

Bidirectional Channel-selective Semantic Interaction for Semi-Supervised Medical Segmentation

Kaiwen Huang¹, Yizhe Zhang¹, Yi Zhou², Tianyang Xu³, Tao Zhou^{1*}

¹School of Computer Science and Engineering, Nanjing University of Science and Technology, China.

²School of Computer Science and Engineering, Southeast University, China.

³School of Artificial Intelligence and Computer Science, Jiangnan University, China.
taozhou.ai@gmail.com

Abstract

Semi-supervised medical image segmentation is an effective method for addressing scenarios with limited labeled data. Existing methods mainly rely on frameworks such as mean teacher and dual-stream consistency learning. These approaches often face issues like error accumulation and model structural complexity, while also neglecting the interaction between labeled and unlabeled data streams. To overcome these challenges, we propose a Bidirectional Channel-selective Semantic Interaction (BCSI) framework for semi-supervised medical image segmentation. First, we propose a Semantic-Spatial Perturbation (SSP) mechanism, which disturbs the data using two strong augmentation operations and leverages unsupervised learning with pseudo-labels from weak augmentations. Additionally, we employ consistency on the predictions from the two strong augmentations to further improve model stability and robustness. Second, to reduce noise during the interaction between labeled and unlabeled data, we propose a Channel-selective Router (CR) component, which dynamically selects the most relevant channels for information exchange. This mechanism ensures that only highly relevant features are activated, minimizing unnecessary interference. Finally, the Bidirectional Channel-wise Interaction (BCI) strategy is employed to supplement additional semantic information and enhance the representation of important channels. Experimental results on multiple benchmarking 3D medical datasets demonstrate that the proposed method outperforms existing semi-supervised approaches.

Code — <https://github.com/taozh2017/BCSI>

Introduction

Medical image segmentation plays a crucial role in clinical diagnosis, enabling precise identification of lesion areas, organs, and other anatomical structures (Jiao et al. 2024; Mei et al. 2025). However, current models (Cao et al. 2022; Zhou et al. 2023) rely heavily on large annotated datasets for training. This is particularly challenging for medical images, as annotating them requires expert knowledge and significant manual effort. To mitigate this issue, Semi-Supervised Learning (SSL) methods (Chen et al. 2021; Tarvainen and Valpola 2017; Yao, Hu, and Li 2022; Zhang, Wang, and

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

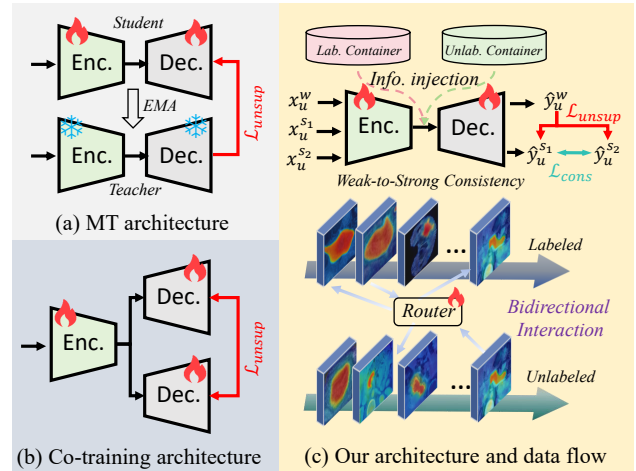


Figure 1: A comparison of semi-supervised learning paradigms: (a) Mean Teacher framework with student and teacher networks, (b) Co-training framework with two sub-nets, and (c) Our proposed framework with bidirectional data-stream interaction and semantic-spatial perturbation under weak-to-strong consistency.

Li 2025) leverage limited labeled data alongside abundant unlabeled data, reducing annotation costs while enhancing model performance.

Consistency regularization is a commonly used framework in semi-supervised learning. For instance, as shown in Fig. 1(a), the Mean Teacher (MT) (Tarvainen and Valpola 2017) method involves a student network to update parameters using gradient propagation, and a teacher network to update parameters using exponential moving average (EMA). However, in the MT framework, the teacher model is prone to being affected by the accumulated prediction errors of the student model, especially when processing unlabeled data. Another widely used strategy is mutual consistency regularization (Wu et al. 2021), which employs multi-decoder or dual-branch architectures (Fig. 1(b)) to improve model generalization. This is achieved by enforcing consistent predictions across different network heads for the same input. While these methods (Peiris et al. 2023; Wang et al. 2023) can enhance model stability, they often converge towards

similar decision boundaries when processing highly noisy or structurally complex medical images. Moreover, excessive reliance on multiple decoders increases computational complexity and lacks sufficient feature diversity in inter-network knowledge transfer, limiting its effectiveness in mitigating overfitting to specific anatomical structures.

In addition to the challenges mentioned, current semi-supervised methods often train labeled and unlabeled data separately, neglecting their potential interaction. As AllSpark (Wang et al. 2024) highlighted that this separation leads to the dominance of labeled data, resulting in low-quality pseudo labels. To address this, this method introduces a channel-wise cross-attention mechanism that regenerates labeled features from unlabeled data. Similarly, SKCDF (Zhang et al. 2025b) proposes decoupling data streams by separating encoder and decoder roles, enhancing semantic representation learning while reducing the adverse effects of unlabeled data on segmentation. However, existing approaches lack bidirectional data interaction and neglect the dual role of feature-level interaction, which acts as both an enhancement and a perturbation between labeled and unlabeled data. Crucially, not all feature channels benefit from “Reborn” process. Excessive manipulation of channels may introduce redundant information, reducing the model’s representational capacity. Without distinguishing each channel’s contribution, excessive intervention can introduce noise, compromising model stability and accuracy. Moreover, cross-stream feature interaction remains largely unexplored in 3D medical image segmentation.

To this end, we propose a novel Bidirectional Channel-selective Semantic Interaction (BCSI) framework for semi-supervised medical segmentation. As shown in Fig. 1(c), our method employs a paradigm based on weak-to-strong consistency learning, which not only avoids the issues of error accumulation and model structural complexity but also allows for effective interaction between labeled and unlabeled data using a single model structure. Specifically, we employ color jitter and copy-paste as two strong augmentation strategies, perturbing the data in the semantic and spatial domains, respectively. To mitigate the excessive noise and interference from irrelevant features during feature interaction between labeled and unlabeled data, we present the Channel-selective Router (CR) and the Bidirectional Channel-wise Interaction (BCI) strategy. The router dynamically selects which feature channels should interact by learning the importance of each channel, ensuring that only relevant features are activated. The BCI strategy facilitates bidirectional interaction across data streams for the selected channels. Moreover, we construct a feature queue for labeled and unlabeled data, enhancing the model’s long-term memory ability during the training process. Overall, our contributions are summarized as follows:

- We propose BCSI, a novel framework for semi-supervised medical image segmentation that facilitates the interaction between labeled and unlabeled data to enhance the ability of feature representations. BCSI employs a spatial-semantic weak-to-strong consistency learning paradigm, mitigating error accumulation while maintaining architectural simplicity.

