

IDEAL: IMPROVED DENSE LOCAL CONTRASTIVE LEARNING FOR SEMI-SUPERVISED MEDICAL IMAGE SEGMENTATION

Hritam Basak¹, Soumitri Chattopadhyay^{*2}, Rohit Kundu^{*2}, Sayan Nag^{*3}, Rammohan Mallipeddi⁴

¹Stony Brook University, ²Jadavpur University, ³University of Toronto, ⁴Kyungpook National University

ABSTRACT

Due to the scarcity of labeled data, Contrastive Self-Supervised Learning (SSL) frameworks have lately shown great potential in several medical image analysis tasks. However, the existing contrastive mechanisms are sub-optimal for dense pixel-level segmentation tasks due to their inability to mine local features. To this end, we extend the concept of metric learning to the segmentation task, using a dense (dis)similarity learning for pre-training a deep encoder network, and employing a semi-supervised paradigm to fine-tune for the downstream task. Specifically, we propose a simple convolutional projection head for obtaining dense pixel-level features, and a new contrastive loss to utilize these dense projections thereby improving the local representations. A bidirectional consistency regularization mechanism involving two-stream model training is devised for the downstream task. Upon comparison, our IDEAL method outperforms the SoTA methods by fair margins on cardiac MRI segmentation. Our source codes are publicly accessible at: <https://github.com/Rohit-Kundu/IDEAL-ICASSP23>.

Index Terms— Semi-supervised learning, Segmentation, MRI, Contrastive learning

1. INTRODUCTION

The success of supervised deep learning approaches can be attributed to the availability of large quantities of labeled data that is essential for network training [1]. However, in the biomedical domain [2, 3], it is difficult to acquire such large quantities of annotated data as the annotations are to be done by trained medical professionals. Although supervised learning has been used extensively in the past decade in biomedical imaging [1], Self-supervised Learning (SSL) [4, 5, 6, 7, 8] provides more traction for sustaining deep learning methods in the medical vision domain [9, 10, 11]. SSL-based pre-training alleviates the data annotation problem by utilizing only unstructured data to learn distinctive information which can further be utilized in downstream applications, typically in a semi-supervised fashion [12, 13, 9]. Typically, SSL have shown to be exceedingly promising in domains with enormous amounts of data such as natural images [6, 7] with a

focus on both contrastive [14, 7, 6, 15] and non-contrastive variants [16, 17]. Recently, SSL has begun to be employed for medical imaging as well, evident from the survey [2].

One of the most popular ways of employing SSL is contrastive learning [14, 7, 6], which owing to its prolific performance across various vision tasks has almost become a de facto standard for self-supervision. The intuition of contrastive learning is to *pull* embeddings of semantically similar objects closer and simultaneously *push* away the representations of dissimilar objects. In medical image segmentation, contrastive learning has been leveraged in a few prior works [13, 18]. However, naive contrastive learning frameworks such as SimCLR [6] and MoCo [7] learn a global representation and thus, are not suitable to be applied directly to segmentation, a pixel-level task. The seminal work by Chaitanya *et. al.* [13] utilized a contrastive pre-training strategy to learn useful representations for a segmentation task in a low-data regime. Despite having success, their method has a few limitations, which we attempt to address and improve upon in this paper. In particular, we take a cue from the aforementioned work [13] to propose a novel contrastive learning-based medical image segmentation framework.

Although our work is built upon [13], there are several salient points of difference between the two models that render our contributions non-trivial and unique. First of all, our proposed method preserves spatial information and constructs a dense output projection that preserves locality. This is in contrast to the projection used in [13] where a global pooling is applied to the projection used in [13] where a global pooling is applied to the backbone encoder, thereby obtaining a single global feature representation vector for every input image. In other words, [13] employed global contrastive learning for encoder pre-training and the local contrastive learning is only restricted to the decoder fine-tuning. In contrast, our method involves pre-training of the encoder on dense local features. The intuitive difference between our proposed framework and that of [13] has been depicted in Fig. 1. We argue that this will benefit pixel-level informed downstream tasks such as segmentation (supported by our findings in Section 3).

