

# DualFete: Revisiting Teacher-Student Interactions from a Feedback Perspective for Semi-supervised Medical Image Segmentation

Le Yi<sup>1</sup>, Wei Huang<sup>1\*</sup>, Lei Zhang<sup>1\*</sup>, Kefu Zhao<sup>1</sup>, Yan Wang<sup>2</sup>, Zizhou Wang<sup>2</sup>

<sup>1</sup>College of Computer Science, Sichuan University, China

<sup>2</sup>Institute of High Performance Computing, A\*STAR, Singapore

{yile, kefuzhao}@stu.scu.edu.cn, {weihuang, leizhang}@scu.edu.cn, {wangyan, wang\_zizhou}@a-star.edu.sg

## Abstract

The teacher-student paradigm has emerged as a canonical framework in semi-supervised learning. When applied to medical image segmentation, the paradigm faces challenges due to inherent image ambiguities, making it particularly vulnerable to erroneous supervision. Crucially, the student’s iterative reconfirmation of these errors leads to self-reinforcing bias. While some studies attempt to mitigate this bias, they often rely on external modifications to the conventional teacher-student framework, overlooking its intrinsic potential for error correction. In response, this work introduces a feedback mechanism into the teacher-student framework to counteract error reconfirmations. Here, the student provides feedback on the changes induced by the teacher’s pseudo-labels, enabling the teacher to refine these labels accordingly. We specify that this interaction hinges on two key components: the feedback attributor, which designates pseudo-labels triggering the student’s update, and the feedback receiver, which determines where to apply this feedback. Building on this, a dual-teacher feedback model is further proposed, which allows more dynamics in the feedback loop and fosters more gains by resolving disagreements through cross-teacher supervision while avoiding consistent errors. Comprehensive evaluations on three medical image benchmarks demonstrate the method’s effectiveness in addressing error propagation in semi-supervised medical image segmentation.

## Introduction

Medical image segmentation, which provides quantitative profiles for inner-body anatomical structures, plays a vital role in clinical practice and has emerged as a rapidly evolving subfield of AI for medicine (Luo et al. 2022a; Chen et al. 2024; Li et al. 2025; Lan et al. 2025a,b). However, annotating medical images requires specialized expertise and is particularly labor-intensive at the voxel level. As a result, segmentation models often suffer from performance degradation due to the limited availability of labeled training data.

Semi-supervised learning (SSL) overcomes this limitation by leveraging additional supervision from unlabeled data. One line of SSL work follows the *smoothness* assumption, which posits that if two samples are close in data space, then so should be their corresponding outputs (Chapelle et al.

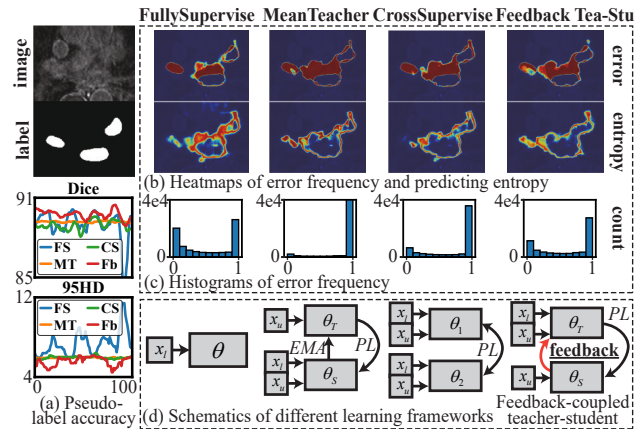


Figure 1: Pre-experiments on the LA dataset. 16 labeled samples are used. FullySupervise runs with only the labeled set, while others use both. Results are evaluated on the first unlabeled sample and averaged from the last 100 training steps. It can be seen that medical image segmentation is susceptible to continual errors, so the confirmation bias issue is fairly problematic. Consistent errors can be reduced by the feedback interaction [highlighted in (d)].

