

---

# Uncertainty-aware multi-view co-training for semi-supervised medical image segmentation and domain adaptation<sup>1</sup>

Yingda Xia<sup>a</sup>, Dong Yang<sup>b</sup>, Zhiding Yu<sup>b</sup>, Fengze Liu<sup>a</sup>, Jinzheng Cai<sup>c</sup>, Lequan Yu<sup>d</sup>, Zhuotun Zhu<sup>a</sup>, Daguang Xu<sup>b</sup>, Alan Yuille<sup>a</sup>, Holger Roth<sup>b,2</sup>,

<sup>a</sup> Johns Hopkins University

<sup>b</sup> Nvidia

<sup>c</sup> University of Florida

<sup>d</sup> The Chinese University of Hong Kong

## Abstract

Although having achieved great success in medical image segmentation, deep learning-based approaches usually require large amounts of well-annotated data, which can be extremely expensive in the field of medical image analysis. Unlabeled data, on the other hand, is much easier to acquire. Semi-supervised learning and unsupervised domain adaptation both take the advantage of unlabeled data, and they are closely related to each other. In this paper, we propose **uncertainty-aware multi-view co-training** (UMCT), a unified framework that addresses these two tasks for volumetric medical image segmentation. Our framework is capable of efficiently utilizing unlabeled data for better performance. We firstly rotate and permute the 3D volumes into multiple views and train a 3D deep network on each view. We then apply co-training by enforcing multi-view consistency on unlabeled data, where an uncertainty estimation of each view is utilized to achieve accurate labeling. Experiments on the NIH pancreas segmentation dataset and a multi-organ segmentation dataset show state-of-the-art performance of the proposed framework on semi-supervised medical image segmentation. Under unsupervised domain adaptation settings, we validate the effectiveness of this work by adapting our multi-organ segmentation model to two pathological organs from the Medical Segmentation Decathlon Datasets. Additionally, we show that our UMCT-DA model can even effectively handle the challenging situation where labeled source data is inaccessible, demonstrating strong potentials for real-world applications.

## 1 Introduction

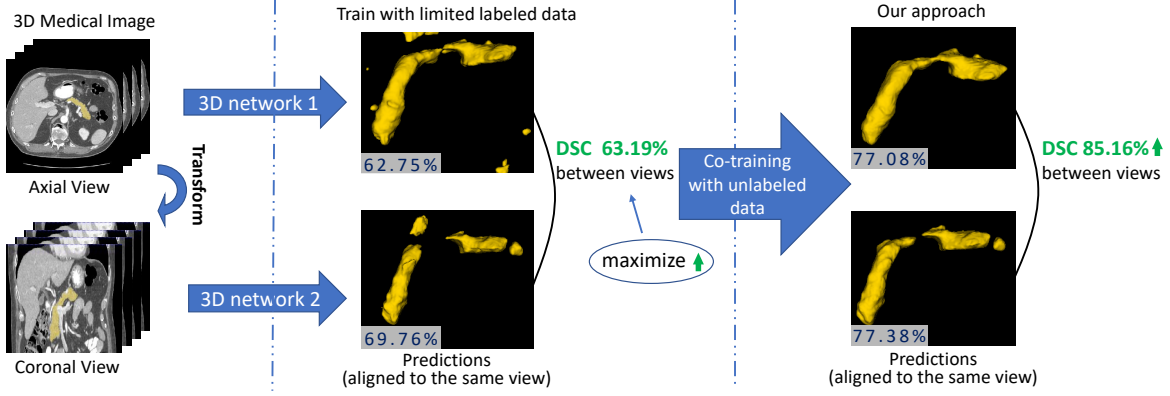
Deep learning has achieved great successes in various computer vision tasks, such as 2D image recognition [22, 26, 31, 56, 59] and semantic segmentation [7, 8, 39, 67]. However, deep networks usually rely on large-scale labeled datasets for training. When it comes to medical volumetric data, human labeling can be extremely costly and often requires expert domain knowledge. Medical image segmentation (i.e. the labeling tissues and organs in CTs and MRIs) plays a critical role in biomedical image analysis and surgical planning. Deep learning-based approaches have been widely adopted for this task and have led to state-of-the-art performance [40, 48, 64, 65]. However, acquiring well-annotated segmentation labels in medical images requires both high-level expertise of radiologists and careful manual labeling of object masks or surface boundaries.

In this paper, we aim to design an approach that can utilize large-scale unlabeled data to improve volumetric medical image segmentation, and is applicable to the scenarios of both semi-supervised learning (SSL) and unsupervised domain adaptation (UDA). SSL and UDA share a common setting by assuming the availability of a labeled training set (denoted as  $\mathcal{S}$ ), as well as an unlabeled one (denoted as  $\mathcal{U}$ ). The

---

<sup>1</sup>Preprint, to appear in Medical Image Analysis (<https://doi.org/10.1016/j.media.2020.101766>)

<sup>2</sup>Corresponding author, email: hroth@nvidia.com

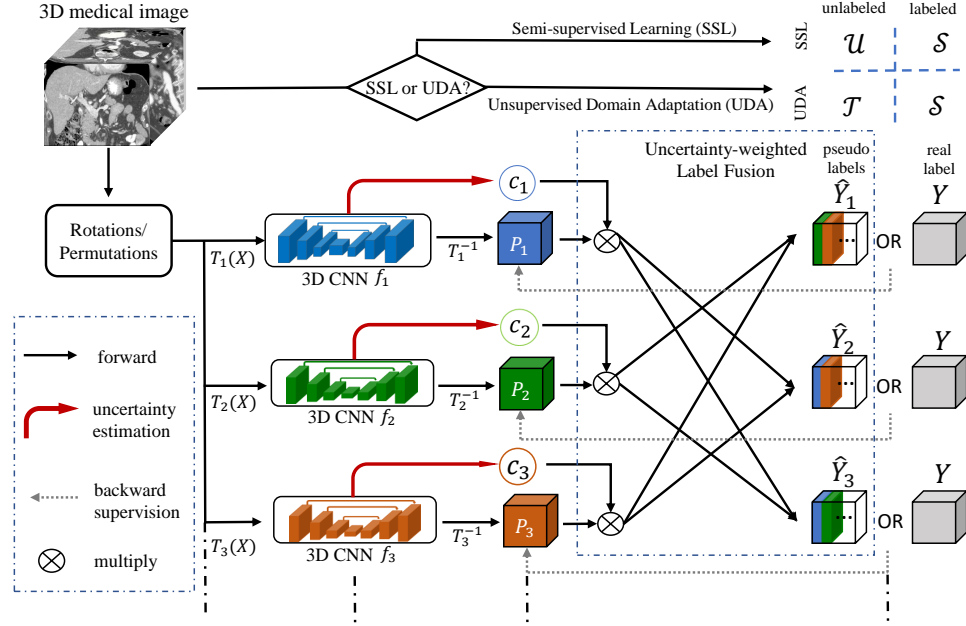


**Figure 1.** An example of our approach for pancreas segmentation (best viewed in color). With limited training data, two 3D networks which are trained on axial and coronal view, respectively, both perform poorly as measured by DSC scores (in dark blue) with ground truth annotations. We observe that the DSC between the two views (in green) is also low, indicating large view differences. With our co-training approach, we minimize the difference between the two predictions on unlabeled data, resulting in significant improvement on each view.

difference between the two tasks is that for  $\mathcal{S}$  and  $\mathcal{U}$  we assume the same distribution in SSL while a larger domain shift is assumed in the UDA setting. Despite such differences, approaches in these two tasks are often closely related. SSL approaches such as self-training [2, 37, 49], co-training [5, 70] and GAN based methods [11, 32] have been widely applied to UDA [9, 14, 24, 55, 61, 72, 73], and vice versa.

Inspired by the success of co-training [5] and its application to single 2D images [46], we further extend this idea to 3D volumetric data. Typical co-training requires at least two views (i.e. sources) of the data, of which either should be sufficient to train a classifier on. Co-training minimizes the disagreement by assigning pseudo labels between each view on unlabeled data. [5] further proved that co-training has PAC-like guarantees on semi-supervised learning with an additional assumption that the two views are conditionally independent given the category. Since most computer vision tasks have only one source of data, encouraging view differences is a crucial factor for successful co-training. For example, *deep co-training* [46] trains multiple deep networks to act as different views by utilizing adversarial examples [21] to address this issue. Another aspect of co-training to emphasize is view confidence estimation. In multi-view settings, with growing differences between each view, the quality of each prediction becomes less and less guaranteed and might result in bad pseudo labels that can be harmful if used in the training process. Co-training could benefit from trusting reliable predictions and degrading the unreliable ones. However, distinguishing reliable and unreliable predictions is challenging for unlabeled data due to lack of ground-truth.