Secondly, our definition of finding positive and negative pairs for contrastive learning is different from that of traditional contrastive [6, 7] methods. This is because, in our case, we find dense correspondences across views and define positives and negatives accordingly, i.e., we perform feature-

^{*}Equal contribution.

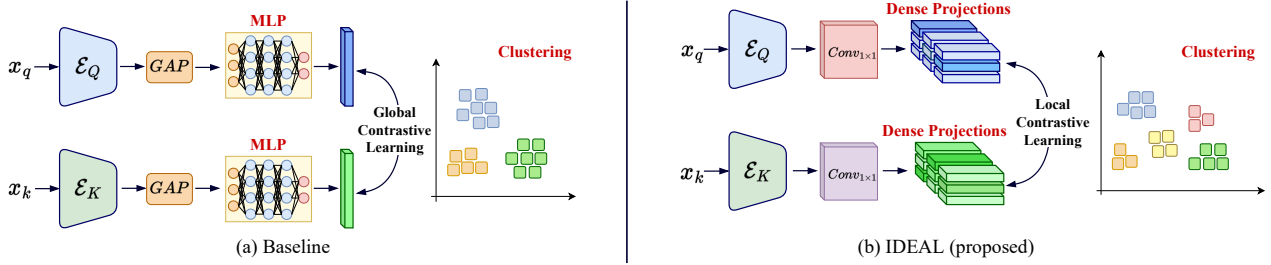


Fig. 1. Schematic representation of contrastive learning-based pre-training of (a) baseline model [13] (b) our IDEAL model. \mathcal{E} represents the encoder backbone; GAP = Global Average Pooling; x_q and x_k are the query and key images respectively. The convolution projection head in the IDEAL model generates dense projections that give better local clusters than using a traditional global pooling followed by a Multi-Layer Perceptron (MLP).

level Contrastive Learning (refer to Section 2). Further, we adopt a novel bidirectional consistency regularization mechanism that perturbs the network instead of image perturbation used by most frameworks [7, 19, 9]. By this approach, we enforce both streams of data flow to learn better features competitively, unlike the traditional image augmentation-based consistency regularization approaches where only one model stream learns through a single backpropagation. The empirical evaluation suggests that our proposed consistency regularization yields a superior segmentation performance, outperforming existing state-of-the-art semi-supervised approaches in literature [20, 21]. Our contributions are as follows:

1. We propose an SSL strategy that leverages dense projection head representations for contrastive learning for learning robust local features.
2. We redefine 'positive' and 'negative' samples in contrastive learning and extend the InfoNCE loss to adapt to dense representations during the pre-training phase.
3. Our work also employs a uniquely devised cross-consistency regularization for fine-tuning the network to the downstream segmentation task.

2. METHODOLOGY

Our proposed method consists of two parts: (1) Self-Supervised pre-training of the encoder, and (2) Semi-Supervised fine-tuning of the network using consistency regularization. Taking reference from [13], we propose a pairwise contrastive (dis)similarity learning, employing global and local feature exploration for pre-training the encoder. The pre-trained encoder is thereafter transferred for the downstream task of cross-supervision-based consistency regularization, which aims to fine-tune the network.

2.1. Background: Global Contrastive Learning

The widely employed variant of self-supervised representation learning is contrastive learning [7, 6, 14], where the

images and their perturbations are fed into a cascade of the shared encoder and projection head. The obtained features (global representations) of each of these images in the mini-batches are then utilized for a contrastive loss function, which is designed considering contrastive learning as a dictionary lookup task [15]. Out of all the encoded keys in the set k for each encoded query q , a single positive key k^+ is matched, whilst the remainder of the keys (negative keys) represent different images of the mini-batch. A contrastive loss function is represented as follows:

$$\mathcal{L}_g = -\log \frac{\exp(q \cdot k^+ / \tau)}{\exp(q \cdot k^+ / \tau) + \sum_{k^-} \exp(q \cdot k^- / \tau)} \quad (1)$$

