

CrossMatch: Enhance Semi-Supervised Medical Image Segmentation with Perturbation Strategies and Knowledge Distillation

Bin Zhao, Chunshi Wang, and Shuxue Ding

Abstract—Semi-supervised learning for medical image segmentation presents a unique challenge of efficiently using limited labeled data while leveraging abundant unlabeled data. Despite advancements, existing methods often do not fully exploit the potential of the unlabeled data for enhancing model robustness and accuracy. In this paper, we introduce CrossMatch, a novel framework that integrates knowledge distillation with dual perturbation strategies—image-level and feature-level—to improve the model's learning from both labeled and unlabeled data. CrossMatch employs multiple encoders and decoders to generate diverse data streams, which undergo self-knowledge distillation to enhance consistency and reliability of predictions across varied perturbations. Our method significantly surpasses other state-of-the-art techniques in standard benchmarks by effectively minimizing the gap between training on labeled and unlabeled data and improving edge accuracy and generalization in medical image segmentation. The efficacy of CrossMatch is demonstrated through extensive experimental validations, showing remarkable performance improvements without increasing computational costs. Code for this implementation is made available at <https://github.com/AiEson/CrossMatch.git>.

Index Terms—Semi-supervised segmentation; Self-knowledge distillation; Image perturbation

I. INTRODUCTION

SEMANtic segmentation, as a precise classification technique at the pixel level, plays a vital role in the field of medical image analysis. Especially when dealing with complex three-dimensional CT and MRI data, although fully supervised learning methods can achieve high-precision segmentation results, their application is severely limited by the high cost of manual annotation and the complexity of operation. In

order to overcome this bottleneck, semi-supervised medical image segmentation methods have emerged and demonstrate great potential [1]. The core of these approaches lies in the effective combination of a small amount of annotated data and a large amount of unlabeled data, aiming to reduce the high cost of annotation and achieve accurate segmentation while promoting the widespread application in clinical and other scenarios.

The main challenge in semi-supervised learning (SSL) is how to effectively exploit the potential of unlabeled data. Recent research has shifted from relying on adversarial training mechanisms based on Generative Adversarial Networks (GANs) [2], [3] to incorporating various methods including consistency regularization and self-training [4]–[8]. In particular, collaborative teaching and mutual learning paradigms [9]–[13] have proven to be highly promising strategies, often involving the parallel training of two models. Knowledge distillation strategies have also been widely employed to optimize model structures, enabling efficient training and good performance by simplifying models.

In handling unlabelled image data, the application of both image-level and feature-level perturbations has become a common strategy. Image-level perturbations, such as random rotations, scaling, flipping, and color adjustments, enhance model robustness to input variations through controlled deformations and modifications of the input images. Moreover, more complex image-level perturbations like CutMix [14] and MixUp [15] create new training samples by blending regions between images and combining them at the pixel level, thus simulating a more diverse data distribution and further improving the model's generalization to unseen data. Feature-level perturbations, particularly those applied to features extracted by the Encoder, have not been fully explored and hold substantial potential. This approach introduces weak to strong feature perturbations during the Decoder decoding process, utilizing the model's prediction consistency under various perturbation conditions to train the model, which ensures stability in performance when the model faces the same image segmentation tasks. For example, feature-level perturbations can be achieved by adding random noise, applying various types of Dropout, etc. [8], [16]. These perturbations not only simulate potential variations in the data but also promote generalization in the model's deep feature abstraction and decoding processes, thereby achieving more accurate and

This work is supported in part by the National Natural Science Foundation of China (Grant No.62076077), the Project of Improving the Basic Scientific Research Ability of Young and Middle-Aged Teachers in Universities of Guangxi Province (Grant No.2023KY0223), Youth Science Foundation of Guangxi Natural Science Foundation (Grant No.2023GXNSFBA026018) and the Guangxi Science and Technology Major Project (Grant No.AA22068057). (Corresponding author: Shuxue Ding.)

Bin Zhao, Chunshi Wang and Shuxue Ding are with School of Artificial Intelligence, Guangxi Colleges and Universities Key Laboratory of AI Algorithm Engineering, Guilin University of Electronic Technology, Guilin 541004, Guangxi, China (e-mail: zhaobinnku@mail.nankai.edu.cn, 2101630316@mails.guet.edu.cn, sdting@guet.edu.cn).

Bin Zhao and Chunshi Wang contribute equally to this paper.

robust predictive performance on unlabeled data.

Knowledge Distillation(KD) has demonstrated significant potential in semi-supervised learning for medical image segmentation [13], [17]. Typically, KD involves a pre-trained teacher model and a student model that needs to learn. However, Self-KD methods [18], [19] primarily rely on soft labels generated within a single model to guide the training process, instead of depending on traditional hard labels or an additional teacher model. These methods use the model's self-generated predictions during training as guidance, refining the model's feature extraction and classification capabilities through iterative processes. This self-teaching method not only reduces dependence on costly manually annotated data but also significantly enhances the model's adaptability and prediction accuracy on unlabeled data. Self-KD promotes deeper feature learning and more stable model behavior by reinforcing the model's reliance on its own predictions. Particularly for medical imaging data, this strategy effectively improves model robustness and accuracy when dealing with highly variable and individually distinct medical images.

Inspired by Self-KD and image perturbation, we have designed an innovative self-training consistency regularization framework called CrossMatch for semi-supervised medical image segmentation. This framework employs a range of image-level and feature-level perturbations from weak to strong and explores the potential of unlabeled data through a more systematic and in-depth approach. Specifically, CrossMatch applies two types of image-level and two types of feature-level perturbations to unlabeled data to create four distinct data streams. These data streams vary in accuracy of output prediction depending on the degree of perturbation to which they are subjected, where the stronger streams guide the weaker ones. In this process, image-level perturbations are implemented as applications of different encoders, while feature-level perturbations are used to generate varied outputs for the same decoder. Through these perturbations, CrossMatch engages in internal knowledge distillation by leveraging the model's consistency across different perturbation intensities, which not only optimizes the model's learning from unlabeled data but also enhances its generalization capability. CrossMatch ensures the stability and accuracy of model outputs, thereby exhibiting superior performance in applications requiring high precision, such as medical image segmentation.

In summary, our contributions are fourfold:

- (1) We propose a consistency regularization framework based on knowledge distillation and image perturbations, which focuses on the exploration of unlabeled data and the transfer of self-knowledge.
- (2) We equate different feed-forward flows to different encoders and decoders, applying the concept of knowledge distillation to semi-supervised semantic segmentation.
- (3) We compute adjacent Self-KD losses between the same decoders, which can bridge the capability gap between the teacher and student models.
- (4) Experimental results on two benchmark datasets demonstrate that CrossMatch achieves significant performance improvements compared to previous state-of-the-art methods.