To address the above two important issues, we propose an *uncertainty-aware multi-view co-training* (UMCT) framework, shown in Fig. 2. We introduce view differences by exploring multiple viewpoints of 3D data through spatial transformations, such as rotation and permutation. The permutation here is defined as the rearrangements of the coordinate system, such as transpose and flip, and “view” is defined as the transformed input data after permutation. Hence, our multi-view approach naturally applies to analyzing 3D data and can be integrated with the proposed co-training framework. Fig. 1 gives an example of the intuition of our approach in two-view scenario. On unlabeled data, we propose to maximize the similarity of the predictions between the two views, resulting in improved segmentation performance on each view. Another key component is the view confidence estimation. We propose to estimate the uncertainty of predictions in each view with Bayesian deep networks by adding dropout in the architectures [16]. A confidence score is computed based on epistemic uncertainty [30], which can act as a weight for each prediction. After propagation through this *uncertainty-weighted label fusion module* (ULF), a set of more accurate pseudo labels can be obtained for each view, which is used as supervision



**Figure 2.** Overall framework of **uncertainty-aware multi-view co-training (UMCT)**, best viewed in color. UMCT can be applied to either the semi-supervised learning (SSL) task or the unsupervised domain adaptation (UDA) task, both of which include an unlabeled and a labeled subset of data. The overall pipeline is described as follows. The  $n$  multi-view inputs of  $\mathbf{X}$  are first generated through different transforms  $\mathbf{T}$ , like rotations and permutations, before being fed into  $n$  deep networks with asymmetrical 3D kernels. A confidence score  $c$  is computed for each view by uncertainty estimation and acts as the weights to compute the pseudo labels  $\hat{Y}$  of other views (Eq. 6) after inverse transform  $\mathbf{T}^{-1}$  of the predictions. The pseudo labels  $\hat{Y}$  for unlabeled data and ground truth  $Y$  for labeled data are used as supervisions during training.

signal for unlabeled data.

UMCT was previously published as a conference paper [62], in which we verified its effectiveness under standard semi-supervised settings on individual organs. In this paper, we extensively validate our approach on more challenging tasks, e.g. multi-organ segmentation. Moreover, we apply our approach to the task of unsupervised domain adaptation, with a labeled source domain and unlabeled target domain. In medical image analysis, this is considered as an important task since we should prefer a model or an approach that has the capability to generalize across datasets from different data sources (e.g. types of machines, acquisition protocols, and characteristics of patients). In addition to the original experiments on NIH pancreas dataset [50], we validate our approach on a multi-organ dataset used in [19] with 8 labeled abdominal organs under semi-supervised settings. We then utilize our co-training approach to adapt the multi-organ model to the Medical Image Decathlon (MSD) [57] pathological liver and pancreas datasets. We even push our approach one step further by assuming that we only have the source model in the absence of source data. With very simple modifications, our final model UMCT-DA illustrates strong potential on this challenging scenario.

## 2 Related Work

**Semi-supervised learning** aims at learning models with limited labeled data and a large proportion of unlabeled data [3, 5, 70, 71]. Emerging semi-supervised approaches have been successfully applied to image recognition using deep neural networks [1, 6, 33, 41, 47, 53, 60]. These algorithms mostly rely on additional regularization terms to train the networks to be resistant to some specific noise. A recent

---

approach [46] extended the co-training strategy to 2D deep networks and multiple views, using adversarial examples to encourage view differences to boost performance.

**Semi-supervised medical image analysis.** [10] mentioned that current semi-supervised medical analysis methods fall into 3 types - self-training (teacher-student models), co-training (with hand-crafted features) and graph-based approaches (mostly applications of graph-cut based optimization). [2] introduced a deep network based self-training framework with conditional random field (CRF) based iterative refinements for medical image segmentation. [68] trained three 2D networks from three planar slices of the 3D data and fused them in each self-training iteration to get a stronger student model. [36] extended the self-ensemble approach  $\pi$  model [33] with 90-degree rotations making the network rotation-invariant. Generative adversarial network (GAN) based approaches are also popular recently for medical imaging [12, 28, 43]. Moreover, mixed supervisions [42, 54] combining dense label masks and weak labels like bounding boxes, slice- or image-level labels, etc., is another field of study to alleviate labeling efforts and is related to semi-supervised medical image analysis.

**Uncertainty Estimation.** Traditional approaches include particle filtering and CRFs [4, 23]. For deep learning, uncertainty is more often measured with Bayesian deep networks [15, 16, 30]. In our work, we emphasize the importance of uncertainty estimation in semi-supervised learning, since most of the training data here is not annotated. We propose to estimate the confidence of each view in our co-training framework via Bayesian uncertainty estimation.

**2D/3D hybrid networks.** 2D networks and 3D networks both have advantages and limitations. The former benefits from 2D pre-trained weights and well-studied architectures on natural images, while the latter better explores 3D information utilizing 3D convolution kernels. [35, 63] either uses 2D probability maps or 2D feature maps for building 3D models. [38] proposed a 3D architecture which can be initialized by 2D pre-trained models. Moreover, [51, 69] illustrates the effectiveness of multi-view training on 2D slices, even by simply averaging multi-planar results, indicating complementary latent information exists in the biases of 2D networks. This inspired us to train 3D multi-view networks with 2D initializations jointly using an additional loss function for multi-view networks which encourages each network to learn from one another.

**Unsupervised Domain Adaptation.** Contrary to semi-supervised learning, domain adaptation problems often contain two datasets that have different distribution. Under unsupervised domain adaptation (UDA) settings, networks are trained from a labeled source domain and an unlabeled target domain. Traditional approaches [20, 44, 52, 58] align domains with statistical constraints. Recent works [14, 17, 18, 24, 25, 55, 61, 72, 73] utilizes adversarial training and self-training to adapt feature training between source domain and target domain. In the field of medical image analysis, [13, 29, 45] have investigated this topic with existing approaches, i.e. adversarial training and self-training.

### 3 Problem Definitions

Before we describe our proposed approach, we firstly discuss the definition and relationships of the three problems, i.e. semi-supervised learning (SSL), unsupervised domain adaptation (UDA) and UDA without data from source domain. Table 1 lists the comparison among the three problems.

**Semi-supervised learning.** Under standard semi-supervised learning (SSL) settings, we denote  $\mathcal{S}$  and  $\mathcal{U}$  as the labeled and unlabeled dataset, respectively. Let  $\mathcal{D} = \mathcal{S} \cup \mathcal{U}$  be the whole available dataset. We denote each labeled data pair as  $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$  and unlabeled data as  $\mathbf{X} \in \mathcal{U}$ . We aim to improve performance on a specific task with unlabeled data. When we consider volumetric medical image segmentation,  $\mathbf{X}$  is a three-dimensional tensor and the ground truth  $\mathbf{Y}$  is a densely-labeled voxel-wise 3D segmentation mask.

**Unsupervised domain adaptation (UDA)** assumes a labeled source domain dataset  $\mathcal{S}$  and an unlabeled target domain dataset  $\mathcal{T}$ , where distributions of data are different but tasks are identical. Our goal is to achieve relatively high performance of a specific task on the target domain. In medical image analysis, domain gaps can result from differences in imaging modalities (e.g. CT / MRI / PET), qualities or imaging protocols (e.g. various machine types and doses of radiation), types of patients (e.g. healthy or with disease), and combinations thereof. The difference between UDA and SSL only lies in data

distributions, so SSL approaches can also be applied to solve UDA problems. In our paper, we illustrate that our proposed approach can effectively handle both problems.

**UDA without data from source domain** was barely investigated in the literature but is an important challenge to be addressed in the field of medical imaging. Here we assume an available pre-trained model from the source domain and unlabeled data from target domain. Differently from UDA, data from source domain is absent.