2009). Consistency constraints, thus, are imposed on models to match the prediction of a perturbed sample to that of its vanilla counterpart (Luo et al. 2021; Huang et al. 2025). Another line follows the *clustering* assumption, emphasizing that samples should reside in high-density regions. It motivates pseudo-labeling unlabeled data for entropy minimization (Grandvalet and Bengio 2004); ultimately, samples can be compactly clustered by classes. Notably, accurate supervision is crucial for both lines of work, as models tend to perpetuate historical errors without awareness of their own mistakes, thereby rendering these errors increasingly difficult to correct. This degenerating case is known as *confirmation bias* (Ke et al. 2019).

This bias, unfortunately, incurs more hazards in semi-supervised medical image segmentation (SSMIS) due to image ambiguity; it leads to high regional uncertainty, especially near object boundaries. [Figure 1(b)] (Yi et al. 2025). For unseen samples, models are prone to generate erroneous pseudo-labels. Paradoxically, when trained on such

\*Corresponding author: Wei Huang; Lei Zhang.

error-prone supervision, they exhibit overconfidence in mistakes (low-entropy regions yet with high error frequency). Some studies attempt to prevent models from fitting to unreliable supervision, thereby partially alleviating error reinforcement (Bai et al. 2023; Wu et al. 2023; Wang et al. 2023; Shen et al. 2023; Chi et al. 2024). However, these heuristic-based methods remain insufficient to counteract the intrinsic error-accumulation tendency of the prevailing teacher-student paradigm. This limitation stems from two fundamental issues. First, most methods rely on model-level perturbations to introduce teacher-student discrepancies and promote representation learning by resolving disagreement. However, such strategies are inadequate for highly nonlinear networks, as they tend to degenerate into self-training behavior, failing to sustain meaningful disagreement. Second, and more critically, the conventional teacher-student paradigm lacks inherent mechanisms for error correction. Although pseudo-labels can be influenced by the student using methods like Mean Teacher (Tarvainen and Valpola 2017) or cross supervision (Chen et al. 2021)[Figure 1(d)], the student has no way to verify whether the update induced by pseudo-labels remains aligned with the constraint induced by labeled data. Once convergence, such methods yield vast consistent mistakes [Figure 1(c)] and also derive negligible change in pseudo-labeling accuracy [Figure 1(a)]. Thus, during the degenerative process toward self-training, they continually accumulate errors, exacerbating miscalibrated predictions and intensifying confirmation bias.

In this work, we incorporate a feedback mechanism into the teacher-student framework, wherein the student assesses whether the updates driven by pseudo-labels align with the direction implied by ground-truth supervision. This assessment is then fed back to the teacher to guide refinements of pseudo-labels. Through this interaction, the framework acquires inherent error-correcting capabilities, helping to prevent error accumulation. As shown in Figure 1(c), this method effectively reduces recurring mistakes. To operationalize this feedback loop, we introduce two key components: the feedback attributor, which identifies the pseudo-labels responsible for triggering the student’s update, and the feedback receiver, which determines the pseudo-labeling probabilities to be adjusted. Building on this foundation, we propose a dual-teacher feedback model, where two teachers collaboratively instruct a student and receive individualized feedback from the student, while also improving each other by resolving their mutual disagreements. This dual-teacher framework not only strengthens the feedback dynamics to mitigate persistent errors but also encourages constructive disagreement, fostering more effective teaching curricula. Extensive experiments are conducted on three benchmarks, including LA (Xiong et al. 2021), Pancreas (Roth et al. 2015), and BraTS (Menze et al. 2014), with sufficient justification of our method in these error-prone scenarios.

Contributions are summarized as follows:

- (1) We introduce a feedback mechanism that equips the teacher-student model with error-correction capabilities, fundamentally mitigating vulnerability raised by error propagation in semi-supervised medical image segmentation.
- (2) We develop a dual-teacher feedback model that coor-

dinates: (i) cross-supervision for resolving conflicts, and (ii) feedback loss for preserving teachers’ diversity while suppressing errors. The two aspects enhance feedback dynamics and work collaboratively to teach a superior student.