- We propose a channel-selective Router that dynamically identifies critical feature channels during cross-stream interaction between labeled and unlabeled data, mitigating noise interference.
- We present a bidirectional channel-wise interaction mechanism, primarily focusing on bidirectional interaction between the selected channels and the features stored in the containers.
- Extensive experimental results on multiple benchmark datasets demonstrate that our model significantly improves segmentation accuracy compared to existing semi-supervised segmentation methods.

Related Work

Semi-Supervised Medical Image Segmentation

Semi-supervised medical image segmentation trains with limited labeled and abundant unlabeled data, primarily through pseudo-labeling and consistency learning. Pseudo-labeling methods aim to generate pseudo-labels by utilizing the model’s predictions on unlabeled data (Han et al. 2022; Seibold et al. 2022; Huang et al. 2025a; Zeng et al. 2023). This work (Han et al. 2022) proposes to generate pseudo labels via class-wise feature distances and improving segmentation with prior-guided patch sampling. CoraNet (Shi et al. 2021) estimates uncertainty via segmentation inconsistency under varying misclassification costs and uses a separate self-training strategy to enhance performance. Consistency learning is widely adopted by enforcing prediction invariance under perturbations and pushing decision boundaries toward low-density regions (Huang et al. 2025b). SCOSSL (Xu et al. 2021) enhances prostate segmentation by introducing shadow augmentation to simulate shadow artifacts and shadow dropout to guide boundary inference from shadow-free regions. Tri-U-MT (Wang et al. 2021) enhances segmentation by incorporating auxiliary tasks for semantic and shape guidance and introduces a tripled-uncertainty framework for reliable pseudo label generation. Besides, several studies (Li et al. 2023; Zhang et al. 2025a; Huang et al. 2025c) have been developed to improve the quality of pseudo labels based on foundation models (*e.g.* SAM (Kirillov et al. 2023) and CLIP (Radford et al. 2021)).

Weak-to-Strong Consistency Learning

Weak-to-strong consistency learning promotes segmentation invariance by applying common perturbations such as Gaussian noise, blur, rotation, scaling, and contrast changes to unlabeled data (Sohn et al. 2020; Kim et al. 2022). These perturbations can be applied at both the input (Huang et al. 2022) and feature (Xu et al. 2022) levels. UniMatch (Yang et al. 2023) enhances consistency training by introducing an auxiliary feature perturbation stream and a dual-stream image-level perturbation strategy, enabling stronger and more diverse supervision signals. (Bortsova et al. 2019) proposed enforcing segmentation consistency under elastic deformations using a Siamese network with shared weights and a composite loss that combines supervised and consistency terms. UDC-Net (Li et al. 2021) introduces a dual-

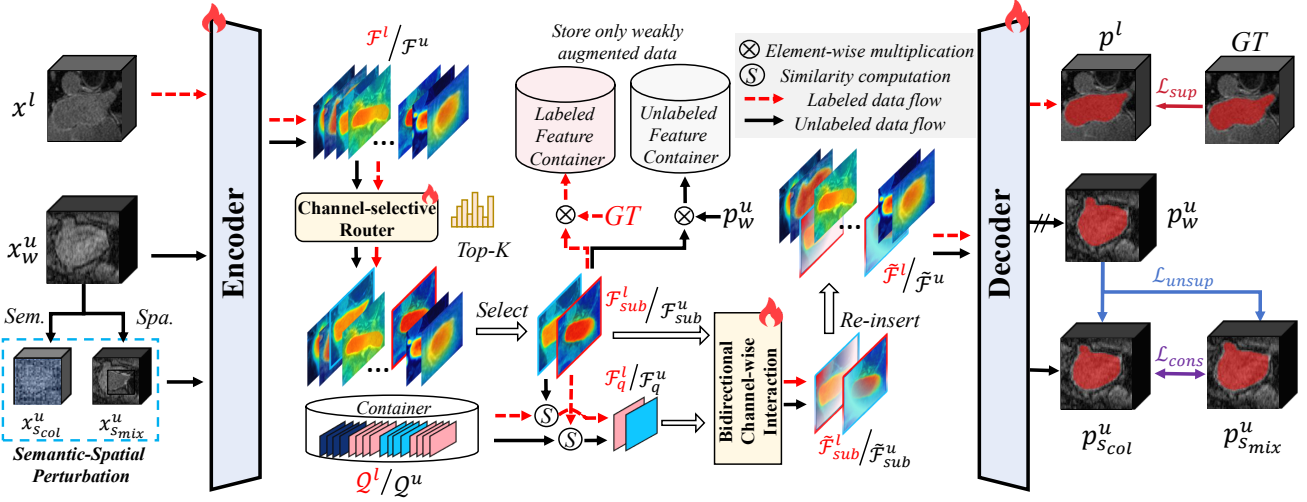


Figure 2: An overview of the proposed BCSI framework, involving an encoder-decoder structure. A channel-selection router then processes the extracted features from labeled and unlabeled data, followed by bidirectional channel-wise interaction for feature enhancement. Finally, the refined features are passed to the decoder under a weak-to-strong consistency learning paradigm.

consistency learning scheme combining image transformation and feature perturbation consistency, guided by quantified uncertainty to improve lesion segmentation.

Methodology

Overview. In the semi-supervised setting, we have a labeled dataset $\mathcal{D}^L = \{(x_i^l, y_i^l)\}_{i=1}^{N_l}$ and an unlabeled dataset $\mathcal{D}^U = \{x_i^u\}_{i=N_l+1}^{N_l+N_u}$, where N_l and N_u denote the numbers of labeled and unlabeled samples, respectively, with $N_l \ll N_u$. $x_i \in \mathbb{R}^{H \times W \times D}$ represents the input volume, and $y_i \in \{0, 1\}^{H \times W \times D}$ denotes the corresponding ground-truth map. Fig. 2 shows the architecture of our framework. Specifically, we first apply two strong augmentations and one weak augmentation to the input data, then pass them through the encoder to extract features. After that, the router selects the top- k channels for cross-data stream interaction. Finally, the features are reinserted into the original channels and stored in their containers.

