 labeled data + 0 unlabeled data) for 16 organs are shown in the last column of Table 1, which suggests significant statistical improvements among almost all organs. Figure 4 shows comparison results of our DMPCT and the fully supervised method by box plots.

It is noteworthy that greater improvements are observed especially for those difficult organs, *i.e.*, organs either small in sizes or with complex geometric characteristics. Table 1 indicates that our DMPCT approach boosts the segmentation performance of these small hard organs by 5.54% (Pancreas), 8.72% (Colon), 9.89% (Duodenum), 8.12% (Small bowels) and 5.48% (Veins), 5.76% (IVC). This promising result indicates that our method distills a reasonable amount of knowledge from the unlabeled data. An example is shown in Figure 5. In this particular case, the DSCs for Celiac AA, Colon, Duodenum, IVC, Pancreas and Veins are boosted from 60.13%, 46.79%, 71.08%, 69.23%, 63.48% to 79.45%, 83.81%, 77.59%, 74.75%, 75.31% respectively.

4.4. Discussion

4.4.1 Amount of labeled data

For ablation analysis, we enlarge the labeled training set to 100 cases and keep the rest of the settings the same. As shown in Figure 6, with more labeled data, the semi-supervised methods (DMPCT, SPSL) still obtain better performance than the supervised method (FCN), while the performance gain becomes less prominent. This is probably

Table 1. The comparison of segmentation accuracy (DSC, %) by using 50 labeled data and varying the number of unlabeled data (e.g., 50-0 indicates 50 labeled data and 0 unlabeled data). We report the mean and standard deviation over 80 cases. The p -values for testing significant difference between DMPCT (50-100) and FCN (50-0) are shown. Significant statistical improvement is shown in *italic* with $p < 0.05$. See Section 4.3 for definitions of FCN, SPSL, and DMPCT (Ours).

Organ Type	FCN	SPSL		DMPCT (Ours)		p -value
	50 - 0	50 - 50	50 - 100	50 - 50	50 - 100	
Aorta	89.14 \pm 7.95	91.10 \pm 5.52	90.76 \pm 5.90	91.43 \pm 4.88	91.54 \pm 4.65	<i>3.32 $\times 10^{-5}$</i>
Adrenal gland	26.45 \pm 12.1	29.92 \pm 14.7	26.93 \pm 15.6	30.58 \pm 12.7	35.48 \pm 11.8	<i>1.98 $\times 10^{-15}$</i>
Celiac AA	35.01 \pm 19.7	37.27 \pm 19.0	39.78 \pm 18.4	36.25 \pm 20.5	40.50 \pm 18.9	<i>1.00 $\times 10^{-5}$</i>
Colon	71.81 \pm 14.9	78.28 \pm 13.0	79.58 \pm 12.9	79.61 \pm 12.3	80.53 \pm 11.6	<i>7.69 $\times 10^{-12}$</i>
Duodenum	54.89 \pm 15.5	57.77 \pm 17.3	62.22 \pm 14.8	66.95 \pm 12.6	64.78 \pm 13.8	<i>1.95 $\times 10^{-19}$</i>
Gallbladder	86.53 \pm 6.21	87.87 \pm 5.45	88.02 \pm 5.83	88.45 \pm 5.07	87.77 \pm 6.29	0.002
IVC	77.67 \pm 9.49	81.28 \pm 8.87	82.63 \pm 7.31	83.49 \pm 6.94	83.43 \pm 7.02	<i>9.30 $\times 10^{-14}$</i>
Kidney (L)	95.12 \pm 5.01	95.59 \pm 4.97	95.88 \pm 3.68	95.82 \pm 3.60	96.09 \pm 3.42	<i>3.69 $\times 10^{-6}$</i>
Kidney (R)	95.69 \pm 2.36	95.77 \pm 4.93	96.14 \pm 2.94	96.17 \pm 2.75	96.26 \pm 2.29	<i>1.74 $\times 10^{-7}$</i>
Liver	95.45 \pm 2.41	96.06 \pm 0.99	96.07 \pm 1.03	96.11 \pm 0.97	96.15 \pm 0.92	0.005
Pancreas	76.49 \pm 11.6	80.12 \pm 7.52	80.93 \pm 6.84	81.46 \pm 6.32	82.03 \pm 6.16	<i>2.97 $\times 10^{-8}$</i>
SMA	52.26 \pm 17.1	51.81 \pm 18.2	51.94 \pm 17.1	49.40 \pm 19.2	52.70 \pm 17.7	0.667
Small bowel	71.13 \pm 13.1	78.93 \pm 12.6	79.97 \pm 12.8	79.49 \pm 12.1	79.25 \pm 12.6	<i>2.53 $\times 10^{-22}$</i>
Spleen	94.81 \pm 2.64	95.46 \pm 2.09	95.58 \pm 1.90	95.73 \pm 2.03	95.98 \pm 1.59	<i>1.83 $\times 10^{-10}$</i>
Stomach	91.38 \pm 3.94	92.62 \pm 3.71	92.92 \pm 3.65	93.33 \pm 3.47	93.42 \pm 3.21	<i>3.30 $\times 10^{-23}$</i>
Veins	64.75 \pm 15.4	70.43 \pm 14.3	69.66 \pm 14.6	69.82 \pm 14.5	70.23 \pm 14.4	<i>4.16 $\times 10^{-15}$</i>
Mean	73.71 \pm 9.97	76.32 \pm 9.58	76.87 \pm 9.08	77.20 \pm 8.75	77.94 \pm 8.51	<i>4.74 $\times 10^{-90}$</i>