II. RELATED WORK

A. Semi-Supervised Learning

In the field of SSL, a key challenge is designing effective supervision signals for unlabeled data. Currently, there are two main strategies to address this issue: entropy minimization [20]–[23] and consistency regularization [6], [8], [24], [25]. Entropy minimization, popular for its simplicity, involves automatically assigning pseudo-labels to unlabeled data and using them for retraining along with labeled data. And consistency regularization is based on the assumption that a model's predictions for the same unlabeled sample under different perturbations should remain consistent. For example, FixMatch [6] combines the advantages of entropy minimization and consistency regularization to apply strong perturbations to unlabeled images and use the predictions from their weakly perturbed versions to guide model training. Advanced methods like FreeMatch [26] further refine this strategy, providing rigorous mathematical justification for its motivation and using thresholds to filter out low-confidence labels, thereby enhancing the model's accuracy and reliability. Our CrossMatch draws on the basic framework of FixMatch, without any bells and whistles. It only uses the most common way to verify the theoretical effectiveness of this method, and also demonstrates its important value in practical applications.

B. Semi-Supervised Semantic Segmentation

Semi-supervised learning based methods have achieved exciting results in classification task, of which several works have been further developed for semantic segmentation. A popular class of methods [9], [27], [28] is based on the Mean Teacher [29] setting. For instance, UA-MT [27] introduces a self-aware model of uncertainty to design thresholds to filter out uncertain regions between teachers and students to get more meaningful and reliable predictions. BCP [28] notes that in semi-supervised learning, the distributions learned in labeled and unlabeled data are not consistent, and proposes a symmetric approach to use both kinds of data so as to maintain the consistency between the two distributions, thus allowing the model to learn common features. CAML [9] pays further attention to the potential of labeled data and proposes Correlation Aware Mutual Learning framework to utilize labeled data to guide the extraction of information from unlabeled data. CPS [4] utilizes a cross-teacher module to simultaneously reduce the coupling among peer networks and the error accumulation between teacher and student networks.

Another mainstream class of semi-supervised segmentation methods is based on the idea of co-training. The networks learn together and transfer knowledge to each other [10], [13]. To transfer knowledge efficiently between networks, knowledge distillation is also a common strategy in semi-supervised semantic segmentation [13]. Besides, some method uses pseudo segmentation map obtained from one network to supervise the other one [30]. MC-Net [11] and MC-Net+ [12] use a shared encoder for feature extraction, and then feed the features into multiple decoders with the same structure but different parameters to get multiple outputs. All these methods require multiple networks, encoders or decoders for training.

Methods based on self-training have begun to evolve rapidly since FixMatch [6] introducing consistency regularisation to self-training, and FixMatch has gradually become the baseline for many methods. DTC [31] uses a dual-task deep network to jointly predict pixel segmentation maps and geometrically-aware level-set representations of a target by introducing dual-task consistency regularisation between level-set derived segmentation maps and directly predicted segmentation maps for labelled and unlabelled data. SASSNet [32] introduces a multi-task deep network that jointly predicts semantic segmentation and symbolic distance maps (SDM) of object surfaces, while introducing adversarial loss in order to capture shape-aware features. URPC [33] enhances pyramid-consistent regularisation using multi-scale uncertainty correction for more efficient semi-supervised medical image segmentation. SS-Net [34] addresses the challenges of semi-supervised medical image segmentation by simultaneously exploring pixel-level smoothness and inter-class separation. UniMatch [8] achieves better segmentation results by consistency regularisation using multiple strongly augmented branches and a dual-stream perturbation feature perturbation. Our CrossMatch also follows this single-stage framework, i.e., there is only one model in our approach. Unlike the above works, our CrossMatch introduces Self-KD and feature perturbation into semi-supervised medical image segmentation, achieving efficient self-knowledge transfer under a broader perturbation space.

III. METHOD

A. Preliminaries

Semi-supervised medical image segmentation aims to fully explore an unlabeled image set $\mathcal{D}^u = \{x_1^u, \dots, x_n^u\}$ and integrate it with a labeled image set $\mathcal{D}^l = \{(x_1^l, y_1^l), \dots, (x_n^l, y_n^l)\}$ that contains limited annotations for precise semantic segmentation. The performance of series methods like FixMatch [6] largely depends on well-designed image-level perturbation strategies. Specifically, each unlabeled input is subjected to two types of perturbations: \mathcal{A}^w denotes a weak perturbation operator, and \mathcal{A}^s denotes a strong perturbation operator. Given an unlabeled input x^u , we have

$$\begin{cases} x^w = \mathcal{A}^w(x^u) \\ x^s = \mathcal{A}^s(\mathcal{A}^w(x^u)), \end{cases} \quad (1)$$

where x^w and x^s represent the weakly perturbed image and the strongly perturbed image, respectively.

B. Knowledge Distillation

In machine learning tasks, Kullback-Leibler (KL) divergence is often used to measure the discrepancy between different probability distributions. In knowledge distillation, it is commonly employed to gauge the performance gap between teacher and student models,

$$\mathcal{L}_{kd}^{KL}(p^{w_i}, p^{w_j}) = KL(\sigma(p^{w_i}/T), \sigma(p^{w_j}/T)), \quad (2)$$

where p^{w_i} and p^{w_j} represent the probability distribution outputs by the teacher and student models for unlabeled samples, respectively. Here, $\sigma(\cdot)$ denotes the softmax function, which transforms logits into normalized probability distributions, and

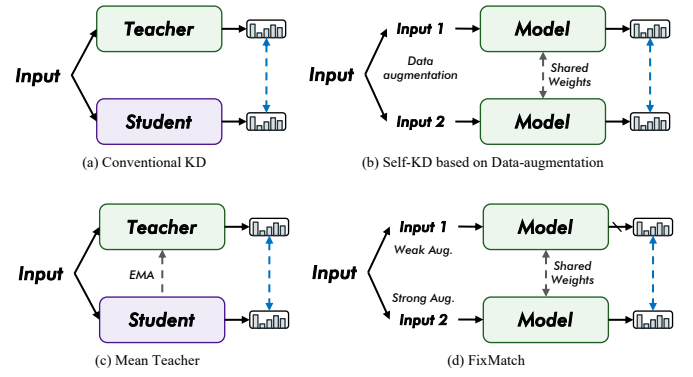


Fig. 1. Comparison of different types of KD and SSL methods. (a) Traditional KD that requires pre-training of the teacher model. (b) Self-KD based on data augmentation. (c) Mean Teacher. (d) FixMatch.