Semantic-Spatial Perturbation

In semi-supervised learning, data augmentation plays a crucial role in enhancing model generalization (Yang et al. 2023). Departing from conventional approaches like multi-decoder or dual-stream architectures, we propose a Semantic-Spatial Perturbation (SSP) method to bolster model robustness. For semantic perturbation, we leverage color jitter (Cubuk et al. 2020), which introduces random variations in brightness, contrast, and Gaussian noise. Formally, this can be expressed as follows:

$$x_{scol}^u = \alpha \cdot x^u + \beta + \mathcal{N}(\mu, \sigma^2), \quad (1)$$

where x_{scol}^u represents the color jitter augmented unlabeled data. α and β denote the randomly sampled contrast factor and brightness offset, respectively. $\mathcal{N}(\mu, \sigma^2)$ represents Gaussian noise with mean μ and standard deviation σ . For

spatial perturbations, we adopt a copy-paste (Bai et al. 2023) strategy by randomly initializing a binary mask \mathcal{M} to spatially blend volumes from both labeled and unlabeled data. This operation introduces structural diversity by disrupting spatial continuity and generating composite inputs, which can be formulated by

$$x_{smix}^u = \mathcal{M} \odot x^l + (1 - \mathcal{M}) \odot x^u, \quad (2)$$

where x_{smix}^u denotes the copy-paste augmented unlabeled data. \odot denotes element-wise multiplication. Then, the strongly augmented images and weakly augmented images x_w^u are fed into the model, which can be expressed as:

$$p_{scol}^u = f(x_{scol}^u; \theta_f), p_{smix}^u = f(x_{smix}^u; \theta_f), p_w^u = f(x_w^u; \theta_f), \quad (3)$$

where p_{scol}^u and p_{smix}^u represent the predictions for the color-jittering and copy-paste data augmentations, respectively. θ_f represents the parameters of $f(\cdot)$. p_w^u is the prediction for the weakly augmented unsupervised data. All the data share the same model parameters. We use p_w^u as pseudo-labels for the unsupervised loss and apply consistency loss between the two strongly augmented predictions p_{scol}^u and p_{smix}^u . The detailed formulation of the losses in this part will be discussed collectively in Sec. ‘‘Loss Function’’.

Channel-selective Router

To mitigate the insufficient feature-level interaction between labeled and unlabeled data while narrowing their distribution gap, we propose the Channel-selective Router (CR). Our key insight is that feature exchange not only enhances representations but also acts as a structured perturbation to improve model robustness. Unlike prior works that indiscriminately mix features across all channels, CR employs a lightweight and learnable routing mechanism to dynamically select and perturb the most informative channels.

Specifically, given the input volumes from the labeled and unlabeled sets, we first obtain the corresponding features $\mathcal{F}^l \in \mathbb{R}^{C \times h \times w \times d}$ and $\mathcal{F}^u \in \mathbb{R}^{C \times h \times w \times d}$ through a shared encoder, where C , h , w , and d represent the number of channels, height, width, and depth, respectively. These features are then fed into a lightweight router $\mathcal{G}(\cdot)$ to generate channel-wise activation scores, denoted as $\mathbf{s} \in \mathbb{R}^C$, which quantify the importance or perturbation sensitivity of each channel. The process can be described as follows:

$$\mathbf{s} = \mathcal{G}(\{\mathcal{F}^l, \mathcal{F}^u\}; \theta_G), \quad (4)$$

where θ_G is the learnable parameter of the router. Subsequently, we construct a sparse channel mask $\mathcal{R} \in \mathbb{R}^C$ to indicate the top- K most activated channels, which can be expressed by

$$\mathcal{R} = \delta(\mathbf{s} \geq \tau_K(\mathbf{s})), \quad (5)$$

where $\delta(\cdot)$ denotes the channel-level threshold function, and $\tau_K(\cdot)$ is the threshold corresponding to the K -largest value in \mathbf{s} . Then, we select the features to be perturbed using the sparse channel mask \mathcal{R} , which can be expressed as follows:

$$\mathcal{F}_{sub}^l = \mathcal{F}^l \odot \mathcal{R}^l, \mathcal{F}_{sub}^u = \mathcal{F}^u \odot \mathcal{R}^u, \quad (6)$$

where $\mathcal{F}_{sub}^l \in \mathbb{R}^{K \times h \times w \times d}$ and $\mathcal{F}_{sub}^u \in \mathbb{R}^{K \times h \times w \times d}$ respectively represent the features selected from labeled and unlabeled data. \mathcal{R}^l and \mathcal{R}^u represent the corresponding sparse channel masks, respectively. Note that $\mathcal{F}_{sub}^l \subset \mathcal{F}^l$ and $\mathcal{F}_{sub}^u \subset \mathcal{F}^u$ are subsets of \mathcal{F}^l and \mathcal{F}^u , respectively, obtained through the masks \mathcal{R}^l and \mathcal{R}^u . Finally, we perform channel-wise interaction on the selected features.

Bidirectional Channel-wise Interaction

To explore the semantic relationships between different data streams, we perform bidirectional feature interaction between labeled and unlabeled data, reactivating features to achieve deep information interaction. Firstly, two feature containers $\mathcal{Q}^l \in \mathbb{R}^{M \times L}$ and $\mathcal{Q}^u \in \mathbb{R}^{M \times L}$ are randomly initialized to store the selected features from the labeled and unlabeled data, where M and L denote the maximum length of the containers and the length of a single channel, respectively. Then, we compute the similarity between the selected channels and the features stored in the containers, from which the most similar features to each selected channel are retrieved. This process can be expressed as follows:

$$\begin{cases} \mathcal{F}_q^l = \{\arg \max_{f_q \in \mathcal{Q}^l} \text{Sim}(\mathcal{F}_{sub,k}^l, f_q)\}_{k=1}^K, \\ \mathcal{F}_q^u = \{\arg \max_{f_q \in \mathcal{Q}^u} \text{Sim}(\mathcal{F}_{sub,k}^u, f_q)\}_{k=1}^K, \end{cases} \quad (7)$$

where $\text{Sim}(\cdot, \cdot)$ represents the cosine similarity function. $\mathcal{F}_{sub,k}^l$ and $\mathcal{F}_{sub,k}^u$ represent the k -th channel feature of \mathcal{F}_{sub}^l and \mathcal{F}_{sub}^u , respectively. $\mathcal{F}_q^l \in \mathbb{R}^{K \times h \times w \times d}$ and $\mathcal{F}_q^u \in \mathbb{R}^{K \times h \times w \times d}$ represent the most similar features selected from the containers, respectively. Subsequently, we perform mutual feature perturbation between the labeled and unlabeled data flow, which can be expressed as follows:

$$\begin{cases} \tilde{\mathcal{F}}_{sub}^l = \sigma(\mathbf{Q}(\mathcal{F}_{sub}^l) \cdot \mathbf{K}(\mathcal{F}_q^u)^\top / \sqrt{d}) \cdot \mathbf{V}(\mathcal{F}_q^u) + \mathcal{F}_{sub}^l, \\ \tilde{\mathcal{F}}_{sub}^u = \sigma(\mathbf{Q}(\mathcal{F}_{sub}^u) \cdot \mathbf{K}(\mathcal{F}_q^l)^\top / \sqrt{d}) \cdot \mathbf{V}(\mathcal{F}_q^l) + \mathcal{F}_{sub}^u, \end{cases} \quad (8)$$

