




Bootstrapping Semi-supervised Medical Image Segmentation with Anatomical-Aware Contrastive Distillation

Chenyu You¹, Weicheng Dai², Yifei Min⁵, Lawrence Staib^{1,3,4},
and James S. Duncan^{1,3,4,5}

¹ Department of Electrical Engineering, Yale University, New Haven, USA

chenyu.you@yale.edu

² Department of Computer Science and Engineering, New York University,
New York, USA

³ Department of Biomedical Engineering, Yale University, New Haven, USA

⁴ Department of Radiology and Biomedical Imaging, Yale University,
New Haven, USA

⁵ Department of Statistics and Data Science, Yale University, New Haven, USA

Abstract. Contrastive learning has shown great promise over annotation scarcity problems in the context of medical image segmentation. Existing approaches typically assume a balanced class distribution for both labeled and unlabeled medical images. However, medical image data in reality is commonly imbalanced (*i.e.*, multi-class label imbalance), which naturally yields blurry contours and usually incorrectly labels rare objects. Moreover, it remains unclear whether all negative samples are equally negative. In this work, we present **ACTION**, an Anatomical-aware **ConT**rastive **dI**stillati**ON** framework, for semi-supervised medical image segmentation. Specifically, we first develop an iterative contrastive distillation algorithm by softly labeling the negatives rather than binary supervision between positive and negative pairs. We also capture more semantically similar features from the randomly chosen negative set compared to the positives to enforce the diversity of the sampled data. Second, we raise a more important question: Can we really handle imbalanced samples to yield better performance? Hence, the *key innovation* in ACTION is to learn global semantic relationship across the entire dataset and local anatomical features among the neighbouring pixels with minimal additional memory footprint. During the training, we introduce anatomical contrast by actively sampling a sparse set of hard negative pixels, which can generate smoother segmentation boundaries and more accurate predictions. Extensive experiments across two benchmark datasets and different unlabeled settings show that ACTION significantly outperforms the current state-of-the-art semi-supervised methods.

Keywords: Contrastive Learning · Knowledge Distillation · Active Sampling · Semi-Supervised Learning · Medical Image Segmentation




Fig. 1. Examples of two benchmarks (*i.e.*, ACDC and LiTS) showing the large variations of class distribution.

1 Introduction

Manually labeling sufficient medical data with pixel-level accuracy is time-consuming, expensive, and often requires domain-specific knowledge. To bypass the cost for labeled data, semi-supervised learning (SSL) is one of the promising, conventional ways to train models with weaker forms of supervision, given a large amount of unlabeled data. Existing SSL methods include adversarial training [12, 28, 32, 33, 37], deep co-training [21, 38], mean teacher schemes [23, 36], multi-task learning [4, 11, 16, 31], and contrastive learning [3, 9, 29, 30, 34, 35].

Among the aforementioned methods, contrastive learning [5, 8] has recently prevailed for DNNs to rich visual representations from unlabeled data. The predominant promise of label-free learning is to capture the similar semantic relationship and anatomical structure between neighboring pixels from massive unannotated data. However, going to realistic clinical scenarios will have the following shortcomings. First, different medical images share similar anatomical structures, but prior methods follow the standard contrastive learning [5, 8] in comparing positive and negative pairs by binary supervision. That naturally leads to the issues of false negatives in representation learning [10, 24], which would hurt segmentation performance. Second, the underlying class distribution of medical image data is highly imbalanced, as illustrated in Fig. 1. It is well known that such imbalanced distribution will severely hurt the segmentation quality [14], which may result in blurry contours and mis-classify minority classes due to the occurrence frequencies [39]. That naturally questions whether contrastive learning can still work well in those imbalance scenarios.

In this work, we present a principled framework called **Anatomical-aware ConTrastive dIstillatiON (ACTION)**, for multi-class medical image segmentation. In contrast to prior work [3, 9, 35] which directly distinguish two image samples of the similar anatomical features that are in the negative pairs, the **key innovation** in ACTION is to actively learn more balanced representations by *dynamically* selecting samples that are semantically similar to the queries, and contrasting the model’s *anatomical-level* features with the target model’s in **imbalanced** and **unlabeled** clinical scenarios. Specifically, we introduce two strategies to improve overall segmentation quality: (1) we believe that all negative samples are not equally negative. Thus, we propose relaxed contrastive learning by using soft labeling on the negatives. In other words, we randomly sample a set of image samples as anchor points to ensure **diversity** in the set of