In our work, we aim to propose a unified approach that is capable of solving the three tasks described above.

settings	dataset 1	dataset 2	same domain?	pre-trained model?
SSL	labeled	unlabeled	yes	-
UDA	labeled	unlabeled	no	-
UDA w/o $\mathcal{S}$	N/A	unlabeled	no	on dataset 1

**Table 1.** The relationship among the three settings i.e. semi-supervised learning (SSL), unsupervised domain adaptation (UDA) and UDA without source domain (UDA w/o  $\mathcal{S}$ ).

## 4 Uncertainty-aware Multi-view Co-training

In this section, we introduce our framework of *uncertainty-aware multi-view co-training* (UMCT) for semi-supervised segmentation and domain adaptation. UMCT is designed to effectively utilize unlabeled data, which is firstly targeted at semi-supervised segmentation of volumetric medical images. In the following sections, we will explain how they are achieved in our 3D framework: a general mathematical formulation of the approach is shown in Sec 4.1; then we demonstrate how to encourage view differences in Sec 4.2, and how to compute the confidence of each view by uncertainty estimation in Sec 4.3, which are the two factors to boost the performance of co-training. Last but not least, the UMCT-DA model is introduced in Sec 4.4 for unsupervised domain adaptation.

### 4.1 Overall Framework

We first consider the task of semi-supervised segmentation for 3D data. Recall that  $\mathcal{S}$  and  $\mathcal{U}$  are the labeled and unlabeled set, respectively. Each labeled data pair is denoted as  $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$  and unlabeled data as  $\mathbf{X} \in \mathcal{U}$ . The ground truth  $\mathbf{Y}$  is a voxel-wise segmentation label map which has the same shape as  $\mathbf{X}$ .

Suppose for each input  $\mathbf{X}$ , we can generate  $N$  different views of 3D data by applying a transformation  $T_i$  (rotation or permutation), resulting in multi-view inputs  $T_i(\mathbf{X})$ ,  $i = 1, \dots, N$ . Such operations will introduce a data-level view difference.  $N$  models  $f_i(\cdot)$ ,  $i = 1, \dots, N$  are then trained over each view of data respectively. For  $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$ , a supervised loss function  $\mathcal{L}_{sup}$  is optimized to measure the similarity between the prediction of each view  $p_i(\mathbf{X}) = T_i^{-1} \circ f_i \circ T_i(\mathbf{X})$  and  $\mathbf{Y}$ :

$$\mathcal{L}_{sup}(\mathbf{X}, \mathbf{Y}) = \sum_{i=1}^N \mathcal{L}(p_i(\mathbf{X}), \mathbf{Y}), \quad (1)$$

where  $\mathcal{L}$  is a standard loss function for segmentation tasks and  $\{p_i(\mathbf{X})\}_{i=1}^N$  are the corresponding voxel-wise prediction score maps after inverse rotation or permutation.

For unlabeled data, we make a co-training assumption under a semi-supervised setting. The co-training strategy assumes the predictions on each view should reach a consensus. So the prediction of each model can act as a pseudo label to supervise other views in order to learn from unlabeled data. However, since the prediction of each view is expected to be diverse after encouraging the view differences, the quality of each view’s prediction needs to be measured before generating trustworthy pseudo labels. This is

---

**Algorithm 1** Uncertainty-aware Multi-view Co-training

---

**Input:**Labeled dataset  $\mathcal{S}$  & Unlabeled dataset  $\mathcal{U}$ *uncertainty-weighted label fusion module* (ULF)  $U_{f_1, \dots, f_n}(\cdot)$ **Output:**Model of each view  $f_1, \dots, f_n$ 

- 1: **while** stopping criterion not met **do**
  - 2:   Sample batch  $b_l = (x_l, y_l) \in \mathcal{S}$  and batch  $b_u = (x_u) \in \mathcal{U}$
  - 3:   Generate multi-view inputs  $T_i(x_l)$  and  $T_i(x_u)$ ,  $i \in \{1, \dots, N\}$
  - 4:   **for**  $i$  in all views **do**
  - 5:     Compute predictions for each view and apply inverse rotation or permutation  
       $p_i(x_l) \leftarrow T_i^{-1} \circ f_i \circ T_i(x_l)$   
       $p_i(x_u) \leftarrow T_i^{-1} \circ f_i \circ T_i(x_u)$
  - 6:   **for**  $i$  in all views **do**
  - 7:     Compute pseudo labels for  $x_u$  with ULF  
       $\hat{y}_i \leftarrow U_{f_1, \dots, f_n}(p_1(x_u), \dots, p_{i-1}(x_u), p_{i+1}(x_u), \dots, p_n(x_u))$
  - 8:    $\mathcal{L}_{sup} = \frac{1}{|b_l|} \sum_{(x_l, y_l) \in b_l} [\sum_i^N \mathcal{L}(p_i(x_l), y_l)]$
  - 9:    $\mathcal{L}_{cot} = \frac{1}{|b_u|} \sum_{(x_u) \in b_u} [\sum_i^N \mathcal{L}(p_i(x_u), \hat{y}_i)]$
  - 10:    $\mathcal{L} = \mathcal{L}_{sup} + \lambda_{cot} \mathcal{L}_{cot}$
  - 11:   Compute gradient of loss function  $\mathcal{L}$  and update network parameters  $\{\theta_i\}$  by back propagation
  - 12: **return**  $f_1, \dots, f_n$
- 

accomplished via *uncertainty-weighted label fusion module* (ULF) introduced in Sec 4.3. With ULF, the co-training loss for unlabeled data can be formulated as:

$$\mathcal{L}_{cot}(\mathbf{X}, \hat{\mathbf{Y}}_i) = \sum_i^N \mathcal{L}(p_i(\mathbf{X}), \hat{\mathbf{Y}}_i), \quad (2)$$

where

$$\hat{\mathbf{Y}}_i = U_{f_1, \dots, f_n}(p_1(\mathbf{X}), \dots, p_{i-1}(\mathbf{X}), p_{i+1}(\mathbf{X}), \dots, p_n(\mathbf{X})) \quad (3)$$

is the pseudo label for the  $i^{th}$  view,  $U_{f_1, \dots, f_n}$  is the ULF computational function, which we will further explain in Sec 4.3.

Overall, the combined loss function is:

$$\sum_{(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}} \mathcal{L}_{sup}(\mathbf{X}, \mathbf{Y}) + \lambda_{cot} \sum_{\mathbf{X} \in \mathcal{U}} \mathcal{L}_{cot}(\mathbf{X}, \hat{\mathbf{Y}}_i). \quad (4)$$

where  $\lambda_{cot}$  is a tunable weight coefficient.

## 4.2 Encouraging View Differences

A successful co-training requires the “views” to be different in order to learn complementary information in the training procedure. In our framework, several techniques are proposed to encourage view differences, both at the data level and the feature level of the neural networks.

**3D multi-view generation.** As stated above, in order to generate multi-view data, we transpose  $\mathbf{X}$  into multiple views by rotations or permutations<sup>3</sup>  $\mathbf{T}$ . For three-view co-training, these can correspond to the coronal, sagittal and axial views in medical imaging, which matches the multi-planar reformatted views that radiologists typically use to analyze the image. Such operation is a natural way to introduce data-level view difference.

---

<sup>3</sup>A permutation rearranges the dimensions of an array in a specific order.



**Asymmetric 3D kernels and 2D initialization.** The co-training assumption encourages models to make similar predictions on both  $\mathcal{S}$  and  $\mathcal{U}$ , which potentially can lead to collapsed neural networks mentioned in [46], a phenomenon that results in a sudden and significant drop in validation accuracy during training of co-training algorithms. In our multi-view settings, this could also happen when the models from different views only learn the permutation or rotation of the kernels, resulting in exactly the same learned feature representation despite the view-point difference. To address this problem, we further encourage view difference at the feature level by designing a task-specific model. We propose to use asymmetric 3D models initialized with 2D pre-trained weights as the backbone network of each view to encourage diverse features for each view learning. In practice, we modify the symmetric 3D convolutional kernels  $n \times n \times n$  into  $n \times n \times 1$  for each branch after the permutation to avoid learning symmetrical representations among views. This structure also makes the model convenient to be initialized with 2D pre-trained weights but fine-tuned in a 3D fashion.