Figure 4. Performance comparison (DSC, %) in box plots of 16 organs by using 50 labeled data and varying the number of unlabeled data (e.g., 50-0 indicates 50 labeled data and 0 unlabeled data). See Section 4.3 for definitions of FCN and DMPCT (Ours).

because the network is already trained well when large training set is available. We believe that if much more unlabeled data can be provided the performance should go up considerably. In addition, we find that DMPCT outperforms SPSL in every setting, which further demonstrates the usefulness of multi-planar fusion in our co-training framework.

4.4.2 Comparison with 3D network-based self-training

Various previous studies [25, 41] demonstrate that 2D multi-planar fusion outperforms directly 3D learning in the fully supervised setting. 3D CNNs come with an increased number of parameters, significant memory and computational requirements. Due to GPU memory restrictions,

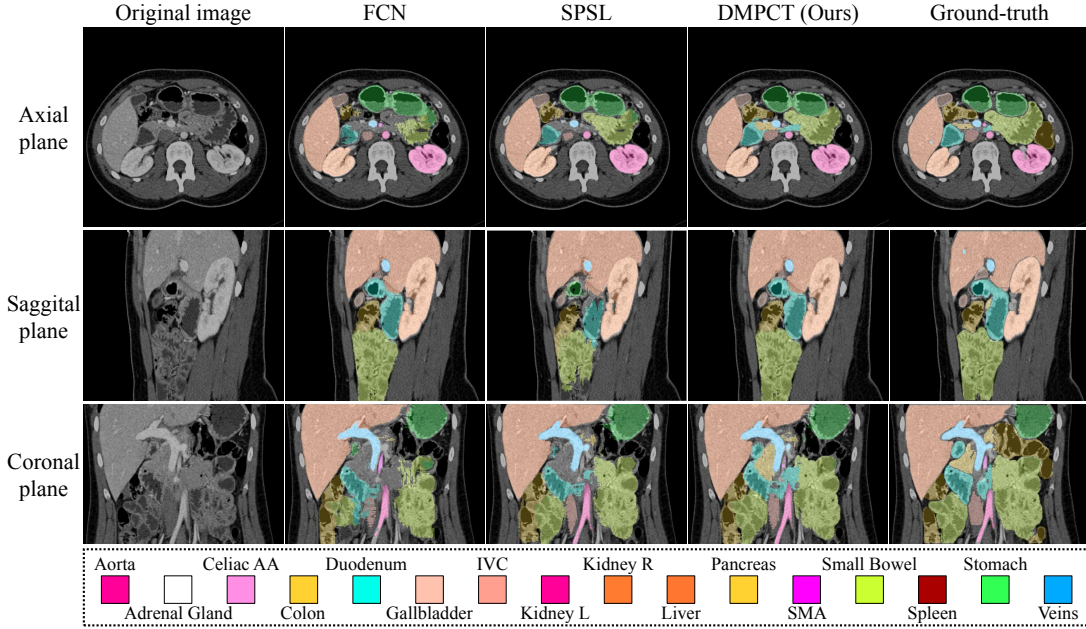


Figure 5. Comparisons among FCN, SPSL, and DMPCT (Ours) viewed from multiple planes. 50 labeled cases are used for all methods. 100 unlabeled cases are used for the SPSL and DMPCT. For this particular case, FCN obtains an average DSC of 72.75%, SPSL gets 78.87%, and DMPCT (Ours) gets 80.75%. See Section 4.3 for definitions of FCN, SPSL, and DMPCT (Ours). Best viewed in color.




Figure 6. Ablation study on numbers of labeled data and unlabeled data. Mean DSC of all testing cases under all settings (*e.g.*, 50-0 indicates 50 labeled data and 0 unlabeled data). See Section 4.3 for definitions of FCN, SPSL, and DMPCT (Ours).

these 3D CNN approaches which adopt the sliding-window strategy do not act on the entire 3D CT volume, but instead on local 3D patches [12, 15, 17]. This results in the lack of holistic information and low efficiency. In order to prove that DMPCT outperforms direct 3D learning in the semi-supervised setting, we also implement a patch-based 3D UNet [12]. 3D UNet gets 69.66% in terms of mean DSC using 50 labeled data. When adding 100 unlabeled data the

performance even drops to 65.21%. This clearly shows that in 3D learning the teacher model is not trained well, thus the errors of the pseudo-labels are reinforced during student model training.

Table 2. Cross-dataset generalization results.

Organ	Spleen	Kidney (R)	Kidney (L)	Gall Bladder	Liver