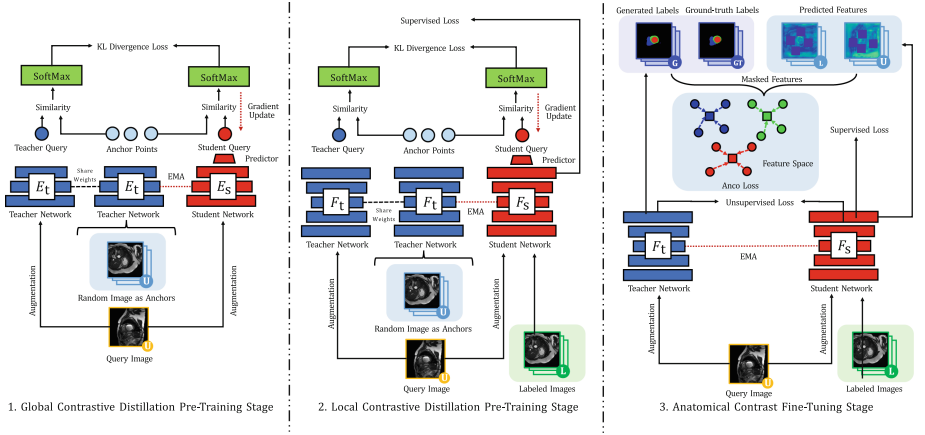


Fig. 2. Overview of the ACTION framework including three stages: (1) global contrastive distillation pre-training used in existing works, (2) our proposed local contrastive distillation pre-training, and (3) our proposed anatomical contrast fine-tuning.

sampled examples. Then the teacher model predicts the underlying probability distribution over neighboring samples by computing the anatomical similarities between the query and the anchor points in the memory bank, and the student model tries to learn from the teacher model. Such a strategy is much more regularized by mincing the same neighborhood anatomical similarity to improve the quality of the anatomical features; (2) to create strong contrastive views on anatomical features, we introduce **AnCo**, another new contrastive loss designed at the anatomical level, by sampling a set of pixel-level representation as queries, and pulling them closer to the mean feature of all representations in a class (positive keys), and pulling other representations apart from other class (negative keys). In addition to reducing the high memory footprint and computation complexity, we use active sampling to dynamically select a sparse set of queries and keys during the training. We apply ACTION on two benchmark datasets under different unlabeled settings. Our experiments show that ACTION can dramatically outperform the state-of-the-art SSL methods. We believe that our proposed ACTION can be a strong baseline for the related medical image analysis tasks in the future.

2 Method

Framework Overview. The workflow of our proposed ACTION is illustrated in Fig. 2. By default, ACTION is built on the BYOL pipeline [7] which is originally designed for image classification tasks, and for a fair comparison, we also follow the setting in [3] such as using 2D U-Net [22] as the backbone and non-linear projection heads H . The **main differences** between our proposed ACTION and [3,9] are as follows: (1) the addition of a predictor $g(\cdot)$ to the student network to avoid collapsed solutions; (2) the utilization of a slow-moving average

of the student network as the teacher network for more semantically compact representations; (3) the use of the output probability rather than logits effectively and semantically constrains the distance between the anatomical features from the imbalanced data (*i.e.*, multi-class label imbalance cases); (4) we propose to contrast the query image features with other random image features at the global and local level, rather than only two augmented versions of the same image features; and (5) we design a novel unsupervised anatomical contrastive loss to provide additional supervision on hard pixels.

Let (X, Y) be a training dataset including N labeled image slices and M unlabeled image slices, with training images $X = \{x_i\}_{i=1}^{N+M}$ and the C -class segmentation labels $Y = \{y_i\}_{i=1}^N$. Our backbone $F(\cdot)$ (2D U-Net) consists of an encoder network $E(\cdot)$ and a decoder network $D(\cdot)$. The training procedure of ACTION includes three stages: (i) global contrastive distillation pre-training, (ii) local contrastive distillation pre-training, and (iii) anatomical contrast fine-tuning. In the first two stages, we use global contrastive distillation to train E on unlabeled data to learn global-level features, and use local contrastive distillation to train E and D on labeled and unlabeled data to learn local-level features.

Global Contrastive Distillation Pre-training. We follow a similar setting in [24]. Given an input query image $q \in \{x_i\}_{i=N+1}^{N+M}$ with the spatial size $h \times w$, we first apply two different augmentations to obtain q_t and q_s , and randomly sample a set of augmented images $\{x_j\}_{j=1}^n$ from a set of unlabeled image slices $\{x_i\}_{i=N+1}^{N+M}$. We believe that such relaxation enables the model to capture more rich semantic relationships and anatomical features from its neighboring images instead of only learning from the different version of the same query image. We then feed $\{x_j\}_{j=1}^n$ to the teacher encoder E_t , and followed by the nonlinear projection head H_t^g to generate their projection embeddings $\{H_t^g(E_t(x_j))\}_{j=1}^n$ as anchor points, and also feed q_t and q_s to the teacher and student (*i.e.*, E and H), creating $z_t = H_t^g(E_t(q_t))$ and $z_s = H_s^g(E_s(q_s))$. Here we utilize the probabilities after SoftMax instead of the feature embedding:

$$p_t(j) = -\log \frac{\exp(\text{sim}(z_t, a_j)/\tau_t)}{\sum_{i=1}^n \exp(\text{sim}(z_t, a_i)/\tau_t)}, \quad (1)$$

where τ_t is a temperature hyperparameter of the teacher, and $\text{sim}(\cdot, \cdot)$ is the cosine similarity. Then inspired by [7], in order to avoid collapsed solutions in an unsupervised scenario, we use a shallow multi-layer perceptron (MLP) predictor $H_p^g(\cdot)$ to obtain the prediction $z_s^* = H_p^g(z_s)$. Of note, $\{a_i\}_{i=1}^n$, z_t , z_s , z_s^* can be generated embedding from a set of randomly chosen augmented images, teacher’s projection embeddings, student’s projection embeddings, and student’s prediction embeddings in either Stage- i or ii . Therefore, we can calculate the similarity distance between the student’s prediction and the anchor embeddings by converting them to probability distribution.

$$p_s(j) = -\log \frac{\exp(\text{sim}(z_s^*, a_j)/\tau_s)}{\sum_{i=1}^n \exp(\text{sim}(z_s^*, a_i)/\tau_s)}, \quad (2)$$

where τ_s refers to a temperature hyperparameter of the student. The unsupervised contrastive loss is computed as follows:

$$\mathcal{L}_{\text{contrast}} = \text{KL}(p_t || p_s). \quad (3)$$

Local Contrastive Distillation Pre-training. After training the teacher’s and student’s encoder to learn global-level image features, we attach the decoders and tune the entire models to perform pixel-level contrastive learning in a semi-supervised manner. The **distinction in the training strategy** between ours and [9] lies in Stage-*ii* and *iii*: [9] only use labeled data in training, while we use both labeled and unlabeled data in training. Considering the training procedure of Stage-*ii* is similar to Stage-*iii*, we briefly describe it here as illustrated in Fig. 2. For the labeled data, we train our model by minimizing the supervised loss (the linear combination of cross-entropy loss and dice loss) in Stage-*ii* and Stage-*iii*. As for the unlabeled input images q and $\{x_j\}_{j=1}^n$, we first apply two different augmentations to q , creating two different versions $[q_t^l, q_s^l]$, and then feed them to F_t and F_s , and their output features $[f_t, f_s]$ are fed into H_t^l and H_s^l . The student’s projection embedding is subsequently fed into H_p^l to obtain the student’s prediction embedding to enforce the similarity between the teacher and the student under the same loss as Eq. 3. We also include the randomly selected images to enforce such similarity because intuitively, it may be beneficial to ensure **diversity** in the set of sampled examples. It is important to note that ACTION will re-use the well-trained weight of the models F_t and F_s as initialization for Stage-*iii*.

Anatomical Contrast Fine-Tuning. Broadly speaking, in medical images, the same tissue types may share similar anatomical information in different patients, but different tissue types often show different class, appearance, and spatial distributions, which can be described as a complicated form of **imbalance** and **uncertainty** in real clinical data, as shown in Fig. 1. This motivates us to efficiently incorporate more useful features so the representations can be more balanced and better discriminated in such multi-class label imbalanced scenarios. Inspired by [15], we propose AnCo, a new unsupervised contrastive loss designed at the anatomical level. Specifically, we additionally attach a representation decoder head H_r to the student network, parallel to the segmentation head, to decode the multi-layer hidden features by first using multiple up-sampling layers for outputting dense features with the same spatial resolution as the query image and then mapping them into high m -dimensional query, positive key, and negative key embeddings: r_q, r_k^+, r_k^- . The AnCo loss is then defined as:

$$\mathcal{L}_{\text{anco}} = \sum_{c \in \mathcal{C}} \sum_{r_q \sim \mathcal{R}_q^c} -\log \frac{\exp(r_q \cdot r_k^{c,+} / \tau_{an})}{\exp(r_q \cdot r_k^{c,+} / \tau_{an}) + \sum_{r_k^- \sim \mathcal{R}_k^c} \exp(r_q \cdot r_k^- / \tau_{an})}, \quad (4)$$

where \mathcal{C} is a set of all available classes in a mini-batch, and τ_{an} denotes a temperature hyperparameter for AnCo loss. \mathcal{R}_q^c and $r_k^{c,+}$ are a set of query embeddings in class c and the positive key embedding, which is the mean representation of

class c , respectively. \mathcal{R}_k^c is a set of negative key embeddings which are not in class c . Suppose \mathcal{P} is a set including all pixel coordinates with the same resolution with x_i , these queries and keys are then defined as:

$$\mathcal{R}_q^c = \bigcup_{[m,n] \in \mathcal{P}} \mathbb{1}(y_{[m,n]} = c) r_{[m,n]}, \quad \mathcal{R}_k^c = \bigcup_{[m,n] \in \mathcal{P}} \mathbb{1}(y_{[m,n]} \neq c) r_{[m,n]}, \quad r_k^{c,+} = \frac{1}{|\mathcal{R}_q^c|} \sum_{r_q \in \mathcal{R}_q^c} r_q. \quad (5)$$