where \mathbf{Q} , \mathbf{K} , and \mathbf{V} represent the linear projections. \sqrt{d} is a scaling factor, and $\sigma(\cdot)$ represents the softmax activation function. Finally, the perturbed features are re-inserted into the original features, guided by the sparse channel mask \mathcal{R} , which can be expressed as follows:

$$\begin{cases} \tilde{\mathcal{F}}^l = \tilde{\mathcal{F}}_{sub}^l \odot \mathcal{R}^l + \mathcal{F}^l \odot (1 - \mathcal{R}^l), \\ \tilde{\mathcal{F}}^u = \tilde{\mathcal{F}}_{sub}^u \odot \mathcal{R}^u + \mathcal{F}^u \odot (1 - \mathcal{R}^u), \end{cases} \quad (9)$$

where $\tilde{\mathcal{F}}^l \in \mathbb{R}^{C \times h \times w \times d}$ and $\tilde{\mathcal{F}}^u \in \mathbb{R}^{C \times h \times w \times d}$ represent the features of the labeled and unlabeled data after interaction, respectively. This bidirectional channel-wise interaction strategy enables targeted interaction between the feature distributions of labeled and unlabeled data, with the router learning to modulate only the most relevant semantic features, avoiding unnecessary interference. It is worth noting that \mathcal{Q}^l and \mathcal{Q}^u are feature queues operating on a first-in-first-out basis. When full, they discard the oldest entries. At each iteration, features from labeled data are multiplied by ground-truth class labels while unlabeled data features are multiplied by pseudo-labels, and the obtained features are appended to the respective queues.

Loss Function

The total loss function includes supervised loss, unsupervised loss, and consistency loss, which can be expressed by

$$\mathcal{L}_{total} = \mathcal{L}_{sup} + \mathcal{L}_{cons} + \lambda_u \mathcal{L}_{unsup} \quad (10)$$

where λ_u employs a Gaussian warm-up function $\lambda_u(t) = \beta * e^{-5(1-t/t_{max})^2}$ with $\beta = 0.1$, and t_{max} denotes the maximum number of iterations. To encourage the model to focus more attention on challenging segmentation regions, such as boundaries between different classes, we introduce uncertainty values as weighting guidance (Zhao et al. 2024) in the segmentation loss calculation. In the following losses, we apply the weighted segmentation loss \mathcal{L}_{seg} throughout, which can be expressed as follows:

$$\mathcal{L}_{seg}(p, y, \mathcal{W}) = \mathcal{L}_{BCE} + \mathcal{L}_{IoU}, \quad (11)$$

where p , y , and \mathcal{W} represent the prediction, ground-truth, and uncertainty weight map, respectively.

We apply two different strong augmentations and one weak augmentation to the unsupervised data, and take the predictions from the weak augmentation as pseudo-labels to supervise the strongly augmented data. The unsupervised loss with the weak-to-strong consistency strategy can be represented as follows:

$$\mathcal{L}_{unsup} = \mathcal{L}_{seg}(p_{scol}^u, p_w^u, \mathcal{W}_{scol}^u) + \mathcal{L}_{seg}(p_{smix}^u, p_w^u, \mathcal{W}_{smix}^u), \quad (12)$$

where \mathcal{W}_{scol}^u and \mathcal{W}_{smix}^u represent their corresponding uncertainty weights. Note that for the strong augmentation operation involving spatial perturbation, we perform spatial restoration before calculating the loss. In addition, we perform cognitive consistency learning on the unlabeled data with different strong augmentations, which can be expressed as follows:

$$\mathcal{L}_{cons} = \frac{1}{H \times W} \sum_{i \in H, j \in W} (p_{scol}^u(i, j) - p_{smix}^u(i, j))^2. \quad (13)$$

Method	10% / 8 labeled data				20% / 16 labeled data				10% / 25 labeled data				20% / 50 labeled data			
	Dice ↑	Jaccard ↑	95HD ↓	ASD ↓	Dice ↑	Jaccard ↑	95HD ↓	ASD ↓	Dice ↑	Jaccard ↑	95HD ↓	ASD ↓	Dice ↑	Jaccard ↑	95HD ↓	ASD ↓
VNet (SupOnly)	82.74	71.72	13.35	3.26	84.93	75.87	14.50	4.30	74.43	61.86	37.11	2.79	80.16	71.55	22.68	3.43
UAMT (Yu et al. 2019)	87.79	78.39	8.68	2.12	88.97	80.38	8.18	1.92	79.49	69.22	11.93	1.93	79.72	69.46	11.26	1.97
DTC (Luo et al. 2021)	87.51	78.17	8.23	2.36	88.32	79.34	8.72	2.02	77.54	66.91	12.16	2.84	82.38	72.15	11.00	2.16
MC-Net (Wu et al. 2021)	87.62	78.25	10.03	1.82	90.37	82.52	6.77	1.71	81.52	71.89	13.26	4.56	83.79	73.69	9.65	1.76
MC-Net+ (Wu et al. 2022)	88.89	80.15	8.01	1.90	90.31	82.41	6.61	1.62	79.27	68.97	14.89	4.91	83.47	72.89	9.69	1.92
URPC (Luo et al. 2022)	86.92	77.03	11.13	2.28	87.68	78.36	9.39	3.52	82.59	72.11	13.88	3.72	82.93	72.57	15.93	4.19
MCF (Wang et al. 2023)	87.06	77.83	7.81	2.67	88.71	80.41	6.32	1.90	78.83	68.49	12.25	1.92	80.07	69.55	10.65	2.00
BCP (Bai et al. 2023)	89.62	81.31	6.81	1.76	90.38	82.57	6.68	1.76	78.64	68.59	12.88	2.81	80.20	69.66	10.52	2.08
BS-Net (He et al. 2024)	82.53	71.27	12.48	3.30	90.43	-	6.21	1.63	79.93	69.37	10.88	2.05	82.03	71.87	10.26	1.96
MLRP (Su et al. 2024)	89.86	81.68	6.91	1.85	<u>91.02</u>	<u>83.62</u>	5.78	1.66	84.29	74.74	9.57	2.55	<u>85.47</u>	<u>76.32</u>	<u>7.76</u>	2.00
SKCDF (Zhang et al. 2025b)	89.42	80.46	8.18	2.38	90.65	82.98	6.70	<u>1.61</u>	83.00	73.08	10.28	<u>1.78</u>	84.40	74.78	9.31	<u>1.68</u>
UnCo (Zeng et al. 2025)	<u>90.37</u>	<u>82.54</u>	<u>6.11</u>	<u>1.74</u>	90.91	83.40	<u>5.36</u>	<u>1.61</u>	<u>85.09</u>	<u>75.64</u>	<u>8.63</u>	1.89	85.16	75.04	8.41	1.74
Ours	91.07	83.67	5.57	1.53	91.84	84.98	5.06	1.38	86.17	77.10	8.43	1.37	86.86	77.80	7.62	1.26

Table 1: Results on the LA and BraTS-2019 datasets (Best and second-best results are highlighted in **Bold** and Underline).