T is a hyperparameter known as the temperature coefficient. This coefficient smooths the probability distributions, allowing the student model not only to focus on the most likely predicted category but also to learn more about the relative information between categories from the teacher model. Thus, within the framework of knowledge distillation, the KL divergence loss function \mathcal{L}_{kd}^{KL} aims to minimize the difference between the probability distributions of the teacher model w_i and the student model w_j that after softmax processing and temperature reduction. Through this way, the student model can emulate the teacher model's 'soft' output predictions, thereby facilitating effective transfer of complex and high-quality knowledge.

DMD [13] delves into knowledge distillation methods specifically for semi-supervised medical image segmentation, and proposes to use Dice loss as an alternative to KL divergence loss. This approach effectively addresses the common issue of foreground and background class imbalance in segmentation tasks. Compared to KL divergence loss, Dice loss can more aptly handle such imbalances, thereby enhancing the model's segmentation performance,

$$\mathcal{L}_{kd}^{Dice}(p^{w_i}, p^{w_j}) = \text{Dice}(\sigma(p^{w_i}/T), \sigma(p^{w_j}/T)). \quad (3)$$

As illustrated in Figure 1, a careful comparison of KD methods and SSL methods reveals remarkable similarities in the structure, design and development of the networks. Based on this observation, we hypothesize that KD methods can be readily adapted to SSL tasks.

C. Feature Perturbation

The performance of FixMatch [6] and its related works, such as UniMatch [8] and ReMixMatch [24], largely depends on the effectiveness of the well-designed image-level perturbation strategies. As mentioned earlier, pseudo-labels generated from the weakly perturbed images x^w are used to supervise the strongly perturbed images x^s to achieve consistency learning. The greater the difference in the degree of perturbation between x^w and x^s , the larger the perturbation space during training. Generally, the perturbation space should be within an appropriate range according to [35]: too small a difference may diminish the effect of consistency regularization, while

excessive perturbation can have a catastrophic impact on the clean data distribution.

Although image-level perturbations have widely used in numerous methods, their performance in semi-supervised image segmentation tasks highly depends on how researchers meticulously tailor perturbation schemes for specific datasets to ensure an appropriate perturbation space is constructed. This process often involves a high demand for expert knowledge and trial-and-error costs, especially in the field of medical image processing, where finding suitable perturbation strategies can become one of the main challenges demanding significant effort.

To mitigate the aforementioned issues, the literature [8], [16] suggests perturbing the high-dimensional features of x^w at the bottleneck section of the segmentation network by using different levels of perturbation to create varied feed-forward flows. Segmentation models typically employ an encoder-decoder structure, where e denotes the encoder and d denotes the decoder. For FixMatch, the weak perturbation feed-forward flow for an unlabeled sample x^u can be represented as:

$$x^u \rightarrow \mathcal{A}^w \rightarrow e \rightarrow d \rightarrow p^w, \quad (4)$$

where $x^u \rightarrow \mathcal{A}^w = x^w$. Based on this format, we can consider inserting a new perturbation \mathcal{P}^r between $e \rightarrow d$ to achieve a larger perturbation space and obtain a new perturbation feed-forward flow:

$$x^u \rightarrow \mathcal{A}^w \rightarrow e \rightarrow \mathcal{P}^r \rightarrow d \rightarrow p_r^w, \quad (5)$$

where r differentiates the intensity of feature perturbations, which will be detailed later. Similarly, the feature perturbation flow for the strongly perturbed input can be represented as:

$$x^u \rightarrow \mathcal{A}^w \rightarrow \mathcal{A}^s \rightarrow e \rightarrow \mathcal{P}^r \rightarrow d \rightarrow p_r^s. \quad (6)$$

For consistent notation, the aforementioned flows can be succinctly expressed as follows:

$$\begin{aligned} x^w &\rightarrow e \rightarrow \mathcal{P}^r \rightarrow d \rightarrow p_r^w \\ x^s &\rightarrow e \rightarrow \mathcal{P}^r \rightarrow d \rightarrow p_r^s \end{aligned} \quad (7)$$

Let \mathcal{P}^n denote no feature perturbation applied, then p^w can be obtained through the flow $x^w \rightarrow e \rightarrow \mathcal{P}^n \rightarrow d \rightarrow p^w$. Based on this, we can compute the loss function for FixMatch with feature perturbation as:

$$\frac{1}{B_u} \sum \mathbb{1}(\max(p^w) \geq \tau) (H(p^w, p_r^s) + H(p^w, p_r^w)), \quad (8)$$

where $H(\cdot)$ denotes the entropy minimizing the discrepancy between two probability distributions. $\mathbb{1}(\cdot > \tau)$ is the indicator function for confidence-based thresholding with the threshold τ . B_u is the batch size for unlabeled data.

D. Multiple Encoders and Decoders

In Eq. 5 and Eq. 6, we re-examine $\mathcal{A}^w \rightarrow e$ and $\mathcal{A}^w \rightarrow \mathcal{A}^s \rightarrow e$, whose aim is to achieve consistency in model predictions by applying varying degrees of image-level perturbations. In contrast, Mean Teacher (MT) [29] and UA-MT [27] achieve a similar effect by the utilization of Exponential Moving Average (EMA), which allows a model to derive one or more

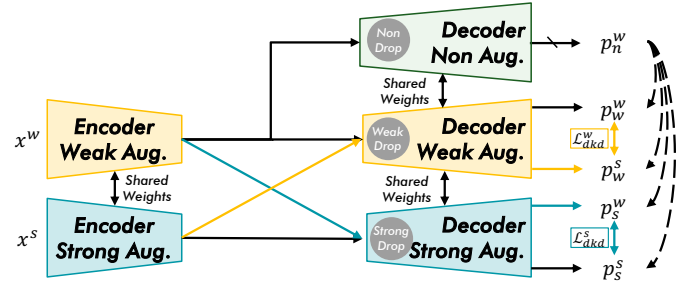


Fig. 2. Overview of our proposed CrossMatch. CrossMatch integrates the core ideas of Self-KD and SSL by enhancing performance through the derivation and mutual distillation of multiple encoder-decoder architectures.

other models for training. By integrating this different method of achieving consistency with knowledge distillation, we can consider different levels of perturbations as encoders with varying capabilities. Hence, let $e^w = \mathcal{A}^w \rightarrow e$ represent the weak perturbation encoder and $e^s = \mathcal{A}^w \rightarrow \mathcal{A}^s \rightarrow e$ represent the strong perturbation encoder, where e^w is clearly outperforms e^s , that is, e^w is less perturbed and its resulting prediction is clearly more accurate.