### 4.3 Compute Reliable Psuedo Labels for Unlabeled Data with Uncertainty Estimation

Encouraging view difference means enlarging the variance of each view’s prediction  $\text{var}(p_i(\mathbf{X}))$ . This raises the question of which view we should trust most on unlabeled data during co-training. Bad predictions from one view may hurt the training procedure of other views through pseudo-label assignments. Meanwhile, encouraging to trust a good prediction as a “strong” label from co-training will boost the performance, and lead to improved performance of overall semi-supervised learning. Instead of assigning a pseudo-label for each view directly from the predictions of other views, we propose an adaptive approach, namely *uncertainty-weighted label fusion module* (ULF), to fuse the outputs of different views. ULF is built up of all the views, takes the predictions of each view as input, and then outputs a set of pseudo labels for each view.

Motivated by uncertainty measurements in Bayesian deep networks, we measure the uncertainty of each view branch for each training sample after turning our model into a Bayesian deep network by adding dropout layers. Between the two types of uncertainty candidates – aleatoric and epistemic uncertainties, we choose to compute the epistemic uncertainty that is driven by the lack of training data [30]. Such measurement fits the semi-supervised learning goal: to improve the model’s generalizability by exploring unlabeled data. Suppose  $y$  is the output of a Bayesian deep network, then the epistemic uncertainty can be estimated as:

$$U_e(y) \approx \frac{1}{K} \sum_{k=1}^K \hat{y}_k^2 - \left( \frac{1}{K} \sum_{k=1}^K \hat{y}_k \right)^2, \quad (5)$$

where  $\{\hat{y}_k\}_{k=1}^K$  are a set of sampled outputs. These sampled outputs are obtained by feeding the same input volume into the sub-network defined by  $K$  different random dropout configurations [30]. The voxel-wise epistemic uncertainty is estimated as the statistical variance of the  $K$  predictions. More details are available in Sec 4.5.

With a transformation function  $\mathbf{h}(\cdot)$ , we can transform the uncertainty score into a confidence score  $\mathbf{c}(y) = \mathbf{h}(U_e(y))$ . In practice, we simply define  $\mathbf{h}(U_e(y)) = 1/U_e(y)$ . After normalization over all views, the confidence score will act as the weight for each prediction to assign as a pseudo label for other views. The pseudo label  $\hat{\mathbf{Y}}_i$  assigned for a single view  $i$  can be formulated as

$$\hat{\mathbf{Y}}_i = \frac{\sum_{j \neq i}^N \mathbf{c}(p_j(\mathbf{X})) p_j(\mathbf{X})}{\sum_{j \neq i}^N \mathbf{c}(p_j(\mathbf{X}))}. \quad (6)$$

Thus the pseudo label  $\hat{\mathbf{Y}}_i$  for view  $i$  is computed from predictions from all the other views.

### 4.4 UMCT-DA model for unsupervised domain adaptation

#### Standard unsupervised domain adaptation (UDA)

We extensively validate our approach on unsupervised domain adaptation setting, where the labeled source domain  $\mathcal{S}$  and the unlabeled target domain  $\mathcal{T}$  are available for training. The task is shared between the two domains and the ultimate goal is to achieve good performance on the target domain test data. Despite the domain shift in labeled and unlabeled data, the overall settings of semi-supervised learning (SSL) and unsupervised domain adaptation (UDA) are the same. Hence, we can directly apply our UMCT to solve this problem. The optimization objective can be modified as the follows :

$$\mathcal{L}_{UDA} = \sum_{(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}} \mathcal{L}_{sup}(\mathbf{X}, \mathbf{Y}) + \lambda_{cot} \sum_{\mathbf{X} \in \mathcal{T}} \mathcal{L}_{cot}(\mathbf{X}). \quad (7)$$

where  $\mathcal{S}$  is the labeled source domain and  $\mathcal{T}$  is the unlabeled target domain.

#### UDA without source domain data

Standard UDA methods usually require the existence of source domain data to allow joint training while doing adaptation to the target domain. Here we consider a more challenging setting where source domain data is unavailable and only deep network model (denoted as  $\mathbf{M}_{\mathcal{S}}$ ) pre-trained on source domain is available. In our co-training framework, when source data is unavailable, we can still finetune  $\mathbf{M}_{\mathcal{S}}$  with  $\mathcal{L}_{cot}$  by iteratively refining pseudo labels. The objective function for UDA without source domain data, namely **UMCT-DA**, can be formulated as:

$$\mathcal{L}_{UDA} = \sum_{\mathbf{X} \in \mathcal{T}} \lambda_{cot} \mathcal{L}_{cot}(\mathbf{X}). \quad (8)$$

## 4.5 Implementation Details

**Network Structure.** In practice, we build an encoder-decoder network based on ResNet-18 [22], and modify it into a 3D version. For the encoder part, the first  $7 \times 7$  convolution layer is extended to  $7 \times 7 \times 3$  kernels for low-level 3D feature extraction similar to [38]. All other  $3 \times 3$  convolution layers are simply changed into  $3 \times 3 \times 1$  that can be trained as a 3D convolution layer. In the decoder part, we adopt 3 skip connections from the encoder followed by 3D convolutions to give low-level cues for more accurate boundary prediction needed in segmentation tasks.

**Uncertainty-weighted Label Fusion.** In terms of view confidence estimation, we modify the network into a Bayesian deep network by adding dropouts. We sample  $K = 10$  outputs for each view and compute voxel-wise epistemic uncertainty. Since we are using Dice loss [40], a common loss function for medical image segmentation which is computed on the image level, an image-wise uncertainty estimation is most suitable. We thus sum over the whole volume to estimate the uncertainty for each view. We then simply use the reciprocal for the confidence transformation function  $\mathbf{h} = \frac{1}{\mathbf{c}}$  to compute the confidence score. The resulting pseudo label assigned for each view is a weighted average of all predictions of multiple views based on the normalized confidence score.

**Loss Function.** We extend the Dice loss [40] for multi-class targets as our training objective function:

$$\mathcal{L}_{Dice} = \frac{1}{D} \sum_{d=0}^D \left( 1 - \frac{2 \sum_{i=1}^N y_i^d \hat{y}_i^d}{\sum_{i=1}^N (y_i^d)^2 + \sum_{i=1}^N (\hat{y}_i^d)^2} \right), \quad (9)$$

**Data Pre-Processing.** All the training and testing data are firstly re-sampled to an isotropic volume resolution of  $1.0 \text{ mm}$  for each axis. Data intensities are normalized to have zero mean and unit variance. We adopt patch-based training, and sample training patches of size  $96^3$  with 1:1 ratio between foreground and background.

**Training.** Our training algorithm is shown in Algorithm 1. We firstly train the views separately on the labeled data and then conduct our co-training by fine-tuning the weights. The stochastic gradient descent (SGD) optimizer is used in both stages. In the view-wise training stage, a constant learning rate policy at  $7 \times 10^{-3}$ , momentum at 0.9 and weight decay of  $4 \times 10^{-5}$  for 20k iterations is used. In the co-training stage, we adopt a constant learning rate policy at  $1 \times 10^{-3}$  and train for 5k iterations. The parameter  $\lambda_{cot} = 0.2$  resulted in the best performance which we report here. The batch size is 20



Method	Backbone	10% lab	20% lab
Supervised	3D ResNet-18	66.75	75.79
DMPCT [68]	2D ResNet-101	63.45	66.75
DCT [46] (2v)	3D ResNet-18	71.43	77.54
TCSE [36]	3D ResNet-18	73.87	76.46
Ours (2 views)	3D ResNet-18	75.63	79.77
Ours (3 views)	3D ResNet-18	77.55	80.14
Ours (6 views)	3D ResNet-18	77.87	80.35
Ours (ensemble)	3D ResNet-18	<b>78.77</b>	<b>81.18</b>

**Table 2.** Comparison to other semi-supervised approaches on NIH dataset (DSC, %). Note that we use the same backbone network as [36] [46]. Here, “2v” means two views. For our approach, we report the average of all single views’ DSC score for a fair comparison (2 views to 6 views), as well as multi-view ensemble results. “10% lab” and “20% lab” mean the percentage of labeled data used for training.

in co-training, among which 4 images are labeled and 16 are unlabeled, maintaining a ratio of labeled and unlabeled to be 1:4. Our framework is implemented in PyTorch. For 3D ResNet-18 on NIH dataset, the whole co-training procedure takes  $\sim 24$  hours on one single NVIDIA Titan RTX GPU with 24 GB memory. In our implementation, training occupies  $\sim 15$ GB GPU memory in total.