- (3) Comprehensive experiments on three benchmarks confirm our method’s effectiveness in combating this error-prone scenario, with sufficient justification through detailed quantitative and qualitative analysis.

## Related Work

### Semi-supervised medical image segmentation

Advanced models are marching fastly and have made great breakthroughs in real-world segmentation tasks (Dosovitskiy et al. 2020; Hu et al. 2023). Such advancements, however, are lagging in the medical field mainly due to the scarcity of annotations. As a solution, researchers turn to SSL to excavate gains from limited labeled data and massive unlabeled data. Methods are categorized into two lines: consistency regularization (Chapelle et al. 2009) and entropy minimization (Grandvalet and Bengio 2004). The former argues the output consistency between a perturbed data and its vanilla input. Such perturbations can be conducted at input (Bai et al. 2023; Chi et al. 2024), feature (Yang et al. 2023; Huang et al. 2024), and model (Luo et al. 2022b) levels. The latter argues for separating samples into clusters, driving to minimize entropy via pseudo labels. Representative methods, such as cross supervision (Chen et al. 2021) and Co-training (Qiao et al. 2018), employ multiple learnable models to facilitate mutual learning.

Notably, medical image ambiguity leads to error-prone supervision and inconsistent model predictions. While prior work has shown that models can benefit from learning courses designed to reduce ambiguity (Xu et al. 2023) and disagreement (Shen et al. 2023; Wang et al. 2023), they inherently lack the ability to handle erroneous pseudo-labels in the teacher-student paradigm. Without explicit error-correction, this paradigm tends to accumulate more irreversible mistakes. Some studies address prediction disagreements while improving supervision accuracy, either by filtering unreliable targets using confidence thresholds (Shen et al. 2023; Huang et al. 2025) and predictive uncertainty (Yu et al. 2019; Shi et al. 2021), or by competitively generating pseudo-labels from multiple models (Wu et al. 2023; Wang et al. 2023; Su et al. 2024). However, these solutions introduce additional architectural constraints that often fail to maintain meaningful disagreement due to the models’ tendency to degenerate into trivial self-training behavior.

This paper introduces a feedback mechanism within the teacher-student framework, endowing it with intrinsic error correction capabilities. With this foundation, we propose a dual-teacher feedback model that not only prevents error accumulation but actively benefits from disagreements.

### Confirmation bias in semi-supervised learning

In SSL, confirmation bias (Nickerson 1998) creates an error-reinforcing cycle where the model grows increasingly overconfident in its mistakes and resistant to correction (Arazo et al. 2020). Rollwage et al. (2020) claim that metacognitive

interventions are one kind of method to combat the high confidence in the shaping process of confirmation bias in psychology. Such findings hint at least two routes to address this bias in SSL: (1) confidence reduction, which disrupts the reinforcing cycle of high-confidence errors, and (2) active interventions, which devise strategies to regularize the model toward expected learning direction. Arazo et al. (2020) and Liu et al. (2022a) supervise the model using soft pseudo-labels yielded by mixup techniques (Zhang et al. 2018); they argue for the alleviation of high-confidence mistakes. Some studies impose a minimum labeled-data ratio, either per mini-batch (Arazo et al. 2020) or per image (Bai et al. 2023), to avoid the distributional shift far from the labeled data. Likewise, Shen et al. (2023) and Chi et al. (2024) substitute uncertain patches with more reliable ones, and thus avoid accumulating too many errors. Some studies also propose competitive methods to ensure the correctness of pseudo-labels (Wu et al. 2023; Wang et al. 2023).

While the above heuristically-driven solutions mitigate confirmation bias to some extent, the feedback interaction proposed in this work can address this issue more straightforwardly, as it inherently acts like a metacognitive intervention. Moreover, in a dual-teacher framework, we can enhance this feedback to be more dynamic.