Similarly, based on the relationship between the encoders and perturbations mentioned above, the newly introduced feature perturbations can be viewed as perturbations to the decoder's capabilities, where $d^r = \mathcal{P}^r \rightarrow d$ represents the high-dimensional features entering the decoder being perturbed by \mathcal{P}^r . Consistent with the form of the encoders, here we propose using both strong and weak feature perturbations, namely \mathcal{P}^w and \mathcal{P}^s , thus yielding three different decoders d^n , d^w and d^s . Consequently, we now have two equivalent encoders and three equivalent decoders.

E. CrossMatch

Fig. 1 demonstrates various KD and SSL methods, revealing a high similarity between them, thus prompting the idea of integrating knowledge distillation into SSL tasks. The purpose of knowledge distillation is to transfer more accurate knowledge to another network model. Self-KD represents a unique distillation mode where the student model learns from knowledge generated from its own outputs, typically involving the backpropagation of deep information to guide the training of earlier layers. This approach incorporates image-level perturbations to achieve varying capabilities, as depicted in Fig. 1 (b), a process very similar to that in Fig. 1 (c).

Based on this, we propose CrossMatch, whose overall structure is depicted in Fig. 2. CrossMatch employs multiple different encoders and decoders, namely e^w , e^s , d^n , d^w and d^s as mentioned in Sec. III-D, to generate diverse outputs. These combinations produce outputs denoted as p_j^i . Specifically, x^u passes through the following feed-forward flow to form different outputs:

$$x^u \rightarrow e^i \rightarrow d^j \rightarrow p_j^i, \quad (9)$$

where $e^i \in \{e^w, e^s\}$ and $d^j \in \{d^n, d^w, d^s\}$. Notably, p_n^w experiences the least perturbation and is the most accurate.

Here, we designate p_n^w as the Teacher, with all other outputs, which have undergone feature perturbations, acting as students. The Teacher is required to impart knowledge to all students, leading to the following teacher distillation loss:

$$\mathcal{L}_{tkd} = \frac{1}{B_u} \sum \mathbb{1}(\max(p_n^w) \geq \tau) \sum_i \sum_j H(p_n^w, p_j^i). \quad (10)$$

Observing that a decoder outputs two segmentation results with varying degrees of perturbation, we can facilitate mutual distillation between these outputs. Specifically, we consider p_w^w and p_s^w as teaching assistants, each imparting knowledge to p_n^s and p_n^s , respectively. These teaching assistants are relative to the same decoder, hence this is referred to as decoder distillation loss:

$$\mathcal{L}_{dkd} = \frac{1}{B_u} \sum \mathbb{1}(\max(p_j^w) \geq \tau) \sum_j H(p_j^w, p_j^s). \quad (11)$$

Eq. 10 and Eq. 11 correspond to the black and colored arrows in Fig. 2, respectively.

In practical implementation, as shown in Fig. 3, we also introduce two image-level strong perturbations (x^{s1} and x^{s2}), which are applied with the same degree of perturbation but differ in perturbation parameters. This aligns with the principles of contrastive learning and has been proven meaningful for our tasks in previous works [8], [36], [37].

Finally, by combining the supervised loss \mathcal{L}_{sup} , the image-level perturbation loss \mathcal{L}_{ip} , we can derive the total loss:

$$\mathcal{L}_{total} = \mathcal{L}_{sup} + \mathcal{L}_{ip} + (1 - \eta)\mathcal{L}_{tkd} + \eta\mathcal{L}_{dkd}, \quad (12)$$

where \mathcal{L}_{sup} consists of Dice and CrossEntropy losses, and \mathcal{L}_{ip} denotes the supervision of p_n^w over p_n^{s1} and p_n^{s2} as shown in Fig. 3, which involves calculating Dice for the two strongly perturbed unlabeled predictions and averaging them. η is used to balance the proportions between the two distillation loss.

F. Performance-Friendly Implementation

The multi-encoder and decoder architecture of CrossMatch is straightforward and intuitive, and it conveniently allows for the introduction of knowledge distillation at the output stage. However, each image-level and feature-level perturbation requires multiple forward propagations, significantly increasing training costs. In response, we introduce methods for equivalent and performance-friendly implementations.

1) *Equivalent Multi-encoders*: For image-level perturbations, we follow the practice of most existing works by applying data augmentation operations before the input reaches the encoder stage. Specifically, we perform weak perturbation operator \mathcal{A}^w and strong perturbation operator \mathcal{A}^s in parallel on CPUs to ensure that the augmentation process does not occupy additional computational time resources during model iterations, efficiently utilizing hardware parallel capabilities.

2) *Equivalent Multi-decoders*: As shown in Fig. 2, our CrossMatch requires four feature perturbations to produce four different outputs from two decoders. For ease of explanation, let $h^i = x^u \rightarrow e^i$ denote the intermediate features produced by different encoders. Considering that each Mini Batch operates independently during gradient computation, suppose $h^i \in$

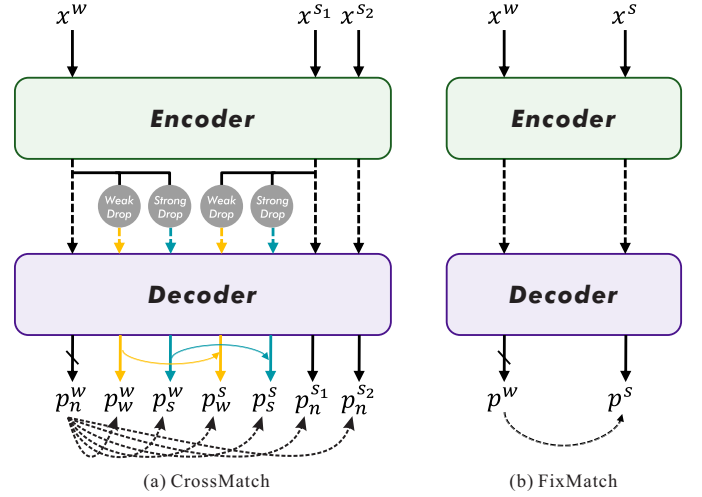


Fig. 3. (a) Our proposed CrossMatch method with Weak Drop denoting \mathcal{P}^w and Strong Drop denoting \mathcal{P}^s . (b) FixMatch.

$\mathbb{R}^{B \times H \times W \times C}$, we can denote $h_{w,s}^i \in \mathbb{R}^{2B \times H \times W \times C}$, thereby stacking each feature perturbation in the Batch dimension to achieve more efficient computation while maintaining the independence between different feature perturbation outcomes.