**Testing.** In the testing phase, there are two choices to finalize the output results: either to choose one single view prediction or to ensemble the predictions of the multi-view outputs with majority voting. We will report both results in subsequent sections for fair comparisons with the baselines since the multiple view networks can be thought of being similar to the ensemble of several single view models. The experimental results show that our model improves the performance in both settings (single view and multi-view ensemble) over all the other approaches. We use sliding-window testing and re-sample our testing results back to the original image resolution to obtain the final results. Testing time for each case ranges from 1 minute to 5 minutes depending on the size of the input volume.

## 5 Experiments

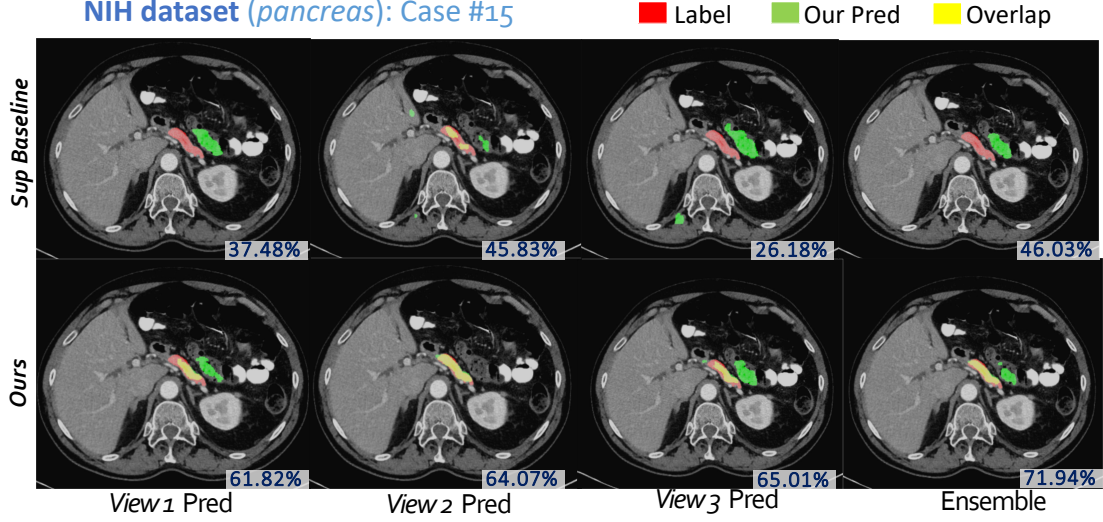
In this section, we first evaluate our framework under semi-supervised settings on the NIH pancreas segmentation dataset [50] with cases from a healthy patient population (e.g. kidney donors<sup>4</sup>); and an multi-organ segmentation dataset [19] with eight abdominal organs [34] with conditions mostly unrelated to the organs of interest (e.g. colorectal cancer or ventral hernia<sup>5</sup>). We will provide detailed experiments, including ablation studies, on the former dataset. Note that the volumes come from different patients in each dataset and were separated at the patient-level for the different training, validation and testing splits. Next, we validate the capability of our approach on the task of unsupervised domain adaptation, which is critical but under-investigated in the field of medical image analysis. The multi-organ segmentation dataset serves as source data. The targets of adaptation include two pathological organ datasets i.e. pancreas and liver datasets in the Medical Segmentation Decathlon (MSD) [57], which both can include tumors in their respective organs. More strictly, we also evaluate our approach under the situation where source data is inaccessible (UDA without source data).

### 5.1 NIH Pancreas Segmentation Dataset

The NIH pancreas segmentation dataset contains 82 abdominal CT volumes. The width and height of each volume are 512, while the axial view slice number can vary from 181 to 466. Under semi-supervised settings, the dataset is randomly split into 20 testing cases and 62 training cases. We report the results

<sup>4</sup><https://wiki.cancerimagingarchive.net/display/Public/Pancreas-CT>

<sup>5</sup><https://www.synapse.org/#!/Synapse:syn3193805/wiki/217789>



**Figure 3.** 2D visualizations for one example of NIH pancreas segmentation dataset 10% labeled data setting. The first row is the supervised baseline and the second row is the prediction after our 3-view co-training. DSC scores are largely improved. Best viewed in color.

of 10% labeled training cases (6 labeled and 56 unlabeled), 20% labeled training cases (12 labeled and 50 unlabeled) and 100% labeled training cases.

### 5.1.1 Results

In Table 2, we first report the average of all single views’ DSC score for a fair comparison (2 views to 6 views, last 2-4 rows), which can be viewed as the average performance of one single view model. Then we report the multi-view ensemble results (6 view ensemble, last row), where we align the multi-view prediction maps to the same view (axial) and average the prediction maps at each pixel to make a final prediction. For 2-view co-training, we use the axial and coronal views. For 3-view co-training, we use the axial, coronal and sagittal view. For 6-view co-training, we use the axial, coronal and sagittal view as well as the horizontal flip version of the three views ( $3 \times 2 = 6$ ). The first row is the supervised training results, using only labeled data and trained on the axial view. The segmentation accuracy is evaluated by Dice-Sørensen coefficient (DSC). A large margin improvement over the fully supervised baselines in terms of single view performance can be observed, proving that our approach effectively leverages the unlabeled data. A Wilcoxon signed-rank test comparing to the supervised baseline’s results (20% labeling) shows significant improvements of our approach with a  $p$ -value of 0.0022. Fig. 3 shows 3 cases in 2D and 3D with ITK-SNAP [66]. In addition, our model is compared with the state-of-the-art semi-supervised approach of deep co-training [46] and recent semi-supervised medical segmentation approaches. In particular, we compare to [36] who extended the  $\pi$  model [33] with transformation consistent constraints; and [68] who extended the self-training procedure by iteratively updating pseudo labels on unlabeled data using a fusion of three 2D networks trained on cross-sectional views. The results reported in Table 2 are based on our careful re-implementations in order to allow a fair comparison.

The implementations of [46] and [36] are operated on the axial view of our single view branch with the same backbone structure (our customized 3D ResNet-18 model). Our co-training approach achieve about 4% gain in the 10% labeled and 90% unlabeled settings. We also find that improvements of other approaches are small in the 20% settings (only 1% compared to the baseline), while ours still is capable to achieve a reasonable performance gain with the growing number of labeled data. For [68] with a 2D approach, their experiment is conducted on 50 labeled cases. We modify their backbone network (FCN [39]) into DeepLab v2 [7], in order to fit our stricter settings (6 and 12 labeled cases).

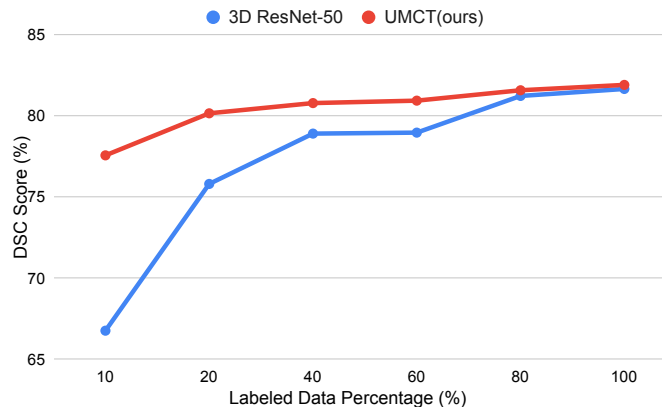
This modification leads to an improvement of 3% in 100% fully supervised training (from 73% to 76%). Their approach outputs the result after using an ensemble with majority voting of three slice-wise 2D models obtained from their semi-supervised training approach..

Since the main difference in two-view learning between our approach and [46] is the way of encouraging view differences, the results illustrate the effectiveness of our multi-view analysis combined with asymmetric feature learning on 3D co-training. With more views, our uncertainty-weighted label fusion can further improve co-training performance. We will report ablation studies later in this section.

### 5.1.2 Analysis and ablation studies

#### Data utilization efficiency

We perform a study on data utilization efficiency of our approach compared to the baseline fully-supervised network (3D ResNet-18). Fig. 4 shows the performance change according to labeled data proportion on NIH pancreas segmentation. From the plot, one can see that when labeled data is over 80%, simple supervised training (with 3D ResNet-18) suffices. Note that our approach with 20% labeled data (DSC 80.35%) performs better than 60% supervised training (DSC 78.95%). At such a performance, our approach can save  $\sim 70\%$  of the labeling efforts.