FCN	71.85 \pm 26.13	54.44 \pm 20.04	54.98 \pm 26.63	48.13 \pm 26.07	85.46 \pm 16.81
DMPCT (Ours)	83.68 \pm 16.53	71.36 \pm 20.85	69.95 \pm 20.50	60.05 \pm 26.91	92.11 \pm 6.46
Organ	Stomach	Aorta	IVC	Veins	Pancreas
FCN	38.89 \pm 23.86	70.43 \pm 19.70	53.67 \pm 18.40	35.54 \pm 18.94	39.40 \pm 25.34
DMPCT (Ours)	54.78 \pm 26.57	76.05 \pm 15.99	68.18 \pm 14.58	37.52 \pm 15.86	60.05 \pm 16.61

4.4.3 Comparison with traditional co-training

In order to show that our DMPCT outperforms traditional co-training algorithm [7], we also select only the most confident samples during each iteration. Here the confidence score is measured by the entropy of probability distribution for each voxel in one slice. Under the setting of 50 labeled cases and 50 unlabeled cases, we select top 5000 samples with the highest confidence in each iteration. The whole training process takes about 6-7 iterations for each plane. The complete training requires more than 50 hours. Compared with our approach, this method requires much more time to converge. It obtains a mean DSC of 76.52%, slightly better than SPSL but worse than our DMPCT, which shows that selecting the most confident samples during training may not be a wise choice for deep network based semi-supervised learning due to its low efficiency.

4.4.4 Cross dataset generalization

We apply our trained DMPCT model (50 labeled data + 100 unlabeled data) and baseline FCN model (50 labeled data + 0 unlabeled data) on a public available abdominal CT datasets¹ with 13 anatomical structures labeled *without any further re-training* on new data cases. 10 out of the 13 structures are evaluated which are also manually annotated in our own dataset and we find that our proposed method improves the overall mean DSC and also reduces the standard deviation significantly, as shown in Table 2. The overall mean DSC as well as the standard deviation for the 10 organs is improved from 59.23 \pm 22.20% to 67.38 \pm 19.64%. We also directly test our models on the NIH pancreas segmentation dataset of 82 cases² and observe that our DMPCT model achieves an average DSC of 66.16%, outperforming the fully supervised method, with an average DSC of 58.73%, by more than 7%. This may demonstrate that our approach, which leverages more unlabeled data from multiple planes, turns out to be much more generalizable than the baseline model.