As shown in Fig. 3, our CrossMatch does not introduce any additional parameter overhead and adheres to the principles of Self-Training and Self-KD. It only uses image-level and feature perturbations to expand the perturbation space of FixMatch, proving to be more efficient than the EMA method and introducing knowledge distillation into semi-supervised learning tasks. Theoretically, multiple encoders and decoders are introduced, but the implementation employs a more efficient coding method, achieving significant performance improvement while ensuring computational friendliness.

G. The pseudocode of CrossMatch

In summary, we present a self-training framework for multiple encoders and decoders based on knowledge distillation and provide a performance-friendly implementation. Algorithm 1 provides pseudocode in PyTorch style.

IV. EXPERIMENTS

1) *Dataset*: In this study, we utilize the 2018 Left Atrium Segmentation Challenge (LA¹) as a platform to evaluate the proposed CrossMatch. The challenge provides data consisting of 3D Magnetic Resonance Imaging (MRI) scans and their corresponding left atrium segmentation masks, divided into training and validation sets in an 80/20 ratio, with an isotropic resolution of $0.625 \times 0.625 \times 0.625 \text{ mm}^3$. Furthermore, we extend our experimental work to the Automatic Cardiac Diagnosis Challenge (ACDC²). To ensure a fair comparison with previous works, we follow the same experimental setup when reporting the performance on the validation set.

¹www.cardiacatlas.org/atriaseg2018-challenge/

²www.creatis.insa-lyon.fr/Challenge/acdc/

Algorithm 1 Pseudocode of CrossMatch.

```

for  $x_w, x_{s1}, x_{s2}$  in loader_u:
    #  $x_w$ :aug_w( $x_u$ ),  $x_{s1}$ :aug_s( $x_w$ ),  $x_{s2}$ :aug_s( $x_w$ )
    # features of perturbed images
     $h_w, h_s = e(x_w), e(x_{s2})$ 
    # perturbed feature
     $pert_{wn} = h_w$  # none drop
     $pert_{ww}, pert_{ws} = P_w(h_w), P_s(h_w)$ 
     $pert_{sw}, pert_{ss} = P_w(h_s), P_s(h_s)$ 
    # concat in the batch dimension
     $perts = \text{cat}(\text{pert}_{wn}, \text{pert}_{ww}, \text{pert}_{ws}, \text{pert}_{sw}, \text{pert}_{ss})$ 
    # decode and split
     $p_{wn}, p_{ww}, p_{ws}, p_{sw}, p_{ss} = d(perts).chunk(5)$ 
     $p_s = d(e(\text{cat}(x_{s1}, x_{s2})))$ 
    # hard (one-hot) pseudo mask
     $\text{mask}_{wn} = \text{pert}_{wn}.argmax(\text{dim}=1).detach()$ 
    # loss from image-level perturbation
     $L_{ip} = \text{criterion}(p_s, \text{mask}_{wn}.repeat(2, 1, 1))$ 
    # loss from feature-level perturbation
     $\text{stus} = \text{cat}(p_{ww}, p_{ws}, p_{sw}, p_{ss})$ 
     $L_{tkd} = \text{criterion}(\text{stus}, \text{mask}_{wn}.repeat(4, 1, 1))$ 
    # loss from TA and student distillation
     $L_{dkd1} = \text{criterion}(p_{sw}, p_{ww}.argmax(\text{dim}=1))$ 
     $L_{dkd2} = \text{criterion}(p_{ss}, p_{ws}.argmax(\text{dim}=1))$ 
     $L_{dkd} = (L_{dkd1} + L_{dkd2}) / 2.0$ 
    # final unsupervised loss
     $L_u = (L_{ip} + (1-\text{eta}) * L_{tkd} + \text{eta} * L_{dkd}) / 2.0$ 

```

2) Implementation Details: The CrossMatch is implemented based on PyTorch and use V-Net and U-Net as the baseline network for experiments on the LA and ACDC dataset, respectively. Specifically on the LA dataset, CrossMatch is optimized using the AdamW [38] optimizer for 9000 iterations, while on the ACDC dataset, it is trained using the SGD optimizer for 300 epochs. Different batch sizes are set for different datasets, with LA at 4 and ACDC at 12, ensuring an equal number of labeled and unlabeled samples per batch. For image preprocessing, images from the LA dataset are randomly cropped to $112 \times 112 \times 80$, and images from the ACDC dataset were cropped to 256×256 . We set $\eta = 0.3$ for both datasets, and set $\tau = 0.85$ and $\tau = 0.95$ for LA and ACDC, respectively. For performance evaluation, the LA dataset uses a sliding window strategy to achieve comprehensive segmentation of the cardiac area, while the ACDC dataset is evaluated by merging predicted slices into a 3D image. The evaluation metrics including Dice, Jaccard, 95% Hausdorff Distance (95HD) and Average Surface Distance (ASD) are used in this paper. In all CrossMatch experiments, feature perturbations are set as standard dropout. The dropout rates for weak and strong perturbations are set at 25% and 75% respectively. The selection of dropout type and discussion on dropout rates are elaborated in Sec. IV-D.

Notably, to ensure the fairness of the experiments, our results are calculated using the final model weights rather than the best weights saved during training, which also demonstrates the stability of our method.

A. Qualitative Comparison

The Fig. 4 presents some 3D visualization examples of all the compared methods and the corresponding ground truth on LA dataset. It can be observed that our CrossMatch

outperforms other methods in terms of segmentation results. Particularly, our segmentation edges are smoother, with fewer misclassified voxels, more closely mirroring ground truth.

B. Quantitative Comparison

Table I summarizes the quantitative results and reveals that CrossMatch surpasses state-of-the-art (SOTA) techniques on LA dataset. When using 5% of the data with labels (4-label setting), although our Dice and Jaccard are close to those of SOTA methods, CrossMatch achieves significantly better results in the remaining evaluation metrics. Furthermore, significant performance improvements are realized in the scenarios with 8 and 16 labels. Especially using only 10% of the data with labels, CrossMatch exceeds the segmentation results obtained by fully supervised learning of V-Net on 100% of the data with labels, achieving a Dice of 91.33%.

Quantitative results on the ACDC dataset summarized in the Table II further demonstrates the effectiveness of CrossMatch. Particularly our method is more outstanding in terms of performance enhancement, where Dice is increased by 3.89% in the setting of 3-label. The experimental setups listed in both Table I and Table II are the same as those in [9], meaning all results are derived from the final iteration outcomes.

C. Computational Performance Analysis

TABLE III

COMPARISON OF ITERATION TIMES FOR DIFFERENT METHODS. TIME IS RECORDED FROM THE BEGINNING OF DATA MIGRATION TO THE CUDA DEVICE TO THE END OF BACKPROPAGATION. THE TIME AVERAGES ARE TAKEN AFTER $1k$ ITERATIONS.