Finally, we apply supervised loss to the labeled data. In this process, we also perform two types of strong augmentations and one weak augmentation on the labeled data to enable the model to adapt to this augmentation strategy, which can be expressed as follows:

$$\begin{aligned} \mathcal{L}_{sup} = & \mathcal{L}_{seg}(p_{scol}^l, y^l, \mathcal{W}_{scol}^l) + \mathcal{L}_{seg}(p_{smix}^l, y^l, \mathcal{W}_{smix}^l) \\ & + \mathcal{L}_{seg}(p_w^l, y^l, \mathcal{W}_w^l), \end{aligned} \quad (14)$$

where p_{scol}^l and p_{smix}^l represent the predictions for the strongly augmented labeled data, and \mathcal{W}_{scol}^l and \mathcal{W}_{smix}^l are their corresponding uncertainty weights. p_w^l and \mathcal{W}_w^l denote the predictions and uncertainty weights for the weakly augmented labeled data, respectively. Additionally, y^l represents the ground truth of the labeled data.

Experiments

Experimental Setup

Datasets. We conduct comparison experiments on three widely used datasets. • The Left Atrium (LA) dataset (Xiong et al. 2021) is the benchmark for the 2018 atrial segmentation challenge. It includes 100 3D gadolinium-enhanced magnetic resonance image volumes, all with corresponding labels. Following the same experimental setup (Bai et al. 2023), we divide the dataset into 80 samples for training and 20 samples for testing in our experiments. • The Pancreas-CT dataset (Roth et al. 2015) consists of 82 3D abdominal CT scans with a resolution of 512×512 pixels, and slice thicknesses ranging from 1.5 mm to 2.5 mm. Following the previous setting (Luo et al. 2021), we partition the dataset into 62 training samples and 20 testing samples. • The BraTS-2019 dataset (Bakas 2020) contains multi-modal MRI scans from 335 glioma patients, with four imaging sequences: T1, T1CE, T2, and FLAIR. Following the configuration in (Luo 2020), we use only FLAIR images for tumor segmentation. The dataset is split into 250 training, 25 validation, and 60 testing samples.

Implementation Details. All experiments are performed in a PyTorch 1.8.1 environment with CUDA 11.2, utilizing

an NVIDIA 4090 GPU. To maintain consistency, we retain the same backbone selection as in previous work, opting for VNet (Milletari, Navab, and Ahmadi 2016). We utilize the SGD optimizer with an initial learning rate of 0.01, a momentum of 0.9, and a weight decay of 0.0005. The batch size is set to 4. The total number of iterations is set to $30k$. During the feature-level interaction process, we select the top 64 channels. The maximum length k of the container is set to 2560. Based on previous settings, during the training process, the cropping size for the LA dataset is $112 \times 112 \times 80$, while for the other datasets, the cropping size is $96 \times 96 \times 96$. During the testing phase, no perturbations or augmentation operations are applied.

Evaluation Metrics. All of our comparative experiments utilize four evaluation metrics, namely Dice coefficient, Hausdorff distance (95HD), Intersection over Union (IoU), and Average Surface Distance (ASD).

Comparison With State-of-the-art Methods

Comparison Methods. The proposed method is evaluated against 11 state-of-the-art semi-supervised medical image segmentation methods, including UAMT (Yu et al. 2019), DTC (Luo et al. 2021), MC-Net (Wu et al. 2021), MC-Net+ (Wu et al. 2022), URPC (Luo et al. 2022), MCF (Wang et al. 2023), BCP (Bai et al. 2023), BS-Net (He et al. 2024), MLRP (Su et al. 2024), SKCDF (Zhang et al. 2025b), and UnCo (Zeng et al. 2025).

Quantitative and Qualitative Comparisons. Table 1 presents the quantitative results of our method on the LA and BraTS-2019 datasets, where our semi-supervised approach achieves satisfactory performance. On the BraTS-2019 dataset with a 10% labeled ratio, our method improves the Dice score from 85.09% to 86.17% compared to the second-best model (UnCo), and reduces the 95HD from 1.89 to 1.37. Additionally, when compared to the SKCDF network, which also uses a dual-stream interaction structure, the Dice score increases by 3.17%. Furthermore, it continues to achieve the best results at other labeled ratios. Table 2 shows the comparison of our method with other semi-supervised approaches on the Pancreas dataset. In the case

Method	10% / 6 labeled data				20% / 12 labeled data			
	Dice \uparrow	Jaccard \uparrow	95HD \downarrow	ASD \downarrow	Dice \uparrow	Jaccard \uparrow	95HD \downarrow	ASD \downarrow
VNet	54.94	40.87	47.48	17.43	75.07	61.96	10.79	3.31
UAMT	66.44	52.02	17.04	3.03	76.10	62.62	10.87	2.43
DTC	66.58	51.79	15.46	4.16	76.27	62.82	8.70	2.20
MC-Net	69.07	54.36	14.53	2.28	78.17	65.22	6.90	1.55
MC-Net+	70.00	55.66	16.03	3.87	79.37	66.83	8.52	1.72
URPC	73.53	59.44	22.57	7.85	80.02	67.30	8.51	1.98
MCF	67.71	53.83	17.17	2.34	75.00	61.27	11.59	3.27
BCP	75.17	60.99	10.49	2.32	82.91	70.97	6.43	2.25
BS-Net	64.61	50.02	24.74	5.28	78.93	65.75	8.49	2.20
MLRP	75.93	62.12	9.07	1.54	81.53	69.35	6.81	<u>1.33</u>
SKCDF	70.71	56.48	19.84	1.37	78.30	65.27	9.34	1.38
UnCo	<u>78.53</u>	<u>65.20</u>	<u>7.36</u>	<u>1.33</u>	<u>81.93</u>	<u>69.77</u>	<u>5.34</u>	1.36
Ours	80.41	67.70	6.33	1.32	83.24	71.61	4.87	1.16

Table 2: Quantitative results on the Pancreas-CT dataset.

Baseline	BCSI			Left Atrium (20%)				BraTS-2019 (20%)			
	SSP	BCI	CR	Dice	Jaccard	95HD	ASD	Dice	Jaccard	95HD	ASD
\checkmark				88.60	79.82	10.98	2.80	82.86	73.24	10.31	1.72
\checkmark	\checkmark			90.58	82.96	6.43	1.58	85.59	76.34	8.51	1.42
\checkmark		\checkmark		89.23	80.79	7.50	2.26	84.12	74.61	9.65	1.62
\checkmark			\checkmark	90.22	82.40	6.49	1.85	84.72	75.44	9.04	1.32
\checkmark	\checkmark	\checkmark	\checkmark	91.84	84.98	5.06	1.38	86.86	77.80	7.62	1.26

Table 3: Ablation study on key components in our method.

of 10% labeled data, compared to UnCo, the Dice score increases from 78.53% to 80.41%, and the 95HD decreases from 7.36 to 6.33. Fig. 3 shows the qualitative results of our method and other comparison methods. It can be observed that our method can still accurately capture the target regions, particularly in handling complex structures and blurred boundaries.