**Figure 4.** Performance plot of our semi-supervised approach over the fully-supervised baseline on different labeled data ratio.

#### Effect of backbone structure

Our backbone selection (2D-initialized, heavily asymmetric 3D architecture) will introduce 2D biases in the training phase while benefiting from such 2D pre-trained models. We have claimed that we can utilize the complementary information from 3-view networks while exploring the unlabeled data with UMCT. We give an ablation study on the network structure, which contains a V-Net [40], a common 3D segmentation network with all symmetrical kernels in all dimensions. Such network also shares a similar amount of parameters with our customized 3D ResNet-18, see Table 3. The results of V-Net show that our multi-view co-training can be generally and successfully applied to 3D networks. Although the results of fully supervised parts are similar, our ResNet-18 outperforms V-Net by more than 1%, illustrating that our asymmetric design, encouraging view differences, brings advantages over traditional 3D deep networks.

#### Uncertainty-weighted label fusion (ULF)

ULF acts as an important role in pruning out bad predictions and keeping good ones as supervision to train other views. Table 4 gives the single view results in multiple views experiments. The performance becomes better with more views. For two views, ULF is not applicable since we can only obtain one view prediction as a pseudo label for the other view. For three views and six views, ULF helps boost the performance, illustrating the effectiveness of our proposed approach for view confidence estimation.

Backbone	Params	MACs	10% Sup	Ours
V-Net	9.44M	41.40G	66.97	76.89
3D ResNet-18	11.79M	17.08G	66.76	77.55
3D ResNet-50	27.09M	23.03G	67.96	<b>78.74</b>

**Table 3.** Ablation studies on backbone structures (3 views UMCT). "Params" is short for parameters and "MACs" is short for multiply-accumulate operations. "10% Sup" means supervised training with 10% labeled data. A Wilcoxon signed-rank test reveals significant improvements ( $p \ll 0.01$ ) of our 3D ResNets over V-Net in the last column, illustrating our asymmetrical design is beneficial for our co-training method.

Views	DSC(%)
2 views	75.63
3 views	76.49
3 views + ULF	<b>77.55</b>
6 views	76.94
6 views + ULF	<b>77.87</b>

**Table 4.** On uncertainty-weighted label fusion (ULF) with difference views in training (10% labeled data, 3D ResNet-18).

## 5.2 Multi-organ Segmentation Dataset

Next, we validate our approach on multi-organ datasets. The dataset we use is a multi-organ re-annotated version from [19], combining two public datasets - The Cancer Image Archive (TCIA) Pancreas-CT data set [50] and Beyond the Cranial Vault (BTCV) Abdomen data set [34]. We perform 4-fold cross validation on 90 cases in total. In each fold, we then randomly split the training cases into our labeled set  $\mathcal{S}$  and unlabeled set  $\mathcal{U}$ . We train our models on different labeled data ratio 10%, 20%, which approximately corresponds to (7,81), (13,75) of (labeled, unlabeled) data pairs and validate on  $\sim 22$  cases in each fold. Results are shown in Table 5 and an example is shown in Fig 5.

Our approach improves consistently over almost every organ under every labeled-unlabeled ratio of data. The results illustrate the ability of our approach to handle the situation of complex multi-organ settings.

## 5.3 Unsupervised domain adaptation from multi-organ segmentation to MSD Dataset

We aim to unsupervisedly adapt a model trained on TCIA multi-organ dataset to pancreas and liver cases in Medical Decathlon Challenge [57] that can exhibit tumors. The target domains contain a shift from the source domain because of the differences in (i) image quality and contrast, and (ii) textures due to the existence of pancreatic/hepatic tumors.

MSD pancreas dataset contains 282 CT scans in portal venous phase, all of which are pathological cases with pancreas and tumor annotation. We randomly split the whole dataset into 200 cases for training (without label) and 81 cases for validation. Since the source domain (multi-organ dataset) only contains healthy pancreas, we aim at segmenting the whole pancreas region (combining pancreas and tumor together). For MSD liver dataset, we aim at segmenting the whole liver region as well, with a random split of 100 training cases (unlabeled) and 31 cases for validation. 118 out of 131 cases contains hepatic tumor.

Table 6 shows the results of unsupervised domain adaptation experiments. The first row is the segmentation performance (in terms of DSC) on pancreas and liver of the original multi-organ validation set. From the second row to the last, the results are DSC scores on MSD liver / pancreas validation set. In standard UDA settings (UMCT w/ source), significant improvements are achieved with our approach (1.12% in liver and 4.70% in pancreas), compared to source only version (direct Test on MSD). Due to

Experiment setups	spleen	l.kidney	gallbladder	esophagus	liver	stomach	pancreas	duodenum
Supervised (upper bound)	94.20	93.90	71.89	66.74	94.78	88.60	81.46	71.29
10% lab	88.46	90.88	42.77	52.41	91.33	76.50	69.63	52.52
10% lab+90% unlabeled (ours)	<b>91.14</b>	<b>92.35</b>	<b>58.29</b>	<b>57.61</b>	<b>92.23</b>	<b>79.67</b>	<b>73.86</b>	<b>57.50</b>
20% lab	92.77	92.29	63.60	61.84	93.95	82.56	75.60	60.26
20% lab+80% unlabeled (ours)	92.80	<b>92.99</b>	<b>66.29</b>	<b>65.01</b>	93.93	<b>83.67</b>	<b>77.91</b>	<b>63.34</b>

**Table 5.** Experimental results for semi-supervised learning on a multi-organ dataset under four fold cross-validation. "lab" is short for "labeled" and "unlab" is short for "unlabeled". Supervised results (first row) uses 100% labeled training data in the training set, which is the upper bound but requires 100% annotation. 10% lab means we only use 10% training data with annotation for supervised training. 10%lab + 90% unlabeled (ours) means we use 10% labeled data and 90% unlabeled data for our co-training method. Results are reported via 4-fold cross-validation. Numbers in **bold** indicate significant improvement over supervised counterparts by Wilcoxon signed rank tests ( $p < 0.01$ ).

Train	Test	Method	liver	pancreas
MO (L)	MO	Supervised	95.59	81.69
MO (L)	MSD	Supervised	92.78	70.23
MO (L) + MSD (U)	MSD	Adv training	93.35	71.23
		Self-training	92.67	71.38
		UMCT	<b>93.90</b>	<b>74.93</b>
MO model + MSD (U)	MSD	UMCT-DA	92.98	74.38

**Table 6.** Experiments of unsupervised domain adaptation (UDA). The source domain is Multi-organ dataset (denoted as "MO") and target domains are MSD liver dataset and pancreas dataset. "L" represents this dataset is labeled and "U" means the opposite.

the superior performance of self-training based approaches [72, 73] for unsupervised domain adaptation, we implement a vanilla self-training method under our settings (denoted as "Self-training"). We first test the model on the unlabeled set and then use the prediction as pseudo labels to train on the whole data set. We iterate these two steps every 1k iterations and trains for 5k iterations in total, which is in line with the proposed co-training scheme. We also implement another baseline approach (AdaptSegNet [61], denoted as "Adv training"), which applies adversarial training onto the predicted masks of semantic segmentation. The segmentation network serves as the generator to output segmentation masks with segmentation loss and tries to fool a patch-based discriminator (a 3D version of the discriminator used in AdaptSegNet [61]) with GAN loss [21]. The discriminator is also trained jointly to distinguish between the predicted mask and the ground-truth mask on unlabeled data. Our approach significantly outperforms all of the baselines. We also show one example for each organ in Fig 6.

The last row gives the results in the absence of source domain data. Under such condition, only a pre-trained source domain model and unlabeled target domain data are available. Our UMCT-DA model (last row) is able to solve this problem by only using the co-training loss  $L_{cot}$  to train on the target dataset, and achieve comparable results with standard UDA settings, even without source domain data.