2.2. Improved Local Contrastive Learning

We employ local contrastive Learning that extends the existing global contrastive framework into a dense paradigm. The key difference between the traditional contrastive framework and the current work lies in the formulation of the loss function and the projection head. The traditional contrastive learning follows the paradigm where features are extracted by a backbone encoder (e.g. ResNet [22]), followed by a projection head. The projection head typically consists of a global pooling operation and an MLP module comprising a few FC layers with ReLU layers in between them. However, in our case, we remove the entire global pooling and MLP part and replace it with a 1×1 convolution operation (\mathcal{G}). This projection head generates dense feature representations, with an equal number of parameters as the traditional ones.

We define a set of keys $\{k_0, k_1, \dots\}$ extracted from the encoder for every query q , where the encoded keys represent the local part of the image, which is very different from traditional contrastive learning where every key corresponds to a single image view. These keys represent individual features of the $D_h \times D_w$ vectors extracted from the encoder followed by 1×1 convolution, where D_h and D_w represent the spatial dimension of the extracted feature maps. The positive keys k^+ are pooled from the extracted correspondence across views, i.e. one of the $D_h \times D_w$ features extracted from different

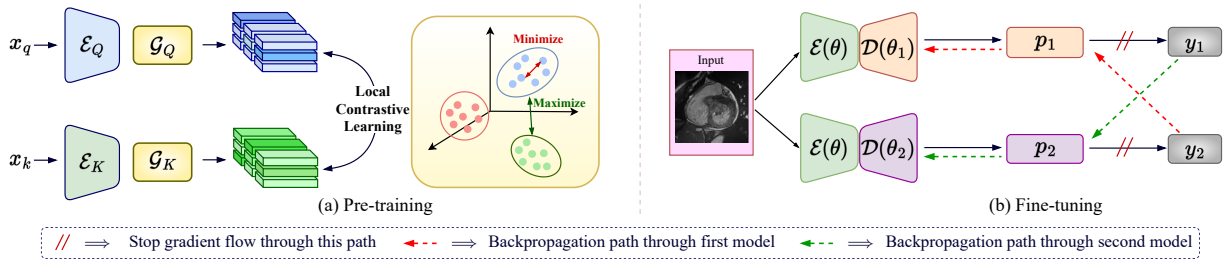


Fig. 2. An overview of the proposed IDEAL framework: (a) *Pre-training*- x_q and x_k are the query and key images, \mathcal{E} and \mathcal{G} represent the encoder and projection head, respectively. The projection head employs a 1×1 convolution layer instead of a traditional MLP for dense feature extraction, resulting in better local clustering of features. (b) *Fine-tuning*- Two perturbed branches with the same input are employed. $\mathcal{E}(\theta)$ is the shared encoder initialized similarly for both streams, $\mathcal{D}(\theta_1)$ and $\mathcal{D}(\theta_2)$ represent two different decoder architectures; p_1 and p_2 are the predicted output segmentation maps which are thresholded to obtain y_1 and y_2 respectively. y_1 backpropagates through the second stream and y_2 backpropagates through the first stream enforcing cross-consistency in segmentation.

views of the same image. The negative key k^- is assigned from features from a view of different images. Having this assignment strategy and taking reference from InfoNCE [14], we define the local contrastive loss as:

$$\mathcal{L}_{loc} = \frac{-1}{D_h D_w} \sum_{i=1}^{D_h D_w} \log \frac{\exp(q_i \cdot k_i^+ / \tau)}{\exp(q_i \cdot k_i^+ / \tau) + \sum_{k_i^-} \exp(q_i \cdot k_i^- / \tau)} \quad (2)$$

2.3. Semi-Supervised Fine-tuning

The SSL pre-training phase is followed by a semi-supervised fine-tuning of the trained encoder network for our downstream segmentation task. Consistency regularization [9, 13], a very popular semi-supervised approach for segmentation, has been employed here, however, with two very unique and non-trivial customization, which are discussed below.