¹30 training data sets at <https://www.synapse.org/#!/Synapse:syn3193805/wiki/217789>

²<https://wiki.cancerimagingarchive.net/display/Public/Pancreas-CT>

4.4.5 Computation time

In our experiments, the teacher model training process takes about 4.94 hours on an NVIDIA TITAN Xp GPU card for 80,000 iterations over all the training cases. The average computation time for generating pseudo-label as well as testing per volume depends on the volume of the target structure, and the average computation time for 16 organs is approximately 4.5 minutes, which is comparable to other recent methods [48, 32] even for single structure inference. The student model training process takes about 9.88 hours for 160,000 iterations.

5. Conclusion

In this paper, we designed a systematic framework DMPCT for multi-organ segmentation in abdominal CT scans, which is motivated by the traditional co-training strategy to incorporate multi-planar information for the unlabeled data during training. The pseudo-labels are iteratively updated by inferencing comprehensively on multiple configurations of unlabeled data with a multi-planar fusion module. We evaluate our approach on our own large newly collected high-quality dataset. The results show that 1) our method outperforms the fully supervised learning approach by a large margin; 2) it outperforms the single planar method, which further demonstrates the benefit of multi-planar fusion; 3) it can learn better if more unlabeled data provided especially when the scale of labeled data is small. Our framework can be practical in assisting radiologists for clinical applications since the annotation of multiple organs in 3D volumes requires massive labor from radiologists. Our framework is not specific to a certain structure, but shows robust results in multiple complex anatomical structures within efficient computational time. We believe that our algorithm may achieve even higher accuracy if a more powerful backbone network or an advanced fusion algorithm is employed, which we leave as the future work.

Acknowledgement. This work was supported by the Lustgarten Foundation for Pancreatic Cancer Research and also supported by NSFC No. 61672336. We thank Prof. Seyoun Park, Dr. Lingxi Xie, Cihang Xie, Zhishuai Zhang, Fengze Liu, Zhuotun Zhu and Yingda Xia for instructive discussions.