In addition, we note that contrastive learning usually benefits from a large collection of positive and negative pairs, but it is usually bounded by the size of GPU memory. Therefore, we introduce two novel active hard sampling methods. To address the *uncertainty* on the most challenging pixels among all available classes (*i.e.*, close anatomical or semantic relationship), we non-uniformly sample negative keys based on relative similarity distance between the query class and each negative key class. For each mini-batch, we build a graph G to measure the pair-wise class relationship to dynamically update G .

$$G[p, q] = (r_k^{p,+} \cdot r_k^{q,+}), \quad \forall p, q \in \mathcal{C}, \text{ and } p \neq q, \quad (6)$$

where $G \in \mathbb{R}^{|\mathcal{C}| \times |\mathcal{C}|}$. Note that this process may be hard to allocate more samples. Thus, to learn a more accurate decision boundary, we first apply SoftMax function by normalizing the pair-wise relationships among all negative classes n from each query class c , yielding a distribution: $\exp(G[c, v]) / \sum_{n \in \mathcal{C}, n \neq c} \exp(G[c, n])$. Then we adaptively sample negative keys from each class v to help learn the corresponding query class c . To alleviate the *imbalance* issue, we sample hard queries based on a defined threshold, to better discriminate the rare classes. The easy and hard queries are computed as follows:

$$\mathcal{R}_q^{c, \text{easy}} = \bigcup_{r_q \in \mathcal{R}_q^c} \mathbb{1}(\hat{y}_q > \theta_s) r_q, \quad \mathcal{R}_q^{c, \text{hard}} = \bigcup_{r_q \in \mathcal{R}_q^c} \mathbb{1}(\hat{y}_q \leq \theta_s) r_q, \quad (7)$$

where \hat{y}_q is the predicted confidence of label c corresponding to r_q after SoftMax function, and θ_s is the user-defined confidence threshold.

3 Experiments


Experimental Setup. We experiment on two benchmark datasets: ACDC 2017 dataset [1] and MICCAI 2017 Liver Tumor Segmentation Challenge (LiTS) [2].

The ACDC dataset includes 200 cardiac cine MRI scans from 100 patients with annotations including three segmentation classes (*i.e.*, left ventricle (LV), myocardium (Myo), and right ventricle (RV)). Following [16, 27], we use 140, 20, and 60 scans for training, validation, and testing, respectively.

The LiTS dataset includes 131 contrast-enhanced 3D abdominal CT volumes with annotations of two segmentation classes (*i.e.*, liver and tumor). Following [13], we use the first 100 volumes for training, and the rest 31 for testing. For pre-processing, we follow the setting in [3] to normalize the intensity of each 3D scans,

Table 1. Comparison of segmentation performance (DSC[%]/ASD[voxel]) on ACDC under two unlabeled settings (3 or 7 labeled). The best results are indicated in **bold**.