## Methodology

Let  $\theta$  be a segmentation model; it inputs an image  $x$  and predicts probabilities  $f^\theta := f(x; \theta)$ . This model is expected to generalize well with a labeled dataset  $\mathcal{D}_l = \{(x^l, y^l)\}^{N_l}$  and an unlabeled dataset  $\mathcal{D}_u = \{x^u\}^{N_u}$ , where  $x^l, x^u$  represents the labeled and unlabeled sample, respectively,  $y^l$  is the label mask of  $x^l$ , each voxel being one of  $C := \{0, \dots, C-1\}$  classes.  $N_l = |\mathcal{D}_l|$ ,  $N_u = |\mathcal{D}_u|$  is the respective dataset size, and it satisfies  $N_l \ll N_u$ .

The teacher-student model is a learning paradigm in the sense that a teacher  $\theta_T$  pseudo-labels unlabeled data  $x^u$ , which is used to supervise the student  $\theta_S$  such that segmentation loss on  $\mathcal{D}_u$  is minimized in terms of pseudo-labels  $\hat{y}^u$ :

$$\hat{y}^u = \arg \max_{c \in C} f_c(x^u; \theta_T) \quad (1)$$

$$\min_{\theta_S} \mathcal{L}_S(\theta_S, \theta_T; \mathcal{D}_u) = \min_{\theta_S} \frac{1}{N_u} \sum_i \ell(f(x_i^u; \theta_S), \hat{y}_i^u) \quad (2)$$

$f_c^\theta = f_c(\cdot; \theta)$  is the  $c$ -th class prediction, and  $\ell$  is the loss function. Due to the ambiguity, medical image segmentation is susceptible to error supervision. However, the current teacher-student model can not perceive whether the update induced by pseudo-labels would further cause error reinforcement. To this end, we propose **DualFete**, which employs a feedback mechanism to equip the teacher-student model with error-correcting abilities. This mechanism will be further enhanced in a dual-teacher framework.

### Feedback-coupled teacher-student model

To avoid error accumulation, current SSMIS work sets a minimum labeled-data ratio per batch and employs some heuristically-driven methods to regularize training towards

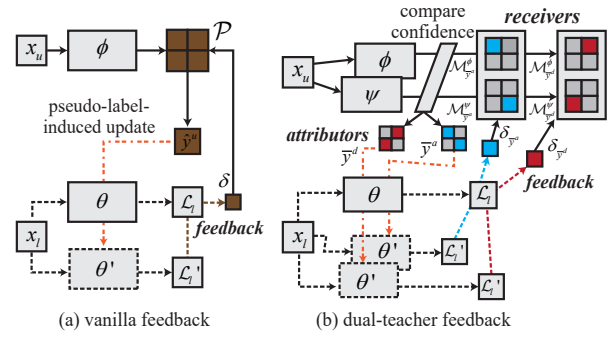


Figure 2: Schematic of the feedback mechanism. (a) Feedback is applied to each unit’s likelihood, leading to a uniform updating direction. (b) The dual-teacher framework enables the feedback more dynamic based on prediction confidence.

the accurate direction (Yu et al. 2019; Luo et al. 2022b). Nonetheless, there is no way to verify if the updates induced by unlabeled data are aligned with those by labeled data.

Motivated by this fact, performance change of the student attributed to the update via pseudo labels is quantified on labeled data. Intriguingly, if performance improved, pseudo labels yielded by the teacher are favorable, prompting further enhancement of pseudo-labeling; otherwise, teacher should reduce the likelihood of yielding such labels. This creates a feedback mechanism within the teacher-student model. In formal, given  $\theta_S$ , the performance on a labeled mini-batch  $\mathcal{D}'_l \subset \mathcal{D}_l$  (with  $N'_l = |\mathcal{D}'_l|$ ) is evaluated by

$$\mathcal{L}_l(\theta_S; \mathcal{D}'_l) = \frac{1}{N'_l} \sum_i \ell(f(x_i^l; \theta_S), y_i^l). \quad (3)$$