References

- [1] P. Aljabar, R. A. Heckemann, A. Hammers, J. V. Hajnal, and D. Rueckert. Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. *NeuroImage*, 46(3):726–738, 2009. 4
- [2] A. J. Asman and B. A. Landman. Non-local statistical label fusion for multi-atlas segmentation. *MIA*, 17(2):194–208, 2013. 4
- [3] S. Bai, X. Bai, Z. Zhou, Z. Zhang, and L. J. Latecki. Gift: A real-time and scalable 3d shape search engine. In *CVPR*, pages 5023–5032, 2016. 3
- [4] S. Bai, X. Bai, Z. Zhou, Z. Zhang, Q. Tian, and L. J. Latecki. Gift: Towards scalable 3d shape retrieval. *TMM*, 19(6):1257–1271, 2017. 3
- [5] W. Bai, O. Oktay, M. Sinclair, H. Suzuki, M. Rajchl, G. Tarroni, B. Glocker, A. King, P. M. Matthews, and D. Rueckert. Semi-supervised learning for network-based cardiac mr image segmentation. In *MICCAI*, pages 253–260, 2017. 2, 5
- [6] A. Blum and S. Chawla. Learning from labeled and unlabeled data using graph mincuts. 2001. 2
- [7] A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, pages 92–100, 1998. 2, 8
- [8] T. Brosch and A. Saalbach. Foveal fully convolutional nets for multi-organ segmentation. In *Medical Imaging 2018: Image Processing*, volume 10574, page 105740U, 2018. 2
- [9] H. Chen, X. Qi, J.-Z. Cheng, P.-A. Heng, et al. Deep contextual networks for neuronal structure segmentation. In *AAAI*, pages 1167–1173, 2016. 1
- [10] S. Chen, H. Roth, S. Dorn, M. May, A. Cavallaro, M. M. Lell, M. Kachelrieß, H. Oda, K. Mori, and A. Maier. Towards automatic abdominal multi-organ segmentation in dual energy ct using cascaded 3d fully convolutional network. *arXiv preprint arXiv:1710.05379*, 2017. 2
- [11] C. Chu, M. Oda, T. Kitasaka, K. Misawa, M. Fujiwara, Y. Hayashi, Y. Nimura, D. Rueckert, and K. Mori. Multi-organ Segmentation based on Spatially-Divided Probabilistic Atlas from 3D Abdominal CT Images. In *MICCAI*, pages 165–172, 2013. 2
- [12] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432, 2016. 7
- [13] D. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. In *NIPS*, pages 2843–2851, 2012. 1
- [14] A. Ciurte, X. Bresson, O. Cuisenaire, N. Houhou, S. Nedevschi, J.-P. Thiran, and M. B. Cuadra. Semi-supervised segmentation of ultrasound images based on patch representation and continuous min cut. *PloS one*, 9(7):e100972, 2014. 1, 2
- [15] Q. Dou, H. Chen, Y. Jin, L. Yu, J. Qin, and P.-A. Heng. 3d deeply supervised network for automatic liver segmentation from ct volumes. In *MICCAI*, pages 149–157, 2016. 7
- [16] 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. 2
- [17] E. Gibson, F. Giganti, Y. Hu, E. Bonmati, S. Bandula, K. Gurusamy, B. R. Davidson, S. P. Pereira, M. J. Clarkson, and D. C. Barratt. Towards image-guided pancreas and biliary endoscopy: automatic multi-organ segmentation on abdominal ct with dense dilated networks. In *MICCAI*, pages 728–736, 2017. 2, 7
- [18] L. Gu, Y. Zheng, R. Bise, I. Sato, N. Imanishi, and S. Aiso. Semi-supervised learning for biomedical image segmentation via forest oriented super pixels (voxels). In *MICCAI*, pages 702–710, 2017. 1, 2
- [19] A. P. Harrison, Z. Xu, K. George, L. Lu, R. M. Summers, and D. J. Mollura. Progressive and multi-path holistically nested neural networks for pathological lung segmentation from ct images. In *MICCAI*, pages 621–629, 2017. 5
- [20] R. A. Heckemann, J. V. Hajnal, P. Aljabar, D. Rueckert, and A. Hammers. Automatic anatomical brain mri segmentation combining label propagation and decision fusion. *NeuroImage*, 33(1):115–126, 2006. 3, 4
- [21] P. Hu, F. Wu, J. Peng, Y. Bao, F. Chen, and D. Kong. Automatic abdominal multi-organ segmentation using deep convolutional neural network and time-implicit level sets. *International journal of computer assisted radiology and surgery*, 12(3):399–411, 2017. 2
- [22] J. E. Iglesias and M. R. Sabuncu. Multi-atlas segmentation of biomedical images: A survey. *MIA*, 24(1):205–219, 2015. 2
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. 1
- [24] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 1, 3, 5
- [25] A. Prason, K. Petersen, C. Igel, F. Lauze, E. Dam, and M. Nielsen. Deep Feature Learning for Knee Cartilage Segmentation Using a Triplanar Convolutional Neural Network. *MICCAI*, 2013. 6
- [26] S. Qiao, W. Shen, Z. Zhang, B. Wang, and A. Yuille. Deep co-training for semi-supervised image recognition. In *ECCV*, 2018. 2
- [27] C. Qin, R. G. Moreno, C. Bowles, C. Ledig, P. Scheltens, F. Barkhof, H. Rhodius-Meester, B. M. Tijms, A. W. Lemstra, W. M. van der Flier, B. Glocker, and D. Rueckert. A semi-supervised large margin algorithm for white matter hyperintensity segmentation. In *MICCAI*, pages 104–112, 2016. 1, 2
- [28] I. Radosavovic, P. Dollr, R. Girshick, G. Gkioxari, and K. He. Data distillation: Towards omni-supervised learning. In *CVPR*, 2018. 2
- [29] T. Rohlfing, R. Brandt, R. Menzel, and C. R. Maurer Jr. Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. *NeuroImage*, 21(4):1428–1442, 2004. 3, 4
- [30] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 1