Time (ms)↓	Method	#Params (M)↓	#Flops (G)↓
22	V-Net	9.443	187.409
273	UA-MT	9.443	187.409
501	SASSNet	20.463	249.194
545	DTC	9.443	187.538
486	MC-Net	12.348	380.394
1269	MC-Net+	15.247	572.229
1057	CAML	19.725	450.677
379	BCP	9.443	187.409
210	ours	9.443	187.409

As described in section III-F, our method also exhibits excellent computational efficiency. For comparative analysis, we have compiled computation performance data from a range of similar works, evaluating them based on their publicly available source codes. The experimental setup is standardized, with all hyperparameters and optimizer configurations identical, and we record the average duration of a single iteration from full data loading onto CUDA devices to the completion of backpropagation, as shown in Table III.

After 1000 iterations, as a straightforward fully supervised learning method, V-Net requires 22 ms per iteration, whereas these semi-supervised learning methods that require a combination of labeled and unlabeled data for training need more time, such as UA-MT [27], MC-Net [11], MC-Net+ [12] and CAML [9] take 273 ms, 486 ms, 1269 ms and 1057 ms

TABLE I
COMPARISONS ON THE LA DATASET. "↑" AND "↓" INDICATE THE LARGER AND THE SMALLER THE BETTER, RESPECTIVELY.

Method	#Scans used		Metrics			
	Lab.	Unlab.	Dice(%)↑	Jaccard(%)↑	95HD(voxel)↓	ASD(voxel)↓
V-Net	4(5%)	0	43.32	31.43	40.19	12.13
V-Net	8(10%)	0	79.99	68.12	21.11	5.48
V-Net	16(20%)	0	86.03	76.06	14.26	3.51
V-Net	80(All)	0	91.14	83.82	5.75	1.52
UA-MT [27] (MICCAI'19)	4(5%)	76(95%)	78.07	65.03	29.17	8.63
SASSNet [32] (MICCAI'20)			79.61	67.00	25.54	7.20
DTC [31] (AAAI'21)			80.14	67.88	24.08	7.18
MC-Net [11] (MICCAI'21)			80.92	68.90	17.25	2.76
URPC [33] (MedIA'22)			80.75	68.54	19.81	4.98
SS-Net [34] (MICCAI'22)			83.33	71.79	15.70	4.33
MC-Net+ [12] (MedIA'22)			83.23	71.70	14.92	3.43
BCP [28] (CVPR'23)			87.52	78.15	8.41	2.64
UniMatch [8] (CVPR'23)			86.08	75.83	12.04	2.85
CAML [9] (MICCAI'23)			87.34	77.65	9.76	2.49
Ours			88.96	80.21	7.75	2.38
UA-MT [27] (MICCAI'19)	8(10%)	72(90%)	85.81	75.41	18.25	5.04
SASSNet [32] (MICCAI'20)			85.71	75.35	14.74	4.00
DTC [31] (AAAI'21)			84.55	73.91	13.80	3.69
MC-Net [11] (MICCAI'21)			86.87	78.49	11.17	2.18
URPC [33] (MedIA'22)			83.37	71.99	17.91	4.41
SS-Net [34] (MICCAI'22)			86.56	76.61	12.76	3.02
MC-Net+ [12] (MedIA'22)			87.68	78.27	10.35	1.85
DMD [13] (MICCAI'23)			89.70	81.42	6.88	1.78
BCP [28] (CVPR'23)			89.55	81.22	7.10	1.69
UniMatch [8] (CVPR'23)			89.09	80.47	12.50	3.59
CAML [9] (MICCAI'23)			89.62	81.28	8.76	2.02
Ours			91.33	84.11	5.29	1.53
UA-MT [27] (MICCAI'19)	16(20%)	64(80%)	88.18	79.09	9.66	2.62
SASSNet [32] (MICCAI'20)			88.11	79.08	12.31	3.27
DTC [31] (AAAI'21)			87.79	78.52	10.29	2.50
MC-Net [11] (MICCAI'21)			90.43	82.69	6.52	1.66
URPC [33] (MedIA'22)			87.68	78.36	14.39	3.52
SS-Net [34] (MICCAI'22)			88.19	79.21	8.12	2.20
MC-Net+ [12] (MedIA'22)			90.60	82.93	6.27	1.58
DMD [13] (MICCAI'23)			90.46	82.66	6.39	1.62
BCP [28] (CVPR'23)			90.18	82.36	6.64	1.61
UniMatch [8] (CVPR'23)			90.77	83.18	7.21	2.05
CAML [9] (MICCAI'23)			90.78	83.19	6.11	1.68
Ours			91.61	84.57	5.36	1.57

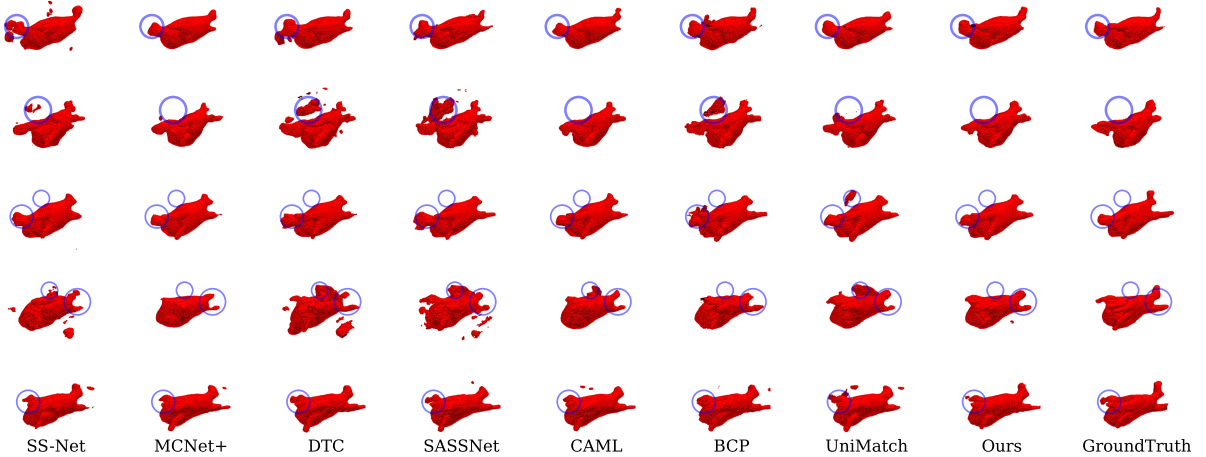


Fig. 4. some visualization examples of several semi-supervised segmentation methods with 10% labeled data and ground truth on LA dataset.