Method	Average	3 Labeled			Average	7 Labeled		
		RV	Myo	LV		RV	Myo	LV
UNet-F [22]	91.5/0.996	90.5/0.606	88.8/0.941	94.4/1.44	91.5/0.996	90.5/0.606	88.8/0.941	94.4/1.44
UNet-L	51.7/13.1	36.9/30.1	54.9/4.27	63.4/5.11	79.5/2.73	65.9/0.892	82.9/2.70	89.6/4.60
EM [26]	59.8/5.64	44.2/11.1	63.2/3.23	71.9/2.57	75.7/2.73	68.0/0.892	76.5/2.70	82.7/4.60
CCT [18]	59.1/10.1	44.6/19.8	63.2/6.04	69.4/4.32	75.9/3.60	67.2/2.90	77.5/3.32	82.9/0.734
DAN [37]	56.4/15.1	47.1/21.7	58.1/11.6	63.9/11.9	76.5/3.01	75.7/2.61	73.3/3.11	80.5/3.31
URPC [17]	58.9/8.14	50.1/12.6	60.8/4.10	65.8/7.71	73.2/2.68	67.0/0.742	72.2/0.505	80.4/6.79
DCT [21]	58.5/10.8	41.2/21.4	63.9/5.01	70.5/6.05	78.1/2.64	70.7/1.75	77.7/2.90	85.8/3.26
ICT [25]	59.0/6.59	48.8/11.4	61.4/4.59	66.6/3.82	80.6/1.64	75.1/0.898	80.2/1.53	86.6/2.48
MT [23]	58.3/11.2	39.0/21.5	58.7/7.47	77.3/4.72	80.1/2.33	75.2/1.22	79.2/2.32	86.0/3.45
UAMT [36]	61.0/7.03	47.8/15.9	65.0/2.38	70.1/2.83	77.6/3.15	70.5/0.81	78.4/4.36	83.9/4.29
CPS [6]	61.0/2.92	43.8/2.95	64.5/2.84	74.8/2.95	78.8/3.41	74.0/1.95	78.1/3.11	84.5/5.18
GCL [3]	70.6/2.24	56.5/1.99	70.7/1.67	84.8/3.05	87.0/0.751	86.9/0.584	81.8/0.821	92.5/0.849
SCS [9]	73.6/5.37	63.5/6.23	76.6/2.42	80.7/7.45	84.2/2.01	81.4/0.850	83.0/2.03	88.2/3.12
• ACTION (ours)	87.5/1.12	85.4/0.915	85.8/0.784	91.2/1.66	89.7/0.736	89.8/0.589	86.7/0.813	92.7/0.804

**Fig. 3.** Visualization of segmentation results on ACDC with 3 labeled data. As is shown, ACTION consistently produces sharper object boundaries and more accurate predictions across all methods. Different structure categories are shown in different colors. (Color figure online)

resample all 2D slices and the corresponding segmentation maps to a fixed spatial resolution (*i.e.*, 256×256 pixels). To quantitatively assess the performance of our proposed method, we report two popular metrics: Dice coefficient (DSC) and Average Surface Distance (ASD) for 3D segmentation results.

Implementation Details. All our models are implemented in PyTorch [19]. We train all methods with SGD optimizer (learning rate = 0.01, momentum = 0.9, weight decay = 0.0001, batch size = 6). All models are trained with two NVIDIA GeForce RTX 3090 GPUs. Stage-*i* and *ii* are trained with 100 epochs, and Stage-*iii* is with 200 epochs. We use the temperature of teacher and student as $\tau_t = 0.01$ and $\tau_s = 0.1$. The teacher is updated using the following

Table 2. Comparison of segmentation performance (DSC [%]/ASD [voxel]) on LiTS under two unlabeled settings (5% or 10% labeled ratio). The best results are in **bold**.

Method	Average	5% Labeled		10% Labeled		
		Liver	Tumor	Average	Liver	Tumor
UNet-F [22]	68.2/16.9	90.6/8.14	45.8/25.6	68.2/16.9	90.6/8.14	45.8/25.6
UNet-L	60.4/30.4	87.5/9.84	33.3/50.9	61.6/28.3	85.4/18.6	37.9/37.9
EM [26]	61.2/33.3	87.7/9.47	34.7/57.1	62.9/38.5	87.4/21.3	38.3/55.7
CCT [18]	60.6/48.7	85.5/27.9	35.6/69.4	63.8/31.2	90.3/7.25	37.2/55.1
DAN [37]	62.3/25.8	88.6/9.64	36.1/42.1	63.2/30.7	87.3/15.4	39.1/46.1
URPC [17]	62.4/37.8	86.7/21.6	38.0/54.0	63.0/43.1	88.1/24.3	38.9/61.9
DCT [21]	60.8/34.4	89.2/12.6	32.5/56.2	61.9/31.7	86.2/19.3	37.5/44.1
ICT [25]	60.1/39.1	86.8/12.6	33.3/65.6	62.5/32.4	88.1/16.7	36.9/48.2
MT [23]	61.9/40.0	86.7/21.6	37.2/58.4	63.3/26.2	89.7/11.6	36.9/40.8
UAMT [36]	61.0/47.0	86.9/22.1	35.2/71.8	62.3/26.0	87.4/7.55	37.3/44.4
CPS [6]	62.1/36.0	87.3/17.9	36.8/54.0	64.0/23.6	90.2/10.6	37.8/36.7
GCL [3]	63.3/20.1	90.7/9.46	35.9/30.8	65.0/37.2	91.3/10.0	38.7/64.3
SCS [9]	61.5/28.8	92.6/7.21	30.4/50.3	64.6/33.9	91.6/5.72	37.6/62.0
• ACTION (ours)	66.8/17.7	93.0/6.04	40.5/29.4	67.7/20.4	92.8/5.08	42.6/35.8