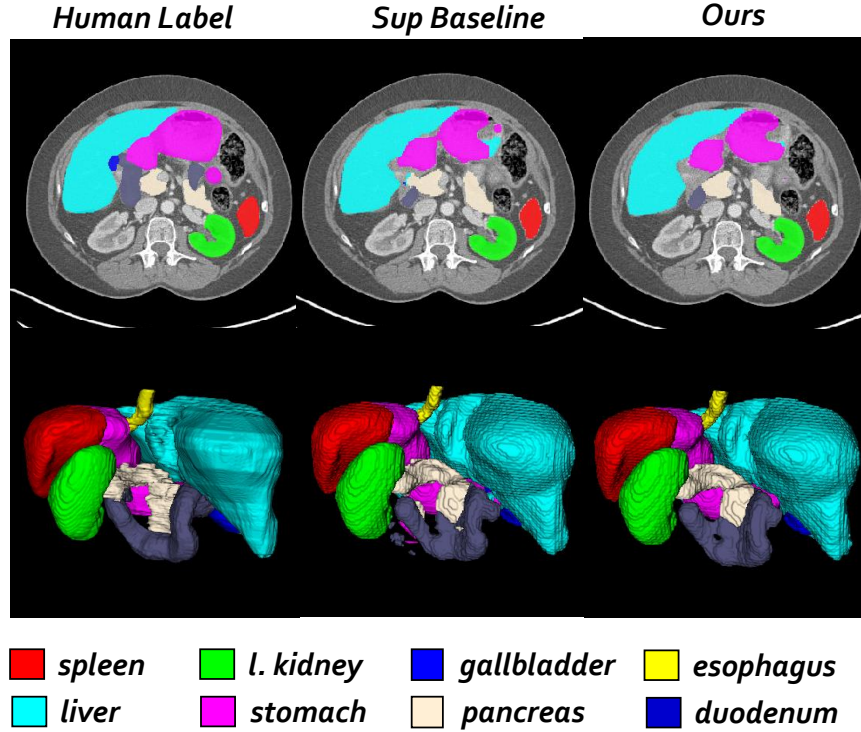
## 6 Discussions

### 6.1 Impact on large-scale benchmarks

Under fully supervised training, our team NVDLMED was ranked the 3<sup>rd</sup> place in the first phase and the **2<sup>nd</sup> place in the final validation phase of Medical Segmentation Decathlon Challenge** [57] (challenge leaderboard available<sup>6</sup>). We applied our 3-view co-training framework taken from axial, coronal

<sup>6</sup><http://medicaldecathlon.com/results.html>





**Figure 5.** An example of semi-supervised multi-organ segmentation.

and sagittal views to ten medical image segmentation tasks simultaneously. The winning team [27] applied heavy model selection and ensemble by cross-validation on the training set, while we used a fixed framework without complicated data augmentation. Although not originally targeted at improving the performance of fully supervised training, our approach still illustrated the effectiveness and robustness of co-training from multiple views.

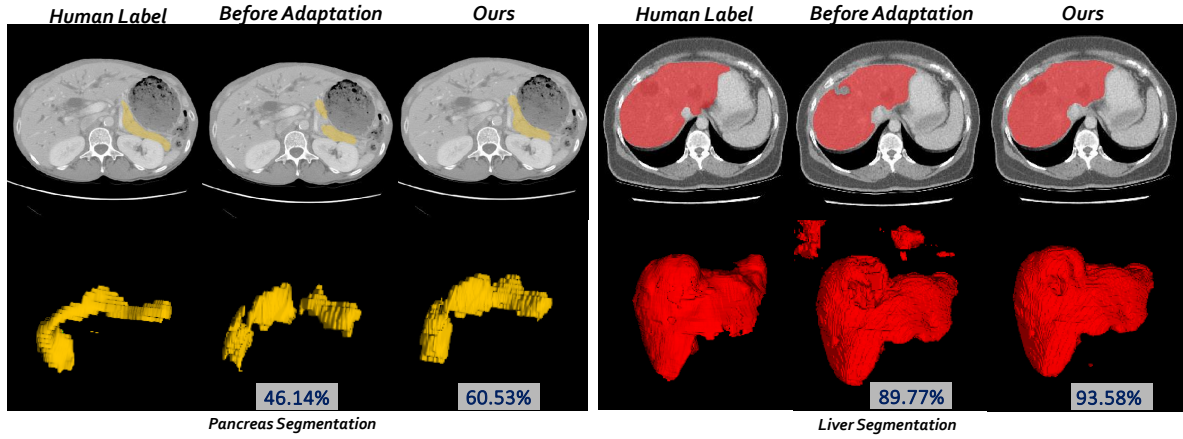
## 6.2 Magnitude of domain shift

In this work, domain shift mainly lies in various sources of CT scans and the pathological/healthy status of abdominal organs. We consider it a reasonable domain shift from different CT datasets originating from different hospitals and patient populations. Typically, this means that when directly transferring a model from one to another (TCIA to MSD in our case), the performance drops significantly (liver 95% to 92%, pancreas 81% to 70% average Dice). While this shift is relatively small compared to, for example, cross modality testing (say CT to MRI), it is unacceptable when considering these models for potential clinical applications. Considering the importance of this topic, we shed light on how well our semi-supervised approach performs on UDA tasks, given their similarity (discussed in Sec 3). Other types of domain shifts of medical images, though not investigated in this paper, are also of great importance. Investigation of domain adaptation under larger domain shifts such as modality changes (e.g. CT to MRI adaptation), contrast and resolution issues remains an active research topic.

## 7 Summary & Conclusion

In this paper, we presented *uncertainty-aware multi-view co-training* (UMCT), aimed at semi-supervised learning and domain adaptation. We extended dual view co-training and deep co-training into 3D volumetric image data by analysing from different view-points, then estimating uncertainty and finally





**Figure 6.** 2D and 3D visualizations for unsupervised domain adaptation of pancreas (left) and liver segmentation (right).

enforcing multi-view consistency on large scale unlabeled data. Our approach was first validated on NIH pancreas dataset, where we outperformed other approaches by a large margin. We further applied our approach to multi-organ datasets and found significant improvements for each organ. Finally, we adapted the multi-organ dataset to MSD pathological pancreas and liver in an unsupervised manner. Our UMCT-DA model achieved good performance even in the absence of source domain data, illustrating strong potential for real-world applications in medical image segmentation.

In the future, we plan to conduct further research in the following aspects. Currently the views of co-training are fixed and pre-defined, so one feasible idea is to incorporate more views and random views. This could increase the robustness of our model and lead to better performance. For domain adaptation, we will also try to explore co-training based approaches on other types of domain shifts including but not limited to image modality changes and contrast variants. We believe co-training based approaches will make a contribution to large scale medical image analysis with limited human annotations.

## References

1. P. Bachman, O. Alsharif, and D. Precup. Learning with pseudo-ensembles. In *Advances in Neural Information Processing Systems*, pages 3365–3373, 2014.
2. 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 *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 253–260. Springer, 2017.
3. M. Belkin, P. Niyogi, and V. Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research*, 7(Nov):2399–2434, 2006.
4. A. Blake, R. Curwen, and A. Zisserman. A framework for spatiotemporal control in the tracking of visual contours. *International Journal of Computer Vision*, 11(2):127–145, 1993.
5. A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100. ACM, 1998.
6. D.-D. Chen, W. Wang, W. Gao, and Z.-H. Zhou. Tri-net for semi-supervised deep learning. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 2014–2020. AAAI Press, 2018.

- 
7. L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE transactions on pattern analysis and machine intelligence, 40(4):834–848, 2018.
  8. L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587, 2017.
  9. M. Chen, K. Q. Weinberger, and J. Blitzer. Co-training for domain adaptation. In Advances in neural information processing systems, pages 2456–2464, 2011.
  10. V. Cheplygina, M. de Bruijne, and J. P. Pluim. Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. arXiv preprint arXiv:1804.06353, 2018.
  11. Z. Dai, Z. Yang, F. Yang, W. W. Cohen, and R. R. Salakhutdinov. Good semi-supervised learning that requires a bad gan. In Advances in neural information processing systems, pages 6510–6520, 2017.
  12. N. Dong, M. Kampffmeyer, X. Liang, Z. Wang, W. Dai, and E. Xing. Unsupervised domain adaptation for automatic estimation of cardiothoracic ratio. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 544–552. Springer, 2018.
  13. Q. Dou, C. Ouyang, C. Chen, H. Chen, and P.-A. Heng. Unsupervised cross-modality domain adaptation of convnets for biomedical image segmentations with adversarial loss. arXiv preprint arXiv:1804.10916, 2018.
  14. G. French, M. Mackiewicz, and M. Fisher. Self-ensembling for visual domain adaptation. 2018.
  15. Y. Gal. Uncertainty in deep learning. University of Cambridge, 2016.
  16. Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In international conference on machine learning, pages 1050–1059, 2016.
  17. Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. The Journal of Machine Learning Research, 17(1):2096–2030, 2016.
  18. Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. The Journal of Machine Learning Research, 17(1):2096–2030, 2016.
  19. 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. IEEE transactions on medical imaging, 37(8):1822–1834, 2018.
  20. B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, pages 2066–2073. IEEE, 2012.
  21. I. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In International Conference on Learning Representations, 2015.
  22. K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
  23. X. He, R. S. Zemel, and M. Á. Carreira-Perpiñán. Multiscale conditional random fields for image labeling. In Computer vision and pattern recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE computer society conference on, volume 2, pages II–II. IEEE, 2004.
-