TABLE II
COMPARISONS ON THE ACDC DATASET. "↑" AND "↓" INDICATE THE LARGER AND THE SMALLER THE BETTER, RESPECTIVELY.

Method		#Scans used		Metrics			
		Lab.	Unlab.	Dice(%)↑	Jaccard(%)↑	95HD(voxel)↓	ASD(voxel)↓
U-Net		3(5%)	0	47.83	37.01	31.16	12.62
U-Net		7(10%)	0	79.41	68.11	9.35	2.70
U-Net		70(All)	0	91.44	84.59	4.30	0.99
UA-MT [27]	(MICCAI'19)	3(5%)	67(95%)	46.04	35.97	20.08	7.75
SASSNet [32]	(MICCAI'20)			57.77	46.14	20.05	6.06
DTC [31]	(AAAI'21)			56.90	45.67	23.36	7.39
MC-Net [11]	(MICCAI'21)			62.85	52.29	7.62	2.33
URPC [33]	(MedIA'22)			55.87	44.64	13.60	3.74
SS-Net [34]	(MICCAI'22)			65.82	55.38	6.67	2.28
DMD [13]	(MICCAI'23)			80.60	69.08	5.96	1.90
UniMatch [8]	(CVPR'23)			84.38	75.54	5.06	1.04
Ours				88.27	80.17	1.53	0.46
UA-MT [27]	(MICCAI'19)	7(10%)	63(90%)	81.65	70.64	6.88	2.02
SASSNet [32]	(MICCAI'20)			84.50	74.34	5.42	1.86
DTC [31]	(AAAI'21)			84.29	73.92	12.81	4.01
MC-Net [11]	(MICCAI'21)			86.44	77.04	5.50	1.84
URPC [33]	(MedIA'22)			83.10	72.41	4.84	1.53
SS-Net [34]	(MICCAI'22)			86.78	77.67	6.07	1.40
DMD [13]	(MICCAI'23)			87.52	78.62	4.81	1.60
UniMatch [8]	(CVPR'23)			88.08	80.10	2.09	0.45
Ours				89.08	81.44	1.52	0.52

per iteration, respectively. In contrast, thanks to its streamlined structure and self-training pipeline, our CrossMatch only requires 210 ms per iteration, which is significantly lower than other semi-supervised segmentation methods, thus highlighting its efficient computational characteristics.

D. Ablation Study

TABLE IV
ABLATION STUDY OF η AT ALL METRICS ON THE LA DATASET. 10% OF LABELED DATA ARE USED FOR TRAINING IN THIS ABLATION STUDY.

η	Dice(%)↑	Jaccard(%)↑	95HD(voxel)↓	ASD(voxel)↓
0.1	90.88%	83.35%	7.80	1.93
0.15	90.55%	82.92%	7.93	2.51
0.2	91.03%	83.62%	6.44	1.99
0.25	91.28%	84.01%	5.77	1.62
0.3	91.33%	84.11%	5.29	1.53
0.35	90.78%	83.21%	6.91	1.88
0.4	87.25%	77.80%	9.80	2.37
0.45	90.53%	82.78%	6.54	1.83
0.5	87.78%	79.22%	6.76	1.90

Table IV displays the results of ablation experiments on the parameter of η in the setting of 10% of the data with labels on LA dataset. The results indicate that an η value of 0.3 yields the best performance, surpassing other values across all metrics. Consequently, we have set η at 0.3 for all experiments in this study.

Table V presents the results of ablation studies on the performance gap between decoders using 10% of the data with labels on LA dataset. It is observed that the optimal model performance is achieved when the performance gap is 0.5. Notably, we also explore the scenario where the decoders are completely consistent, that is, when the performance gap is zero. In this case, the strong perturbation encoder and the

TABLE V
ABLATION STUDY OF THE PERFORMANCE GAP BETWEEN STRONG AND WEAK DECODERS. 10% OF LABELED DATA ARE USED FOR TRAINING IN THIS ABLATION STUDY.

Bottom	Top	Perf. Gap	Dice(%)↑	Jaccard(%)↑	95HD(voxel)↓	ASD(voxel)↓
0.500	0.500	0.00	91.20	83.87	5.41	1.80
0.375	0.625	0.25	91.07	83.65	5.47	1.59
0.250	0.750	0.50	91.33	84.11	5.29	1.53
0.125	0.875	0.75	90.22	82.32	6.59	1.76




Fig. 5. Ablation study on the efficacy of various feature perturbation strategies in our method.

weak perturbation encoder apply identical perturbations, and the model degenerates into a UniMatch with an additional \mathcal{L}_{dkd} .

Fig. 5 shows the results of ablation studies on the type of Dropout used under different training sample ratios on LA dataset. We investigate three different types of Dropout available in PyTorch: standard Dropout3D, AlphaDropout, and FeatureAlphaDropout [39]. It is evident that using the standard Dropout3D as our feature perturbation strategy results in the best performance across all three data splits, followed by

AlphaDropout and FeatureAlphaDropout. This may be due to the latter two inducing stronger feature perturbations, resulting in an increased Performance Gap between decoders, which is detrimental to the correct transfer of knowledge in knowledge distillation.

V. CONCLUSION

We have re-evaluated the role of Self-Knowledge in semi-supervised medical image segmentation and cleverly integrated feature perturbation, consistency regularization and Knowledge Distillation to propose a Self-Training segmentation method named CrossMatch. We rethink the role of perturbations in semi-supervised tasks and suggest using multiple equivalent encoders and decoders to play roles at different learning stages to expand the traditional teacher-student model, aiming to reduce the capability gap between different roles. Specifically, we derive two encoders from image-level perturbations and three decoders from feature-level perturbations, designating the unperturbed feed-forward flow as the teacher, to perform knowledge distillation on the four groups of outcomes produced by the aforementioned encoder and decoder combinations. Additionally, we utilize the properties of Mini Batches to optimize the performance of our method and provide a quantitative iteration time comparison table. Our CrossMatch demonstrate robust performance on two benchmark datasets (LA and ACDC), showing significant improvements over SOTA methods. Extensive ablation studies further validate the assumptions and design of our method.