Fig. 4. Visualization of segmentation results on LiTS with 5% labeled ratio. As is shown, ACTION achieves consistently sharp and accurate object boundaries compared to other SSL methods. Different structure categories are shown in different colors. (Color figure online)

rule $\theta_t \leftarrow m\theta_t + (1 - m)\theta_s$, where θ refers to the model’s parameters and the momentum hyperparameter m is 0.99. The memory bank size is 36. We follow the standard augmentation strategies in [7]. In Stage- i , we train E_s, E_t, H_t^g, H_s^g , and H_p^g on the unlabeled data with global-level $\mathcal{L}_{\text{contrast}}$ in Eq. 3. We follow [9] to use a MLP as heads, and the setting of the predictors is similar to [7], which has a feature dimension of 512. In Stage- ii , we train F_s, F_t, H_t^l, H_s^l , and H_p^l on the labeled and unlabeled data. We train with the supervised loss [36] on labeled data, and local-level $\mathcal{L}_{\text{contrast}}$ in Eq. 3 on unlabeled data. Given the logits output $\hat{y} \in \mathbb{R}^{C \times h \times w}$, we use the 1×1 convolutional layer to project all

pixels into the latent space with the feature dimension of 512, and the output feature dimension of G is also 512. As for Stage-*iii*, we train F_s , F_t , H_t , H_s , and H_r on the labeled and unlabeled data. We use the supervised segmentation loss on labeled data, unsupervised cross-entropy loss (on pseudo-labels generated by a confidence threshold θ_s), and $\mathcal{L}_{\text{ancco}}$ in Eq. 4 on unlabeled data. We then adaptively sample 256 query samples and 512 key samples for each mini-batch, and temperature for the student and confidence thresholds are set to $\tau_s = 0.5$ and $\theta_s = 0.97$, respectively. Of note, the projection heads, the predictor, and the representation decoder head are only utilized during the training, and will be removed during the inference.

Main Results. We compare our proposed method to previous state-of-the-art SSL methods using 2D Unet [22] as backbone, including UNet trained with full/limited supervisions (UNet-F/UNet-L), EM [26], CCT [18], DAN [37], URPC [17], DCT [21], ICT [25], MT [23], UAMT [36], CPS [6], SCS [9], and GCL [3]. Table 1 shows the evaluation results on ACDC dataset under two unlabeled settings (3 or 7 labeled cases). ACTION can substantially improve results on two unlabeled settings, greatly outperforming the previous state-of-the-art SSL methods. Specifically, our ACTION, trained on 3 labeled cases, dramatically improves the previous best averaged Dice score from 73.6% to 87.5% by a large margin, and even matches previous SSL methods using 7 labeled cases. When using 7 labeled cases, ACTION further pushes the state-of-the-art results to 89.7% in Dice. We observe that the gains are more pronounced on the two categories (*i.e.*, RV and Myo), and our ACTION achieves 89.8% and 86.7% in terms of Dice, performing competitive or even better than the supervised baseline (89.2% and 86.7%). As shown in Fig. 3, we can see the clear advantage of ACTION, where the boundaries of different regions are clearly sharper and more accurate such as RV and Myo regions. Table 2 also shows the evaluation results on LiTS dataset under two unlabeled settings (5% or 10% labeled cases). On both two labeled settings, ACTION significantly outperforms all the state-of-the-art methods by a significant margin. As shown in Fig. 4, ACTION achieves consistently sharp and accurate object boundaries compared to other SSL methods.

Ablation on Different Components. We investigate the impact of different components in ACTION. All reported results in this section are based on the ACDC dataset under the 3 labeled setting. Table 3 shows the ablation result of our model. Upon our choice of architecture, we first consider a naïve baseline (BYOL) without any random sampled images (RSI), stage-*ii*, and stage-*iii*, denoted by (1) Vanilla. Then, we consider a wide range of different settings for improved representation learning: (2) incorporating other random sampled images; (3) no stage-*ii*; (4) no other random sampled images and stage-*ii*; (5) no stage-*iii*; since stage-*iii* includes two losses, (6) no $\mathcal{L}_{\text{ancco}}$, (7) no $\mathcal{L}_{\text{unsup}}$, and (8) our proposed ACTION. As shown in Table 3, it is notable that ACTION performs generally better than other evaluated baselines. We find that only applying any single component of ACTION often comes at the cost of performance degradation. The intuitions behind are as follows: (1) incorporating

Table 3. Ablation on (a) model component: w/o Random Sampled Images (RSI); w/o Local Contrastive Distillation (Stage-ii); w/o Anatomical Contrast Fine-tuning (Stage-iii); (b) loss formulation: w/o $\mathcal{L}_{\text{anco}}$; w/o $\mathcal{L}_{\text{unsup}}$; compared to the Vanilla and our proposed ACTION. Note that $\mathcal{L}_{\text{unsup}}$ denotes cross-entropy loss (on pseudo-labels generated by a confidence threshold θ_s) together with $\mathcal{L}_{\text{anco}}$ used in Stage-iii.