Ablation Studies

We conduct ablative studies on the key components and all proposed strategies within the BCSI framework. All experiments are performed on the LA and Pancreas datasets with a 20% labeled ratio.

Effectiveness of Key Components. Table 3 shows the results of our quantitative analysis to explore the effectiveness of the key components in our method. Notably, without the SSP strategy, the weak-to-strong consistency learning framework cannot be applied, and we substitute it with the MT structure. Additionally, applying only the BCI module without the CR module implements a full channel interaction strategy. As shown in Table 3, the SSP structure based on weak-to-strong consistency outperforms traditional semi-supervised structures like MT. Furthermore, applying the BCI module within the MT structure improves performance, validating the effectiveness of the dual-stream interaction strategy. Building on this, the introduction of channel routing further improves performance. On the LA dataset with 20% labeled data, the Dice score increases from 89.23% to 90.22%, and when all strategies are applied, the performance rises to 91.84%.

Settings			Left Atrium (20%)				BraTS-2019 (20%)			
Weak	Col.	Mix.	Dice	Jaccard	IoU	ASD	Dice	Jaccard	95HD	ASD
\checkmark			90.22	82.40	6.49	1.85	84.72	75.44	9.04	1.32
\checkmark	\checkmark		90.87	83.33	6.38	1.70	85.80	76.69	9.18	1.49
\checkmark		\checkmark	90.97	83.50	6.49	1.50	86.16	77.30	8.44	1.27
\checkmark	\checkmark	\checkmark	91.84	84.98	5.06	1.38	86.86	77.80	7.62	1.26

Table 4: Ablation study of different settings in the SSP.

Effectiveness of Different Settings in SSP. We employ an ablation study to quantify the impact of different augmentation settings in our SSP strategy, and Table 4 presents ablative results. Notably, when only weak augmentation is applied, the MT framework is used as a substitute with keeping all other strategies. As shown in Table 4, it can be observed that our SSP with one weak and two strong augmentation strategies achieves the best performance. This indicates our strategy fully leverages changes in the semantic and spatial domains to enhance model performance.

Settings	Left Atrium (20%)				BraTS-2019 (20%)			
	Dice	Jaccard	95HD	ASD	Dice	Jaccard	95HD	ASD
Lab. \times Unlab.	90.58	82.96	6.43	1.58	85.59	76.34	8.51	1.42
Lab. \rightarrow Unlab.	91.42	84.29	5.34	1.50	86.13	77.04	8.17	1.34
Lab. \leftarrow Unlab.	91.39	84.24	5.60	1.55	86.05	76.72	8.30	1.48
Lab. \leftrightarrow Unlab.	91.84	84.98	5.06	1.38	86.86	77.80	7.62	1.26

Table 5: Ablation study of interactive data flow direction.

Effectiveness of Bidirectional Interaction. We investigate the effects of different interaction directions between the labeled and unlabeled data flows. Notably, ‘‘Lab. \rightarrow Unlab.’’ refers to the process of using labeled data to augment unlabeled data, which corresponds to executing only the second half of Eq. (8), with the reverse applying similarly. As shown in Table 5, perturbations from labeled data to unlabeled data perform better than those from unlabeled data to labeled data. This is because the features of labeled data are generally more stable, providing more reliable information to guide the learning process of unlabeled data. More importantly, bidirectional interaction achieves the best performance, by exploiting the potential relationship between labeled and unlabeled data to enhance their complementarity.

Effect of Channel Selection Count. To further investigate the impact of channel selection ratio on model performance, we conduct an ablation study on the channel selection ratio, aiming to ensure that additional information is fully utilized while minimizing unnecessary interference. Note that we only modify the number of selected channels (*i.e.*, Top- K), while the other strategies (*i.e.*, CR and SSP) remain unchanged. As shown in Table 6, full-channel interaction ($K = 256$) yields the lowest performance, with a Dice score of 91.27% on the LA dataset using 20% labeled data. This degradation occurs because the full-channel approach introduces excessive noise and redundant information, which distracts the model from key features. While reducing K to 32 improves the Dice score to 91.66%, the gain remains limited. It can be observed that we can achieve the

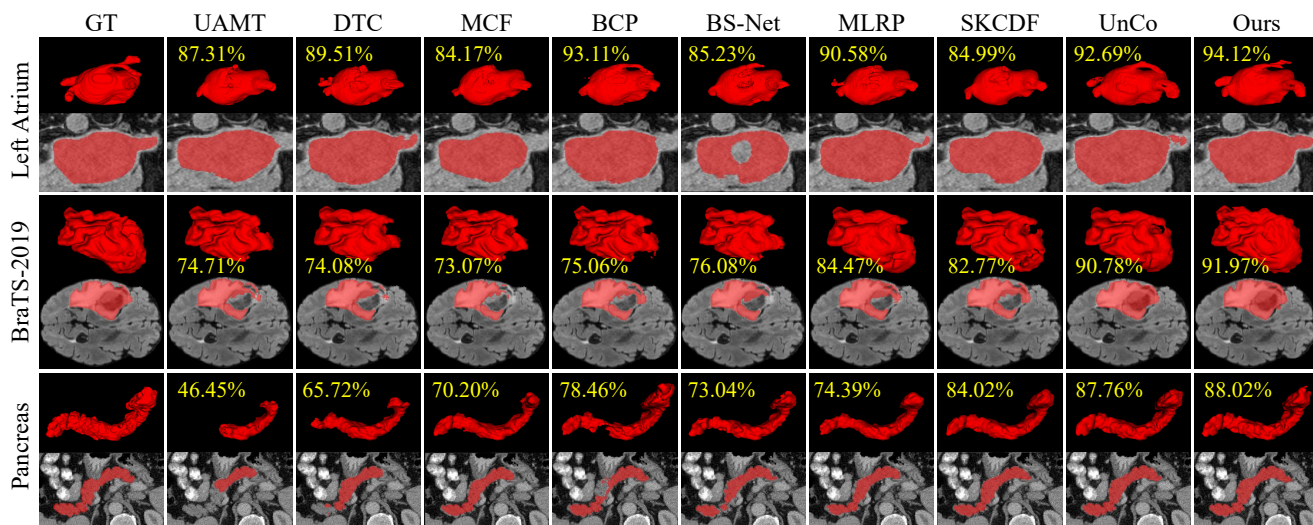


Figure 3: Visualization results of our method compared to other semi-supervised methods on the LA, BraTS-2019, and Pancreas datasets. The yellow numbers represent the Dice score of the currently displayed sample.