REFERENCES

- [1] K. Han, V. S. Sheng, Y. Song, Y. Liu, C. Qiu, S. Ma, and Z. Liu, "Deep semi-supervised learning for medical image segmentation: A review," *Expert Systems with Applications*, p. 123052, 2024.
- [2] N. Souly, C. Spampinato, and M. Shah, "Semi supervised semantic segmentation using generative adversarial network," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5688–5696.
- [3] S. Mittal, M. Tatarchenko, and T. Brox, "Semi-supervised semantic segmentation with high-and low-level consistency," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 4, pp. 1369–1379, 2019.
- [4] X. Chen, Y. Yuan, G. Zeng, and J. Wang, "Semi-supervised semantic segmentation with cross pseudo supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2613–2622.
- [5] H. Hu, F. Wei, H. Hu, Q. Ye, J. Cui, and L. Wang, "Semi-supervised semantic segmentation via adaptive equalization learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 22 106–22 118, 2021.
- [6] K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. A. Raffel, E. D. Cubuk, A. Kurakin, and C.-L. Li, "Fixmatch: Simplifying semi-supervised learning with consistency and confidence," *Advances in neural information processing systems*, vol. 33, pp. 596–608, 2020.
- [7] L. Yang, W. Zhuo, L. Qi, Y. Shi, and Y. Gao, "St++: Make self-training work better for semi-supervised semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4268–4277.
- [8] L. Yang, L. Qi, L. Feng, W. Zhang, and Y. Shi, "Revisiting weak-to-strong consistency in semi-supervised semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7236–7246.
- [9] S. Gao, Z. Zhang, J. Ma, Z. Li, and S. Zhang, "Correlation-aware mutual learning for semi-supervised medical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023, pp. 98–108.
- [10] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4320–4328.
- [11] Y. Wu, M. Xu, Z. Ge, J. Cai, and L. Zhang, "Semi-supervised left atrium segmentation with mutual consistency training," in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part II 24*. Springer, 2021, pp. 297–306.
- [12] Y. Wu, Z. Ge, D. Zhang, M. Xu, L. Zhang, Y. Xia, and J. Cai, "Mutual consistency learning for semi-supervised medical image segmentation," *Medical Image Analysis*, vol. 81, p. 102530, 2022.
- [13] Y. Xie, Y. Yin, Q. Li, and Y. Wang, "Deep mutual distillation for semi-supervised medical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023, pp. 540–550.
- [14] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6023–6032.
- [15] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.
- [16] Y. Liu, Y. Tian, Y. Chen, F. Liu, V. Belagiannis, and G. Carneiro, "Perturbed and strict mean teachers for semi-supervised semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4258–4267.
- [17] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [18] S. Yun, J. Park, K. Lee, and J. Shin, "Regularizing class-wise predictions via self-knowledge distillation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 13 876–13 885.
- [19] L. Zhang, J. Song, A. Gao, J. Chen, C. Bao, and K. Ma, "Be your own teacher: Improve the performance of convolutional neural networks via self distillation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3713–3722.
- [20] Y. Grandvalet and Y. Bengio, "Semi-supervised learning by entropy minimization," *Advances in neural information processing systems*, vol. 17, 2004.
- [21] H. Pham, Z. Dai, Q. Xie, and Q. V. Le, "Meta pseudo labels," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 11 557–11 568.
- [22] Y. Zou, Z. Zhang, H. Zhang, C.-L. Li, X. Bian, J.-B. Huang, and T. Pfister, "Pseudoseg: Designing pseudo labels for semantic segmentation," *arXiv preprint arXiv:2010.09713*, 2020.
- [23] Y. Wang, H. Wang, Y. Shen, J. Fei, W. Li, G. Jin, L. Wu, R. Zhao, and X. Le, "Semi-supervised semantic segmentation using unreliable pseudo-labels," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 4248–4257.
- [24] D. Berthelot, N. Carlini, E. D. Cubuk, A. Kurakin, K. Sohn, H. Zhang, and C. Raffel, "Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring," *arXiv preprint arXiv:1911.09785*, 2019.
- [25] H. Zhang, Z. Zhang, A. Odena, and H. Lee, "Consistency regularization for generative adversarial networks," *arXiv preprint arXiv:1910.12027*, 2019.
- [26] Y. Wang, H. Chen, Q. Heng, W. Hou, Y. Fan, Z. Wu, J. Wang, M. Savvides, T. Shinozaki, B. Raj et al., "Freematch: Self-adaptive thresholding for semi-supervised learning," *arXiv preprint arXiv:2205.07246*, 2022.
- [27] L. Yu, S. Wang, X. Li, C.-W. Fu, and P.-A. Heng, "Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation," in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*. Springer, 2019, pp. 605–613.
- [28] Y. Bai, D. Chen, Q. Li, W. Shen, and Y. Wang, "Bidirectional copy-paste for semi-supervised medical image segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 11 514–11 524.
- [29] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *Advances in neural information processing systems*, vol. 30, 2017.
- [30] H. Xiao, D. Li, H. Xu, S. Fu, D. Yan, K. Song, and C. Peng, "Semi-supervised semantic segmentation with cross teacher training," *Neurocomputing*, vol. 508, pp. 36–46, 2022.
- [31] X. Luo, J. Chen, T. Song, and G. Wang, "Semi-supervised medical image segmentation through dual-task consistency," in *Proceedings of the AAAI*

- conference on artificial intelligence*, vol. 35, no. 10, 2021, pp. 8801–8809.
- [32] S. Li, C. Zhang, and X. He, “Shape-aware semi-supervised 3d semantic segmentation for medical images,” in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I* 23. Springer, 2020, pp. 552–561.
 - [33] X. Luo, G. Wang, W. Liao, J. Chen, T. Song, Y. Chen, S. Zhang, D. N. Metaxas, and S. Zhang, “Semi-supervised medical image segmentation via uncertainty rectified pyramid consistency,” *Medical Image Analysis*, vol. 80, p. 102517, 2022.
 - [34] Y. Wu, Z. Wu, Q. Wu, Z. Ge, and J. Cai, “Exploring smoothness and class-separation for semi-supervised medical image segmentation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2022, pp. 34–43.
 - [35] J. Yuan, Y. Liu, C. Shen, Z. Wang, and H. Li, “A simple baseline for semi-supervised semantic segmentation with strong data augmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 8229–8238.
 - [36] Y. Wang, H. Wang, Y. Shen, J. Fei, W. Li, G. Jin, L. Wu, R. Zhao, and X. Le, “Semi-supervised semantic segmentation using unreliable pseudo-labels,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 4248–4257.
 - [37] S. Liu, S. Zhi, E. Johns, and A. J. Davison, “Bootstrapping semantic segmentation with regional contrast,” *arXiv preprint arXiv:2104.04465*, 2021.
 - [38] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” *arXiv preprint arXiv:1711.05101*, 2017.
 - [39] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, “Self-normalizing neural networks,” *Advances in neural information processing systems*, vol. 30, 2017.