	Method	Metrics	
		Dice [%]	ASD [voxel]
(a)	Vanilla	60.6	6.64
	ACTION (ours)	87.5	1.12
	w/o RSI	82.7	6.66
	w/o Stage-ii	86.4	1.69
	w/o RSI + Stage-ii	82.6	1.77
(b)	w/o Stage-iii	76.7	2.91
	w/o $\mathcal{L}_{\text{anco}}$	86.5	1.30
	w/o $\mathcal{L}_{\text{unsup}}$	83.7	2.51

Table 4. Ablation on augmentation strategies.

Method	Student Teacher		Metrics	
	Aug.	Aug.	Dice[%]	ASD[voxel]
ACTION	Weak	Weak	84.6	1.78
ACTION	Strong	Weak	87.5	1.12
ACTION	Weak	Strong	85.4	2.12
ACTION	Strong	Strong	86.5	1.89

other random sampled images will enforce the diversity of the sampled data, preventing redundant anatomically and semantically similar samples; (2) using stage-ii leads to worse performance without considering local context; (3) using stage-iii enables a robust segmentation model to learn better representations with few human annotations. Using the above components confers a significant advantage at representation learning, and further illustrates the benefit of each component.

Ablation on Different Augmentations. We investigate the impact of using weak or strong augmentations for ACTION on the ACDC dataset under 3 labeled setting. We summarize the effects of different data augmentation strategies in Table 4. We apply *weak* augmentation to the teacher’s input, including rotation, cropping, flipping, and *strong* augmentation to the student’s input, including rotation, cropping, flipping, random contrast, and brightness changes [20]. Empirically, we find that when using weak and strong augmentation strategies on the teacher and student network, the network performance is optimal.

4 Conclusion and Limitations

In this work, we have presented ACTION, a novel anatomical-aware contrastive distillation framework with active sampling, designed specifically for medical image segmentation. Our method is motivated by two observations that all negative samples are not equally negative, and the underlying class distribution of medical images is highly unlabeled and imbalanced. Through extensive experiments across two benchmark datasets and unlabeled settings, we show that ACTION can significantly improve segmentation performance with minimal additional memory requirements, outperforming the previous state-of-the-art by a large margin. For future work, we plan to explore a more advanced contrastive learning approach for better performance when the medical data is unlabeled and imbalanced.

References

1. Bernard, O., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE Trans. Med. Imaging* **37**, 2514–2525 (2018)
2. Bilic, P., et al.: The liver tumor segmentation benchmark (LiTS). *arXiv preprint arXiv:1901.04056* (2019)
3. Chaitanya, K., Erdil, E., Karani, N., Konukoglu, E.: Contrastive learning of global and local features for medical image segmentation with limited annotations. In: *NeurIPS* (2020)
4. Chen, S., Bortsova, G., García-Uceda Juárez, A., van Tulder, G., de Bruijne, M.: Multi-task attention-based semi-supervised learning for medical image segmentation. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11766, pp. 457–465. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_51
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *ICML*, pp. 1597–1607. PMLR (2020)
6. Chen, X., Yuan, Y., Zeng, G., Wang, J.: Semi-supervised semantic segmentation with cross pseudo supervision. In: *CVPR* (2021)
7. Grill, J.B., et al.: Bootstrap your own latent—a new approach to self-supervised learning. In: *NeurIPS* (2020)
8. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *CVPR*, pp. 9729–9738 (2020)
9. Hu, X., Zeng, D., Xu, X., Shi, Y.: Semi-supervised contrastive learning for label-efficient medical image segmentation. In: de Bruijne, M., et al. (eds.) *MICCAI 2021*. LNCS, vol. 12902, pp. 481–490. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87196-3_45
10. Huynh, T., Kornblith, S., Walter, M.R., Maire, M., Khademi, M.: Boosting contrastive self-supervised learning with false negative cancellation. In: *WACV* (2022)
11. Kervadec, H., Dolz, J., Granger, É., Ben Ayed, I.: Curriculum semi-supervised segmentation. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 568–576. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_63
12. Li, S., Zhang, C., He, X.: Shape-aware semi-supervised 3D semantic segmentation for medical images. In: Martel, A.L., et al. (eds.) *MICCAI 2020*. LNCS, vol. 12261, pp. 552–561. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_54