Let  $\theta'_S = \theta_S - \eta \Delta \theta_S$  represents the student with one-step update, where  $\Delta \theta_S = \nabla_{\theta_S} \mathcal{L}_S(\theta_S, \theta_T; \mathcal{D}'_u)$  is the stochastic gradient generated on an unlabeled mini-batch  $\mathcal{D}'_u \subset \mathcal{D}_u$ , and  $\eta$  is the step size. Then, feedback can be derived:

$$\delta = \mathcal{L}_l(\theta_S) - \mathcal{L}_l(\theta'_S). \quad (4)$$

We drop  $\mathcal{D}'_l$  from  $\mathcal{L}_l$  for brevity. Let  $\mathcal{P}(\hat{y}^u | x^u; \theta_T, \mathcal{D}'_u)$  represent the likelihood of the teacher generating pseudo-label as  $\hat{y}^u$  on  $\mathcal{D}'_u$ . The teacher minimizes the following feedback loss  $\mathcal{L}_{fb}$  to refine pseudo labels:

$$\mathcal{L}_{fb}(\theta_T; \mathcal{D}'_u) = -\delta \log \mathcal{P}(\hat{y}^u | x^u; \theta_T, \mathcal{D}'_u). \quad (5)$$

Correspondingly, when  $\delta > 0$ , the teacher is guided to maximize the likelihood  $\mathcal{P}$ ; while, when  $\delta < 0$ , the teacher is guided to reduce the likelihood.

Theoretically, the feedback  $\delta$  is a first-order approximation to the inner product of two gradients, i.e.,  $\Delta \theta_S$  and  $\nabla_{\theta_S} \mathcal{L}_l(\theta'_S)$ , contained by a meta-objective (Pham et al. 2021). This means that when  $\delta > 0$ , the update induced by pseudo-labels aligns with the gradient direction implied by the supervised loss  $\mathcal{L}_l$ ; otherwise, the two gradients are in opposite directions. This meta-objective is in line with metacognitive intervention to address confirmation bias in psychology (Rollwage et al. 2020). With this feedback mechanism, the teacher-student framework is equipped with inherent error-correcting abilities such that error accumulation can be circumvented.

## Dual-teacher feedback

While the feedback loss  $\mathcal{L}_{fb}$  provides signals for adjusting pseudo-labels  $\hat{y}^u$ , it enforces a uniform update direction across all voxel predictions, potentially limiting its corrective capacity [Figure 2(a)]. To address this, we propose an enhanced feedback mechanism within a dual-teacher framework, which not only enriches feedback dynamics but also promotes mutual learning between teachers [Figure 2(b)]. To this end, two key components are identified from  $\mathcal{L}_{fb}$ , namely the feedback attributor and receiver. Conceptually, the attributor refers to a set of pseudo-labels with which the student yields gradient for update, while the receiver indicates to which region feedback is applied to modulate the pseudo-labeling likelihood. In the following, the two components are specified in terms of the dual-teacher model.

Let  $\phi, \psi$  denote the two teachers, and  $\hat{y}^{\phi_u}, \hat{y}^{\psi_u}$  be their predicted label by Eq. 1. Pseudo-labels are set to the consensus if agreement is achieved, and set to the higher-confidence label if predictions conflict:

$$\hat{y}^u = \begin{cases} \hat{y}^{\phi_u} \text{ or equivalently } \hat{y}^{\psi_u}, & \text{if } \hat{y}^{\phi_u} = \hat{y}^{\psi_u}; \\ \arg \max_c \max_{\phi, \psi} \{f_c^\phi, f_c^\psi\}, & \text{if } \hat{y}^{\phi_u} \neq \hat{y}^{\psi_u}. \end{cases} \quad (6)$$

This method is more likely to yield accurate pseudo-labels. Then, two types of feedback are quantified, induced by pseudo-labels of agreement and disagreement regions. We use  $\bar{y}^a = \{\hat{y}^u | \hat{y}^{\phi_u} = \hat{y}^{\psi_u}\}$ ,  $\bar{y}^d = \{\hat{y}^u | \hat{y}^{\phi_u} \neq \hat{y}^{\psi_u}\}$  to denote the attributor in terms of the agreement and disagreement region, respectively. Based on Eq. 4, we have

$$\delta_{\bar{y}} = \mathcal{L}_l(\theta_S) - \mathcal{L}_l\left(\theta_S - \eta \frac{\Delta^{\bar{y}} \theta_S}{\|\Delta^{\bar{y}} \theta_S\|}\right). \quad (7)$$