In typical consistency regularization schemes, such as in MoCo [7], dampened sharing of weights takes place i.e. one branch is learned directly following the gradient updates whereas the other branch is learned by employing an Exponential Moving Average (EMA) weight transfer strategy. Such a learning paradigm limits the feature learning ability of the second branch. To alleviate this, we consider two decoders of different architecture initialization. Furthermore, the up-scaling convolutions corresponding to each of them are made to work differently, i.e., one of them has ConvTranspose (convolution with trainable kernels) whilst the other one has an Up-sampling layer (simple bilinear interpolation). Outputs of each of these decoders are passed through sigmoid activation functions resulting in p_i which are then thresholded to generate one-hot y_i (see Fig. 2). For labeled samples, we compute a simple cross-entropy loss between target mask m and the predicted outputs p_i . Otherwise, a cross-consistency

loss is enforced between the output probabilities p_i from the two decoder branches, guided by the one-hot pseudo-label y_j coming from the other branch. In other words, y_1 acts as a pseudo-supervisory signal for p_2 , and y_2 for p_1 . Thus, one branch learns based on the other branch's output, leading to a *competitive learning* between two branches, resulting more confident predictions. The loss function is defined as:

$$\mathcal{L}_{SSF} = \begin{cases} CE(p_1, m) + CE(p_2, m), & \text{for labeled set} \\ CE(p_1, y_2) + CE(p_2, y_1), & \text{for unlabeled set} \end{cases} \quad (3)$$

3. EXPERIMENTS AND RESULTS

3.1. Experimental Setup

Datasets: The proposed method has been evaluated on two publicly available MRI datasets - (1) the ACDC dataset [23], consisting of 100 cardiac 3D shot-axis MRI volumes with expert manual annotations for three structures, namely, the left and right ventricles and the myocardium; and (2) the MMWHS dataset [24, 25], consisting of 20 cardiac 3D MRI volumes with expert manual annotations for seven structures: the left and right ventricles, the left and right atria, the pulmonary artery, the myocardium, and the ascending aorta. The datasets are split into train, validation, and test sets, where the validation set for each dataset consists of 2 volumes; the test set consists of 20 volumes for the ACDC dataset and 10 volumes for the MMWHS dataset. The rest of the volumes (78 for ACDC and 8 for MMWHS) are used for training. The training and validation set together form the pre-training set, while the testing set was used *solely* for model evaluation.

Implementation Details: IDEAL uses a ResNet-50 [22] encoder backbone with ADAM optimizer and an initial learning rate of $1e - 5$, and a U-Net [3] decoder with 4 upscaling

layers. The IDEAL model was implemented in PyTorch [26] and accelerated by an NVIDIA Tesla K80 GPU.

Evaluation Metrics: Three different metrics have been used to evaluate the IDEAL model on the segmentation tasks: Dice Similarity Coefficient (DSC), Average Symmetric Distance (ASD), and Hausdorff Distance (HD). The average scores over 5 runs for all the metrics have been reported for all the fine-tuning experiments.

3.2. Performance Analysis

3.2.1. Fine-tuning performances on limited annotations

Metric	ACDC				MMWHS			
	L=1.25%	L=2.5%	L=10%	L=100%	L=10%	L=20%	L=40%	L=100%
ASD ↓	0.677	0.614	0.556	0.489	2.397	1.798	1.355	1.317
HD ↓	2.409	2.209	2.094	1.999	5.001	3.499	2.221	2.183
DSC ↑	0.738	0.846	0.879	0.882	0.626	0.791	0.815	0.826

Table 1. Results obtained by the IDEAL framework with varying amounts of labeled data on the ACDC and MMWHS datasets. ‘ L ’ represents the amount of labeled data used.