13. Li, X., Chen, H., Qi, X., Dou, Q., Fu, C.W., Heng, P.A.: H-DenseUNet: hybrid densely connected UNet for liver and tumor segmentation from CT volumes. *IEEE Trans. Med. Imaging* **37**, 2663–2674 (2018)
14. Li, Z., Kamnitsas, K., Glocker, B.: Analyzing overfitting under class imbalance in neural networks for image segmentation. *IEEE Trans. Med. Imaging* **40**, 1065–1077 (2020)
15. Liu, S., Zhi, S., Johns, E., Davison, A.J.: Bootstrapping semantic segmentation with regional contrast. *arXiv preprint [arXiv:2104.04465](https://arxiv.org/abs/2104.04465)* (2021)
16. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: *AAAI* (2020)
17. Luo, X., et al.: efficient semi-supervised gross target volume of nasopharyngeal carcinoma segmentation via uncertainty rectified pyramid consistency. In: de Bruijne, M. (ed.) *MICCAI 2021. LNCS*, vol. 12902, pp. 318–329. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87196-3_30
18. Ouali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: *CVPR* (2020)
19. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: *NeurIPS* (2019)
20. Perez, F., Vasconcelos, C., Avila, S., Valle, E.: Data augmentation for skin lesion analysis. In: Stoyanov, D., et al. (eds.) *CARE/CLIP/OR 2.0/ISIC -2018. LNCS*, vol. 11041, pp. 303–311. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01201-4_33
21. Qiao, S., Shen, W., Zhang, Z., Wang, B., Yuille, A.: Deep co-training for semi-supervised image recognition. In: *ECCV* (2018)
22. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015. LNCS*, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
23. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: *NeurIPS*, pp. 1195–1204 (2017)
24. Tejankar, A., Koohpayegani, S.A., Pillai, V., Favaro, P., Pirsiavash, H.: ISD: self-supervised learning by iterative similarity distillation. In: *ICCV* (2021)
25. Verma, V., Kawaguchi, K., Lamb, A., Kannala, J., Bengio, Y., Lopez-Paz, D.: Interpolation consistency training for semi-supervised learning. In: *IJCAI* (2019)
26. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: ADVENT: adversarial entropy minimization for domain adaptation in semantic segmentation. In: *CVPR*, pp. 2517–2526 (2019)
27. Wu, Y., et al.: Mutual consistency learning for semi-supervised medical image segmentation. *Med. Image Anal.* **81**, 102530 (2022)
28. Yang, L., et al.: NuSeT: a deep learning tool for reliably separating and analyzing crowded cells. *PLoS Comput. Biol.* **16**, e1008193 (2020)
29. You, C., et al.: Mine your own anatomy: revisiting medical image segmentation with extremely limited labels. *arXiv preprint [arXiv:2209.13476](https://arxiv.org/abs/2209.13476)* (2022)
30. You, C., et al.: Rethinking semi-supervised medical image segmentation: a variance-reduction perspective. *arXiv preprint [arXiv:2302.01735](https://arxiv.org/abs/2302.01735)* (2023)
31. You, C., et al.: Incremental learning meets transfer learning: application to multi-site prostate MRI segmentation. In: Albarqouni, S., et al. (eds.) *DeCaF FAIR 2022. LNCS*, vol. 13573, pp. 3–16. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-18523-6_1

32. You, C., Yang, J., Chapiro, J., Duncan, J.S.: Unsupervised wasserstein distance guided domain adaptation for 3D multi-domain liver segmentation. In: Cardoso, J., et al. (eds.) IMIMIC/MIL3ID/LABELS -2020. LNCS, vol. 12446, pp. 155–163. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-61166-8_17
33. You, C., et al.: Class-aware adversarial transformers for medical image segmentation. In: NeurIPS (2022)
34. You, C., Zhao, R., Staib, L.H., Duncan, J.S.: Momentum contrastive voxel-wise representation learning for semi-supervised volumetric medical image segmentation. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. LNCS, vol. 13434, pp. 639–652. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16440-8_61
35. You, C., Zhou, Y., Zhao, R., Staib, L., Duncan, J.S.: SimCVD: simple contrastive voxel-wise representation distillation for semi-supervised medical image segmentation. *IEEE Trans. Med. Imaging* **41**, 2228–2237 (2022)
36. Yu, L., Wang, S., Li, X., Fu, C.-W., Heng, P.-A.: Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 605–613. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_67
37. Zhang, Y., Yang, L., Chen, J., Fredericksen, M., Hughes, D.P., Chen, D.Z.: Deep adversarial networks for biomedical image segmentation utilizing unannotated images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 408–416. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_47
38. Zhou, Y., et al.: Semi-supervised 3d abdominal multi-organ segmentation via deep multi-planar co-training. In: WACV. IEEE (2019)
39. Zhu, X., Anguelov, D., Ramanan, D.: Capturing long-tail distributions of object subcategories. In: CVPR (2014)