Settings	Left Atrium (20%)				BraTS-2019 (20%)			
	Dice	Jaccard	95HD	ASD	Dice	Jaccard	95HD	ASD
$K = 32$	91.66	84.68	5.88	1.43	86.46	77.36	8.00	1.22
$K = 64$	91.84	84.98	5.06	1.38	86.86	77.80	7.62	1.26
$K = 128$	91.36	84.22	5.80	1.45	86.31	77.47	8.07	1.23
$K = 256$	91.27	84.07	5.41	1.50	86.14	77.16	7.91	1.23

Table 6: Ablation study on channel selection count.

Settings	Left Atrium (20%)				BraTS-2019 (20%)			
	Dice	Jaccard	95HD	ASD	Dice	Jaccard	95HD	ASD
Random	90.67	83.11	6.03	1.60	85.88	76.92	8.17	1.30
Router	91.84	84.98	5.06	1.38	86.86	77.80	7.62	1.26

Table 7: Ablation study on the channel selection mechanism.

best performance when $K = 64$.

Effect of Channel Selection Mechanism. We further conduct an ablation study to investigate the impact of the channel selection mechanism on model performance, comparing the router method with the random channel selection method. This study aims to analyze the advantages of the router, which dynamically selects key features in a learnable manner, compared to the random method. Note that the ablation study is conducted based on $K = 64$. Table 7 shows the router mechanism outperforming random selection by 1.17% (LA) and 0.98% (BraTS-2019) in Dice score with 20% labeled data. This performance gain confirms the router’s utility in mitigating redundant information interference and improving model robustness.

In-depth Comparison and Discussion. To more comprehensively evaluate the effectiveness of our method, we compare it with foundation model-based semi-supervised methods (VCLIP (Li et al. 2024) and SFR (Li et al. 2025)) and the




Figure 4: Performance comparison of our method with foundation model-based semi-supervised approaches and the fully supervised VNet, using a 20% labeled data ratio.

fully-supervised VNet. The results in Fig. 4 demonstrate that our method achieves performance comparable to the fully-supervised paradigm. More significantly, on the BraTS-2019 dataset with only 20% labeled data, our method surpasses the fully-supervised VNet in the 95HD metric, underscoring its exceptional precision.

Conclusion

We present BCSI, a novel semi-supervised paradigm that enhances interactions between labeled and unlabeled data. Our framework incorporates a Semantic-Spatial Perturbation (SSP) strategy, which enriches data diversity via strong augmentations while leveraging predictions from weak augmentations as pseudo-labels. Moreover, we propose a Channel-selective Router (CR) and a Bidirectional Channel-wise Interaction (BCI) module to dynamically select informative feature channels and enable bidirectional exchange between labeled and unlabeled data streams. Experimental results demonstrate that BCSI outperforms existing state-of-the-art semi-supervised medical image segmentation methods.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (Nos. 62576153, 62172228, 62201263, and 62476054).

References

- Bai, Y.; Chen, D.; Li, Q.; Shen, W.; and Wang, Y. 2023. Bidirectional copy-paste for semi-supervised medical image segmentation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 11514–11524.
- Bakas, S. S. 2020. Brats MICCAI Brain tumor dataset.
- Bortsova, G.; Dubost, F.; Hogeweg, L.; Katramados, I.; and De Bruijne, M. 2019. Semi-supervised medical image segmentation via learning consistency under transformations. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 810–818. Springer.
- Cao, H.; Wang, Y.; Chen, J.; Jiang, D.; Zhang, X.; Tian, Q.; and Wang, M. 2022. Swin-unet: Unet-like pure transformer for medical image segmentation. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 205–218. Springer.
- Chen, X.; Yuan, Y.; Zeng, G.; and Wang, J. 2021. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2613–2622.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 702–703.
- Han, K.; Liu, L.; Song, Y.; Liu, Y.; Qiu, C.; Tang, Y.; Teng, Q.; and Liu, Z. 2022. An effective semi-supervised approach for liver CT image segmentation. *IEEE Journal of Biomedical and Health Informatics*, 26(8): 3999–4007.
- He, A.; Li, T.; Yan, J.; Wang, K.; and Fu, H. 2024. Bilateral Supervision Network for Semi-supervised Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 43(5): 1715–1726.
- Huang, K.; Zhou, T.; Fu, H.; Zhang, Y.; Zhou, Y.; Gong, C.; and Liang, D. 2025a. Learnable prompting sam-induced knowledge distillation for semi-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 44(5): 2295–2306.
- Huang, K.; Zhou, T.; Fu, H.; Zhang, Y.; Zhou, Y.; and Wu, X.-J. 2025b. Uncertainty-aware Cross-training for Semi-supervised Medical Image Segmentation. *IEEE Transactions on Image Processing*, 34: 5543–5556.
- Huang, K.; Zhou, Y.; Fu, H.; Zhang, Y.; Gong, C.; and Zhou, T. 2025c. Text-Driven Multiplanar Visual Interaction for Semi-supervised Medical Image Segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 604–613. Springer.
- Huang, W.; Chen, C.; Xiong, Z.; Zhang, Y.; Chen, X.; Sun, X.; and Wu, F. 2022. Semi-supervised neuron segmentation via reinforced consistency learning. *IEEE Transactions on Medical Imaging*, 41(11): 3016–3028.
- Jiao, R.; Zhang, Y.; Ding, L.; Xue, B.; Zhang, J.; Cai, R.; and Jin, C. 2024. Learning with limited annotations: a survey on deep semi-supervised learning for medical image segmentation. *Computers in Biology and Medicine*, 169: 107840.
- Kim, J.; Min, Y.; Kim, D.; Lee, G.; Seo, J.; Ryoo, K.; and Kim, S. 2022. Conmatch: Semi-supervised learning with confidence-guided consistency regularization. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 674–690. Springer.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proc. Int. Conf. Comput. Vis. (ICCV)*, 4015–4026.
- Li, L.; Lian, S.; Luo, Z.; Wang, B.; and Li, S. 2024. VCLIPSeg: Voxel-Wise CLIP-Enhanced Model for Semi-supervised Medical Image Segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 692–701. Springer.
- Li, N.; Xiong, L.; Qiu, W.; Pan, Y.; Luo, Y.; and Zhang, Y. 2023. Segment anything model for semi-supervised medical image segmentation via selecting reliable pseudo-labels. In *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 138–149. Springer.
- Li, S.; Qi, L.; Yu, Q.; Huo, J.; Shi, Y.; and Gao, Y. 2025. Stitching, fine-tuning, re-training: A sam-enabled framework for semi-supervised 3d medical image segmentation. *IEEE Transactions on Medical Imaging*, 44(10): 3909–3923.
- Li, Y.; Luo, L.; Lin, H.; Chen, H.; and Heng, P.-A. 2021. Dual-consistency semi-supervised learning with uncertainty quantification for COVID-19 lesion segmentation from CT images. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 199–209. Springer.
- Luo, X. 2020. SSL4MIS. <https://github.com/HiLab-git/SSL4MIS>.
- Luo, X.; Chen, J.; Song, T.; and Wang, G. 2021. Semi-supervised medical image segmentation through dual-task consistency. In *Proc. AAAI Conf. Artif. Intell.*, volume 35, 8801–8809.
- Luo, X.; Wang, G.; Liao, W.; Chen, J.; Song, T.; Chen, Y.; Zhang, S.; Metaxas, D. N.; and Zhang, S. 2022. Semi-supervised medical image segmentation via uncertainty rectified pyramid consistency. *Medical Image Analysis*, 80: 102517.
- Mei, J.; Zhou, T.; Huang, K.; Zhang, Y.; Zhou, Y.; Wu, Y.; and Fu, H. 2025. A survey on deep learning for polyp segmentation: Techniques, challenges and future trends. *Visual Intelligence*, 3(1): 1.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 565–571. IEEE.
- Peiris, H.; Hayat, M.; Chen, Z.; Egan, G.; and Harandi, M. 2023. Uncertainty-guided dual-views for semi-supervised volumetric medical image segmentation. *Nature Machine Intelligence*, 5(7): 724–738.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.;