To study the segmentation performance under limited labels during the fine-tuning phase, we have experimented with 1.25%, 2.5%, and 10% labeled volumes (denoted by ‘ L ’ henceforth) for ACDC and 10%, 20%, and 40% labeled volumes for the MMWHS dataset, following the existing works in literature that have used the same datasets [9, 13]. The amount of labeled data is expressed as a percentage of the training set. The results are tabulated in Table 1. From the results, it can be inferred that on both datasets, the expected trend of improving metrics with an increase in labeled data for the fine-tuning phase. Further, it is promising to observe that even with a fraction of labels available (10% for ACDC and 40% for MMWHS), the respective semi-supervised setups asymptotically approach the fully supervised setup (differs by 0.3% in ACDC and 1.1% in MMWHS), thus depicting the efficacy of our framework under limited annotations.

3.2.2. Comparison to state-of-the-art

We have also compared our proposed framework with several existing state-of-the-art semi-supervised cardiac MRI segmentation methods in literature [13, 6, 18, 10, 27, 21, 20, 28], the results of which are tabulated in Table 2. Additionally, we also provide a visual comparison of our segmentation performances with the SoTA methods in Fig. 3.

ACDC Dataset From Table 2, it can be observed that our IDEAL model performs the best in $L = 1.25\%$ setting, while being highly competitive to PCL [18] in $L = 2.5\%$ (~ 0.004) and $L = 10\%$ (~ 0.006). This shows that even in severely constrained conditions, IDEAL can perform optimally.

MMWHS Dataset Compared to existing state-of-the-art methods, our IDEAL model outperforms all previous frame-

Table 2. Performance Comparison (DSC scores) of the proposed IDEAL framework with SoTA methods in the literature on the ACDC and MMWHS datasets.

Method	Average DSC (ACDC)			Average DSC (MMWHS)		
	L=1.25%	L=2.5%	L=10%	L=10%	L=20%	L=40%
Chaitanya <i>et. al.</i> [13]	0.725	0.789	0.872	0.569	0.694	0.794
Global CL [6]	0.729	–	0.847	0.500	0.659	0.785
PCL [18]	0.671	0.850	0.885	–	–	–
Context Restoration [10]	0.625	0.714	0.851	0.482	0.654	0.783
MC-Net [28]	0.677	0.724	0.855	0.551	0.654	0.798
Label Efficient [27]	–	–	–	0.382	0.553	0.764
Data Augmentation [21]	0.731	0.786	0.865	0.529	0.661	0.785
Self Train [20]	0.690	0.749	0.860	0.563	0.691	0.801
Ours	0.738	0.846	0.879	0.626	0.791	0.815




Fig. 3. Visual comparison of our results with SoTA methods and ground truth, thus qualitatively validating the superiority of IDEAL in terms of segmentation performance.

works by large margins, as inferred from Table 2. Chaitanya *et. al.* [13] achieved the next best performance with DSC values of 56.9% ($L = 10\%$), 69.4% ($\sim 10\%$ less than ours) with 20% labeled data and 79.4% with 40% labeled data.

4. CONCLUSION

In this paper, we propose a semi-supervised approach for cardiac MRI segmentation to tackle the labeling bottleneck in medical imaging. Our framework leverages self-supervised contrastive learning with a novel projection head to capture dense local feature representation during pre-training, followed by employing a unique cross-consistency regularization scheme during model fine-tuning for the downstream segmentation task. Results show that our IDEAL framework outperforms several state-of-the-art methods on two widely used cardiac MRI datasets. We believe that such a paradigm can also be employed in cross-modal and cross-domain scenarios, as well as extended to natural images, something we intend to explore in near future.

5. ACKNOWLEDGEMENT

This work was supported by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2021R111A3049810).