- [31] C. Rosenberg, M. Hebert, and H. Schneiderman. Semi-supervised self-training of object detection models. In *WACV/MOTION*, pages 29–36, 2005. 2
- [32] H. R. Roth, L. Lu, N. Lay, A. P. Harrison, A. Farag, A. Sohn, and R. M. Summers. Spatial aggregation of holistically-nested convolutional neural networks for automated pancreas localization and segmentation. *MIA*, 45:94 – 107, 2018. 1, 8
- [33] H. R. Roth, H. Oda, Y. Hayashi, M. Oda, N. Shimizu, M. Fujiwara, K. Misawa, and K. Mori. Hierarchical 3d fully convolutional networks for multi-organ segmentation. *arXiv preprint arXiv:1704.06382*, 2017. 1, 2
- [34] H. R. Roth, C. Shen, H. Oda, T. Sugino, M. Oda, Y. Hayashi, K. Misawa, and K. Mori. A multi-scale pyramid of 3d fully convolutional networks for abdominal multi-organ segmentation. *arXiv preprint arXiv:1806.02237*, 2018. 2
- [35] W. Shen, B. Wang, Y. Jiang, Y. Wang, and A. Yuille. Multi-stage multi-recursive-input fully convolutional networks for neuronal boundary detection. In *ICCV*, pages 2391–2400, 2017. 1
- [36] R. T. Sousa and J. Gama. Comparison between co-training and self-training for single-target regression in data streams using amrules. 2017. 2
- [37] P. Tang, X. Wang, S. Bai, W. Shen, X. Bai, W. Liu, and A. L. Yuille. Pcl: Proposal cluster learning for weakly supervised object detection. *TPAMI*, 2018. 1
- [38] P. Tang, X. Wang, X. Bai, and W. Liu. Multiple instance detection network with online instance classifier refinement. In *CVPR*, pages 2843–2851, 2017. 1
- [39] P. Tang, X. Wang, A. Wang, Y. Yan, W. Liu, J. Huang, and A. Yuille. Weakly supervised region proposal network and object detection. In *ECCV*, pages 370–386, 2018. 1
- [40] J. Wang, T. Jebara, and S.-F. Chang. Semi-supervised learning using greedy max-cut. *JMLR*, 14(Mar):771–800, 2013. 2
- [41] Y. Wang, Y. Zhou, W. Shen, S. Park, E. K. Fishman, and A. L. Yuille. Abdominal multi-organ segmentation with organ-attention networks and statistical fusion. *arXiv preprint arXiv:1804.08414*, 2018. 2, 4, 5, 6
- [42] Y. Wang, Y. Zhou, P. Tang, W. Shen, E. K. Fishman, and A. L. Yuille. Training multi-organ segmentation networks with sample selection by relaxed upper confident bound. In *MICCAI*, 2018. 1, 2
- [43] R. Wolz, C. Chu, K. Misawa, M. Fujiwara, K. Mori, and D. Rueckert. Automated abdominal multi-organ segmentation with subject-specific atlas generation. *TMI*, 32(9):1723–1730, 2013. 2
- [44] C. Xu, D. Tao, and C. Xu. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*, 2013. 2
- [45] 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 *CVPR*, 2018. 5
- [46] X. Zhou, T. Ito, R. Takayama, S. Wang, T. Hara, and H. Fujita. Three-dimensional ct image segmentation by combining 2d fully convolutional network with 3d majority voting. In *Deep Learning and Data Labeling for Medical Applications*, pages 111–120. 2016. 2
- [47] 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
- [48] 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*, pages 693–701, 2017. 1, 4, 5, 8