Here,  $\bar{y} \in \{\bar{y}^a, \bar{y}^d\}$ ,  $\Delta^{\bar{y}} \theta_S$  is the gradient of  $\mathcal{L}_S$  w.r.t.  $\theta_S$  induced by pseudo-labels  $\bar{y}$ . Without conflicts, we employ  $\delta_a, \delta_d$ , respectively, to simplify the notations  $\delta_{\bar{y}^a}, \delta_{\bar{y}^d}$ .

In our design, the agreement feedback  $\delta_a$  is applied to the lower-confidence side between  $\phi$  and  $\psi$ . This is because when  $\delta_a > 0$ , we would like to maximize the confidence lower-bound of  $\hat{y}^u$ ; when  $\delta_a < 0$ , the lower confidence favors the easier occurrence of disagreement between two teachers. On the contrary, the disagreement feedback  $\delta_d$  is applied to the higher-confidence side. When  $\delta_d > 0$ , the pseudo-label  $\hat{y}^u$  should become more confident, while when  $\delta_d < 0$ ,  $\hat{y}^u$  is more likely to flip to the label predicted by the other teacher. With this idea, feedback receivers can be identified. Given a teacher  $\theta \in \{\phi, \psi\}$ , and let  $\bar{\theta}$  be the other, the receiver mask is defined at first by  $\mathcal{M}_{\bar{y}}^\theta := \mathcal{I} \left[ \hat{y}^{\theta_u} \star \hat{y}^{\bar{\theta}_u}, f_{\hat{y}^{\theta_u}}^\theta \diamond f_{\hat{y}^{\bar{\theta}_u}}^{\bar{\theta}} \right]$ .  $\mathcal{I}[\cdot]$  is the indicator function.  $\star, \diamond$  are two comparative operators depended on  $\bar{y}$ ; they are respectively replaced by  $=, <$  for  $\bar{y}^a$ , while by  $\neq, >$  for  $\bar{y}^d$ . We can then define the receiver in terms of the feedback  $\delta_{\bar{y}}$  as  $\mathcal{P}(\hat{y}^{\theta_u} | x^u; \theta, \mathcal{D}_u, \mathcal{M}_{\bar{y}}^\theta)$ , i.e., the pseudo-labeling likelihood of  $\theta$  on  $\mathcal{D}_u$  after masking regions by  $\mathcal{M}_{\bar{y}}^\theta$ . As a result, the dual-teacher feedback loss is formulated as

$$\mathcal{L}_{df}(\theta) = - \sum_{\bar{y} \in \{\bar{y}^a, \bar{y}^d\}} \delta_{\bar{y}} \log \mathcal{P}(\hat{y}^{\theta_u} | x^u; \theta, \mathcal{D}_u, \mathcal{M}_{\bar{y}}^\theta). \quad (8)$$

This design gives rise to two distinct yet coordinated types of feedback, each exhibiting greater dynamics by adapting

---

## Algorithm 1: DualFete in a stochastic training step

---

**Input:** Randomly sampled mini-batches of labeled data

$\mathcal{D}'_l \subset \mathcal{D}_l$  and unlabeled data  $\mathcal{D}'_u \subset \mathcal{D}_u$ .