- 
24. J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell. Cycada: Cycle-consistent adversarial domain adaptation. arXiv preprint arXiv:1711.03213, 2017.
  25. J. Hoffman, D. Wang, F. Yu, and T. Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. arXiv preprint arXiv:1612.02649, 2016.
  26. G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4700–4708, 2017.
  27. F. Isensee, J. Petersen, A. Klein, D. Zimmerer, P. F. Jaeger, S. Kohl, J. Wasserthal, G. Koehler, T. Norajitra, S. Wirkert, et al. nnu-net: Self-adapting framework for u-net-based medical image segmentation. arXiv preprint arXiv:1809.10486, 2018.
  28. J. Jiang, Y.-C. Hu, N. Tyagi, P. Zhang, A. Rimner, G. S. Mageras, J. O. Deasy, and H. Veeraraghavan. Tumor-aware, adversarial domain adaptation from ct to mri for lung cancer segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 777–785. Springer, 2018.
  29. K. Kamnitsas, C. Baumgartner, C. Ledig, V. Newcombe, J. Simpson, A. Kane, D. Menon, A. Nori, A. Criminisi, D. Rueckert, et al. Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In International conference on information processing in medical imaging, pages 597–609. Springer, 2017.
  30. A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In Advances in neural information processing systems, pages 5574–5584, 2017.
  31. A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
  32. A. Kumar, P. Sattigeri, and T. Fletcher. Semi-supervised learning with gans: Manifold invariance with improved inference. In Advances in Neural Information Processing Systems, pages 5534–5544, 2017.
  33. S. Laine and T. Aila. Temporal ensembling for semi-supervised learning. International Conference on Learning Representations (ICLR), 2017.
  34. B. Landman, Z. Xu, J. Eugenio Igelsias, M. Styner, T. Langerak, and A. Klein. Miccai multi-atlas labeling beyond the cranial vault–workshop and challenge, 2015.
  35. X. Li, H. Chen, X. Qi, Q. Dou, C.-W. Fu, and P. A. Heng. H-denseunet: Hybrid densely connected unet for liver and liver tumor segmentation from ct volumes. IEEE Transactions on Medical Imaging, 2017.
  36. X. Li, L. Yu, H. Chen, C.-W. Fu, and P.-A. Heng. Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model. BMVC, 2018.
  37. Y. Li, C. Guan, H. Li, and Z. Chin. A self-training semi-supervised svm algorithm and its application in an eeg-based brain computer interface speller system. Pattern Recognition Letters, 29(9):1285–1294, 2008.
  38. S. Liu, D. Xu, S. K. Zhou, O. Pauly, S. Grbic, T. Mertelmeier, J. Wicklein, A. Jerebko, W. Cai, and D. Comaniciu. 3d anisotropic hybrid network: Transferring convolutional features from 2d images to 3d anisotropic volumes. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 851–858. Springer, 2018.
  39. J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 3431–3440, 2015.
-

- 
40. F. Milletari, N. Navab, and S.-A. Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In 2016 Fourth International Conference on 3D Vision (3DV), pages 565–571. IEEE, 2016.
  41. T. Miyato, S.-i. Maeda, S. Ishii, and M. Koyama. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. IEEE transactions on pattern analysis and machine intelligence, 2018.
  42. P. Mlynarski, H. Delingette, A. Criminisi, and N. Ayache. Deep learning with mixed supervision for brain tumor segmentation. Journal of Medical Imaging, 6(3):034002, 2019.
  43. D. Nie, Y. Gao, L. Wang, and D. Shen. Asdnet: Attention based semi-supervised deep networks for medical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 370–378. Springer, 2018.
  44. S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang. Domain adaptation via transfer component analysis. IEEE Transactions on Neural Networks, 22(2):199–210, 2010.
  45. C. S. Perone, P. Ballester, R. C. Barros, and J. Cohen-Adad. Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. NeuroImage, 194:1–11, 2019.
  46. S. Qiao, W. Shen, Z. Zhang, B. Wang, and A. Yuille. Deep co-training for semi-supervised image recognition. In Proceedings of the European Conference on Computer Vision (ECCV), pages 135–152, 2018.
  47. A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-supervised learning with ladder networks. In Advances in Neural Information Processing Systems, pages 3546–3554, 2015.
  48. O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention, pages 234–241. Springer, 2015.
  49. C. Rosenberg, M. Hebert, and H. Schneiderman. Semi-supervised self-training of object detection models. WACV/MOTION, 2, 2005.
  50. 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.
  51. 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. Medical image analysis, 45:94–107, 2018.
  52. K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In European conference on computer vision, pages 213–226. Springer, 2010.
  53. M. Sajjadi, M. Javanmardi, and T. Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In Advances in Neural Information Processing Systems, pages 1163–1171, 2016.
  54. M. P. Shah, S. Merchant, and S. P. Awate. Ms-net: mixed-supervision fully-convolutional networks for full-resolution segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 379–387. Springer, 2018.
  55. R. Shu, H. H. Bui, H. Narui, and S. Ermon. A dirt-t approach to unsupervised domain adaptation. 2018.
  56. K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
-

- 
57. A. L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. van Ginneken, A. Kopp-Schneider, B. A. Landman, G. Litjens, B. Menze, et al. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. arXiv preprint arXiv:1902.09063, 2019.
  58. B. Sun and K. Saenko. Subspace distribution alignment for unsupervised domain adaptation. In BMVC, volume 4, pages 24–1, 2015.
  59. C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1–9, 2015.
  60. A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In Advances in neural information processing systems, pages 1195–1204, 2017.
  61. Y.-H. Tsai, W.-C. Hung, S. Schuster, K. Sohn, M.-H. Yang, and M. Chandraker. Learning to adapt structured output space for semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 7472–7481, 2018.
  62. Y. Xia, F. Liu, D. Yang, J. Cai, L. Yu, Z. Zhu, D. Xu, A. Yuille, and H. Roth. 3d semi-supervised learning with uncertainty-aware multi-view co-training. In The IEEE Winter Conference on Applications of Computer Vision, pages 3646–3655, 2020.
  63. Y. Xia, L. Xie, F. Liu, Z. Zhu, E. K. Fishman, and A. L. Yuille. Bridging the gap between 2d and 3d organ segmentation with volumetric fusion net. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 445–453. Springer, 2018.
  64. 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.
  65. 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. 2017.
  66. P. A. Yushkevich, J. Piven, H. Cody Hazlett, R. Gimpel Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. Neuroimage, 31(3):1116–1128, 2006.
  67. H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), pages 2881–2890, 2017.
  68. Y. Zhou, Y. Wang, P. Tang, W. Shen, E. K. Fishman, and A. L. Yuille. Semi-supervised multi-organ segmentation via multi-planar co-training. WACV, 2019.
  69. Y. Zhou, L. Xie, W. Shen, Y. Wang, E. K. Fishman, and A. L. Yuille. A fixed-point model for pancreas segmentation in abdominal ct scans. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 693–701. Springer, 2017.
  70. Z.-H. Zhou and M. Li. Semi-supervised regression with co-training. In IJCAI, volume 5, pages 908–913, 2005.
  71. Z.-H. Zhou and M. Li. Tri-training: Exploiting unlabeled data using three classifiers. IEEE Transactions on knowledge and Data Engineering, 17(11):1529–1541, 2005.
  72. Y. Zou, Z. Yu, X. Liu, B. Kumar, and J. Wang. Confidence regularized self-training. In Proceedings of the IEEE International Conference on Computer Vision, pages 5982–5991, 2019.
  73. Y. Zou, Z. Yu, B. Vijaya Kumar, and J. Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In Proceedings of the European Conference on Computer Vision (ECCV), pages 289–305, 2018.
-