- et al. 2021. Learning transferable visual models from natural language supervision. In *Proc. Int. Conf. Mach. Learn. (ICML)*, 8748–8763. PMLR.
- Roth, H. R.; Lu, L.; Farag, A.; Shin, H.-C.; Liu, J.; Turkbey, E. B.; and Summers, R. M. 2015. Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 556–564. Springer.
- Seibold, C. M.; Reiß, S.; Kleesiek, J.; and Stiefelhagen, R. 2022. Reference-guided pseudo-label generation for medical semantic segmentation. In *Proc. AAAI Conf. Artif. Intell.*, volume 36, 2171–2179.
- Shi, Y.; Zhang, J.; Ling, T.; Lu, J.; Zheng, Y.; Yu, Q.; Qi, L.; and Gao, Y. 2021. Inconsistency-aware uncertainty estimation for semi-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 41(3): 608–620.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 33: 596–608.
- Su, J.; Luo, Z.; Lian, S.; Lin, D.; and Li, S. 2024. Mutual learning with reliable pseudo label for semi-supervised medical image segmentation. *Medical Image Analysis*, 94: 103111.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 30.
- Wang, H.; Zhang, Q.; Li, Y.; and Li, X. 2024. Allspark: Reborn labeled features from unlabeled in transformer for semi-supervised semantic segmentation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3627–3636.
- Wang, K.; Zhan, B.; Zu, C.; Wu, X.; Zhou, J.; Zhou, L.; and Wang, Y. 2021. Triple-uncertainty guided mean teacher model for semi-supervised medical image segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 450–460. Springer.
- Wang, Y.; Xiao, B.; Bi, X.; Li, W.; and Gao, X. 2023. Mcf: Mutual correction framework for semi-supervised medical image segmentation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 15651–15660.
- Wu, Y.; Ge, Z.; Zhang, D.; Xu, M.; Zhang, L.; Xia, Y.; and Cai, J. 2022. Mutual consistency learning for semi-supervised medical image segmentation. *Medical Image Analysis*, 81: 102530.
- Wu, Y.; Xu, M.; Ge, Z.; Cai, J.; and Zhang, L. 2021. Semi-supervised left atrium segmentation with mutual consistency training. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 297–306. Springer.
- Xiong, Z.; Xia, Q.; Hu, Z.; Huang, N.; Bian, C.; Zheng, Y.; Vesal, S.; Ravikumar, N.; Maier, A.; Yang, X.; et al. 2021. A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *Medical Image Analysis*, 67: 101832.
- Xu, M.-C.; Zhou, Y.-K.; Jin, C.; Blumberg, S. B.; Wilson, F. J.; Degroot, M.; Alexander, D. C.; Oxtoby, N. P.; and Jacob, J. 2022. Learning morphological feature perturbations for calibrated semi-supervised segmentation. In *MIDL*, 1413–1429. PMLR.
- Xu, X.; Sanford, T.; Turkbey, B.; Xu, S.; Wood, B. J.; and Yan, P. 2021. Shadow-consistent semi-supervised learning for prostate ultrasound segmentation. *IEEE Transactions on Medical Imaging*, 41(6): 1331–1345.
- Yang, L.; Qi, L.; Feng, L.; Zhang, W.; and Shi, Y. 2023. Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 7236–7246.
- Yao, H.; Hu, X.; and Li, X. 2022. Enhancing pseudo label quality for semi-supervised domain-generalized medical image segmentation. In *Proc. AAAI Conf. Artif. Intell.*, volume 36, 3099–3107.
- Yu, L.; Wang, S.; Li, X.; Fu, C.-W.; and Heng, P.-A. 2019. Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 605–613. Springer.
- Zeng, L.-L.; Gao, K.; Hu, D.; Feng, Z.; Hou, C.; Rong, P.; and Wang, W. 2023. SS-TBN: A semi-supervised tri-branch network for COVID-19 screening and lesion segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 10427–10442.
- Zeng, X.; Xiong, S.; Xu, J.; Du, G.; and Rong, Y. 2025. Uncertainty Co-estimator for Improving Semi-Supervised Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 44(9): 3870–3881.
- Zhang, Q.; Wang, H.; and Li, X. 2025. S&D Messenger: Exchanging Semantic and Domain Knowledge for Generic Semi-Supervised Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 44(11): 4487–4498.
- Zhang, Y.; Lv, B.; Xue, L.; Zhang, W.; Liu, Y.; Fu, Y.; Cheng, Y.; and Qi, Y. 2025a. SemiSAM+: Rethinking Semi-Supervised Medical Image Segmentation in the Era of Foundation Models. *arXiv preprint arXiv:2502.20749*.
- Zhang, Z.; Yin, G.; Zhang, B.; Liu, W.; Zhou, X.; and Wang, W. 2025b. A Semantic Knowledge Complementarity based Decoupling Framework for Semi-supervised Class-imbalanced Medical Image Segmentation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 25940–25949.
- Zhao, Y.; Zhou, Y.; Zhang, Y.; Wu, Y.; and Zhou, T. 2024. TextPolyp: Point-Supervised Polyp Segmentation with Text Cues. In *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, 711–722. Springer.
- Zhou, H.-Y.; Guo, J.; Zhang, Y.; Han, X.; Yu, L.; Wang, L.; and Yu, Y. 2023. nnFormer: volumetric medical image segmentation via a 3D transformer. *IEEE Transactions on Image Processing*, 32: 4036–4045.