**Models:** Two teachers  $\phi, \psi$ , and a student  $\theta_S$

- 1: Pseudo-label  $\mathcal{D}'_u$  by  $\phi, \psi$ , respectively.  $\triangleright$  [Eq. 1]
- 2: Get pseudo labels  $\hat{y}^u$  and attributors  $\bar{y}^a, \bar{y}^d$ .  $\triangleright$  [Eq. 6]
- 3: Evaluate feedback  $\delta_{\bar{y}^a}, \delta_{\bar{y}^d}$ .  $\triangleright$  [Eq. 7]
- 4: Get receiver masks  $\mathcal{M}_{\bar{y}^a}^\phi, \mathcal{M}_{\bar{y}^d}^\phi, \mathcal{M}_{\bar{y}^a}^\psi$ , and  $\mathcal{M}_{\bar{y}^d}^\psi$ .
- 5: Compute dual-teacher feedback loss  $\mathcal{L}_{df}$ .  $\triangleright$  [Eq. 8]
- 6: Update  $\theta_S$  by  $\mathcal{L}_S$  using  $\mathcal{D}'_u$  and  $\hat{y}^u$ .  $\triangleright$  [Eq. 2]
- 7: Update  $\{\phi, \psi\}$  by  $\mathcal{L}_T(\phi) + \mathcal{L}_T(\psi)$  using both  $\mathcal{D}'_l$  and  $\mathcal{D}'_u$ .  $\triangleright$  [ $\mathcal{L}_{df}$  and Eq. 9, 10]

**return**  $\phi, \psi$ , and  $\theta_S$

---

to changes in prediction confidence – thereby breaking the limits of uniform updates. Meanwhile, the individualized feedback not only helps prevent teachers from making consistent errors, but also fosters productive disagreement between them, enabling mutual improvement through cross-supervision [Eq. 10] (Qiao et al. 2018; Shen et al. 2023).

## Holistic framework

In total, the training procedure for a stochastic training is presented in Algorithm 1. There exist differences in the training data and objectives between the teacher and the student.

**Student Model.** The student updates its parameters only using  $\mathcal{D}_u$  with pseudo-label  $\hat{y}^u$  yielded by the dual-teacher model [Eq. 6, Eq. 2]. Except for learning on  $\mathcal{D}_u$ , the student takes the duty of feeding back the performance change on labeled data  $\mathcal{D}_l$  per-step of training [Eq. 7]. (Optionally) The student can be further fine-tuned with  $\mathcal{D}_l$ , particularly when the gap between  $N_l$  and  $N_u$  is relatively small.

**Dual-teacher Model.** Feedback from the student alone cannot guarantee superior teachers; they need to improve themselves through both individual and mutual learning. Thus, in addition to  $\mathcal{L}_{df}$ , the teachers are fully-supervised on  $\mathcal{D}_l$  and cross-supervised on  $\mathcal{D}_u$ . Specifically, given a teacher  $\theta \in \{\phi, \psi\}$ , the objective is formulated as

$$\min_{\theta} \mathcal{L}_T(\theta) = \min_{\theta} \mathcal{L}_l(\theta) + \mathcal{L}_{df}(\theta) + \lambda \mathcal{L}_{cs}^A(\theta; \bar{\theta}, \mathcal{A}). \quad (9)$$

Here,  $\lambda$  is a ramp-up weighted factor, and  $\mathcal{A}$  is a strong augmentator [e.g., copy-paste (Ghiasi et al. 2021); color-jitter (Cubuk et al. 2020)].  $\mathcal{L}_l$  is the fully-supervised loss [Eq. 3], and the cross-supervised loss  $\mathcal{L}_{cs}^A$  is defined by

$$\mathcal{L}_{cs}^A(\theta; \bar{\theta}, \mathcal{A}) = \frac{1}{N_u} \sum_i^{N_u} \ell \left( f(\mathcal{A}(x_i^u); \theta), \mathcal{A}(\hat{y}^{\bar{\theta}_u}) \right) \quad (10)$$

Here,  $\hat{y}^{\bar{\theta}_u}$  is the pseudo-label yielded by the teacher  $\bar{\theta}$ , and  $\mathcal{A}(y)$  only operates positional transformations for the label  $y$ . In line with prior work, a confidence threshold is employed to filter possibly unreliable targets (Yang et al. 2023; Huang et al. 2025). Besides, we use  $\mathcal{L}_{cs}$  to refer cross-supervised loss without weak-to-strong consistency.