6. REFERENCES

- [1] Hritam Basak, Rohit Kundu, and Ram Sarkar, "Mfs-net: A multi focus segmentation network for skin lesion segmentation," *Pattern Recognition*, 2022.
- [2] Saeed Shurrab and Rehab Duwairi, "Self-supervised learning methods and applications in medical imaging analysis: A survey," *PeerJ Computer Science*, 2022.
- [3] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*, 2015.
- [4] Carl Doersch, Abhinav Gupta, and Alexei A Efros, "Un-supervised visual representation learning by context prediction," in *ICCV*, 2015.
- [5] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros, "Context encoders: Feature learning by inpainting," in *CVPR*, 2016.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, 2020.
- [7] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick, "Momentum contrast for unsupervised visual representation learning," in *CVPR*, 2020.
- [8] Siladittya Manna, Soumitri Chattopadhyay, Saumik Bhattacharya, and Umapada Pal, "Swis: Self-supervised representation learning for writer independent offline signature verification," in *IEEE ICIP*, 2022.
- [9] Hritam Basak, Rajarshi Bhattacharya, Rukhshanda Husain, and Agniv Chatterjee, "An exceedingly simple consistency regularization method for semi-supervised medical image segmentation," in *IEEE ISBI*, 2022.
- [10] Liang Chen, Paul Bentley, Kensaku Mori, Kazunari Misawa, et al., "Self-supervised learning for medical image analysis using image context restoration," *Medical image analysis*, 2019.
- [11] Aiham Taleb, Winfried Loetzsch, Noel Danz, Julius Severin, Thomas Gaertner, Benjamin Bergner, and Christoph Lippert, "3d self-supervised methods for medical imaging," in *NeurIPS*, 2020.
- [12] Hritam Basak, Sagnik Ghosal, and Ram Sarkar, "Addressing class imbalance in semi-supervised image segmentation: A study on cardiac mri," in *MICCAI*, 2022.
- [13] Krishna Chaitanya, Ertunc Erdil, Neerav Karani, and Ender Konukoglu, "Contrastive learning of global and local features for medical image segmentation with limited annotations," in *NeurIPS*, 2020.
- [14] Aaron van den Oord, Yazhe Li, and Oriol Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [15] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li, "Dense contrastive learning for self-supervised visual pre-training," in *CVPR*, 2021.
- [16] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny, "Barlow twins: Self-supervised learning via redundancy reduction," in *ICML*, 2021.
- [17] Adrien Bardes, Jean Ponce, and Yann LeCun, "Vicreg: Variance-invariance-covariance regularization for self-supervised learning," in *ICLR*, 2022.
- [18] Dewen Zeng, Yawen Wu, Xinrong Hu, Xiaowei Xu, et al., "Positional contrastive learning for volumetric medical image segmentation," in *MICCAI*, 2021.
- [19] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, et al., "Pseudoseg: Designing pseudo labels for semantic segmentation," in *ICLR*, 2020.
- [20] Wenjia Bai, Ozan Oktay, Matthew Sinclair, et al., "Semi-supervised learning for network-based cardiac mr image segmentation," in *MICCAI*, 2017.
- [21] Krishna Chaitanya, Neerav Karani, Christian F Baumgartner, Anton Becker, Olivio Donati, and Ender Konukoglu, "Semi-supervised and task-driven data augmentation," in *IPMI*, 2019.
- [22] Kaiming He, Xiangyu Zhang, et al., "Deep residual learning for image recognition," in *CVPR*, 2016.
- [23] Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, et al., "Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved?," *IEEE TMI*, 2018.
- [24] Xiahai Zhuang, "Challenges and methodologies of fully automatic whole heart segmentation: a review," *Journal of healthcare engineering*, 2013.
- [25] Xiahai Zhuang and Juan Shen, "Multi-scale patch and multi-modality atlases for whole heart segmentation of mri," *Medical image analysis*, 2016.
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, et al., "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*, 2019.
- [27] Xinrong Hu, Dewen Zeng, Xiaowei Xu, and Yiyu Shi, "Semi-supervised contrastive learning for label-efficient medical image segmentation," in *MICCAI*, 2021.
- [28] Yicheng Wu, Zongyuan Ge, Donghao Zhang, Minfeng Xu, Lei Zhang, Yong Xia, and Jianfei Cai, "Mutual consistency learning for semi-supervised medical image segmentation," *Medical Image Analysis*, 2022.