Methods	LA			Pancreas		BraTS	
	5% (4)	10% (8)	20% (16)	10% (6)	20% (12)	10% (25)	20% (50)
Fully Supervise	52.55, 47.1	82.74, 13.4	86.96, 11.9	55.60, 45.3	72.38, 19.4	74.43, 37.1	80.16, 22.7
UA-MT [Y. (2019)]	82.26, 13.7	86.28, 18.7	88.74, 8.39	66.44, 17.0	76.10, 10.8	84.64, 10.5	85.32, 8.68
SASSNet [L. (2020)]	81.60, 16.2	85.22, 11.2	89.16, 8.95	68.97, 18.8	76.39, 11.1	84.73, 9.88	85.64, 9.17
DTC [L. (2021)]	81.25, 14.9	87.51, 8.23	89.52, 7.07	66.58, 15.5	76.27, 8.70	-	-
URPC [L. (2022b)]	82.48, 14.7	85.01, 15.4	88.74, 12.7	73.53, 22.6	80.02, 8.51	84.53, 9.79	85.38, 8.36
PS-MT [L. (2022b)]	88.49, 8.12	89.72, 6.94	90.02, 6.74	76.94, 13.1	80.74, 7.41	84.88, 9.93	85.91, 8.63
MC-Net+ [W. (2022a)]	83.59, 14.1	88.96, 7.93	91.07, 5.84	70.00, 16.0	80.59, 6.47	84.96, 9.45	86.02, 8.74
SS-Net [W. (2022b)]	86.33, 9.97	88.55, 7.49	89.28, 7.29	71.76, 17.6	78.98, 8.86	-	-
BCP [B. (2023)]	88.02, 7.90	89.62, 6.81	91.26, 5.76	73.83, 12.7	82.91, 6.43	85.14, 9.89	86.13, 8.99
UniMatch [Y. (2023)]	-	89.04, 7.26	90.99, 6.07	-	82.35, 7.66	85.03, 9.50	85.84, 8.68
MutRel [S. (2024)]	87.20, 9.90	89.86, 6.91	91.02, 5.78	75.93, 9.07	81.53, 6.81	84.29, 9.57	85.47, <b>7.76</b>
AD-MT [Z. (2024)]	89.63, 6.56	90.55, 5.81	-	80.21, 7.18	82.61, 4.94	-	-
TraCoCo [L. (2024)]	-	89.86, 6.81	91.51, 5.63	79.22, 8.46	83.36, 7.34	85.71, 9.20	<b>86.69</b> , 8.04
DualFete [ours]	<b>90.35</b> , <u>6.42</u>	<b>91.28</b> , <u>5.51</u>	<u>91.89</u> , <u>5.24</u>	<u>81.99</u> , <b>5.34</b>	<u>83.49</u> , <u>4.76</u>	<u>86.13</u> , <u>9.02</u>	85.83, 8.12
DualFete w.ft. [ours]	<u>90.22</u> , <b>5.89</b>	<u>91.12</u> , <b>5.44</b>	<b>91.91</b> , <b>5.22</b>	<b>82.45</b> , <u>5.96</u>	<b>83.85</b> , <b>4.43</b>	<b>86.25</b> , <b>8.94</b>	86.46, 7.80

Table 1: Comparison with SOTAs on LA, Pancreas, and BraTS19 datasets. The **best** and second best results are highlighted.




Figure 3: Visualizations of several methods (10% labels).

## Experiments

### Experimental settings

**Datasets.** Experiments are conducted on LA (Xiong et al. 2021), Pancreas (Roth et al. 2015), and BraTS19 (Menze et al. 2014) datasets. We follow preprocessing steps and data split used in prior work (Yu et al. 2019; Li et al. 2020; Luo et al. 2021; Liu et al. 2024; Zhao et al. 2024). **LA** contains 100 MRIs, with 80 for training and 20 for testing. Three commonly used label-settings are involved, i.e., 5%, 10%, and 20%. **Pancreas** contains 82 CT scans, in which 62 for training and 20 for testing. Two label settings are used 10%, and 20%. **BraTS** contains 335 brain MRIs, with 250 for training, 25 for validation, and 60 for testing. Label settings are the same as the Pancreas dataset.

**DualFete details.** We employ V-Net for LA and Pancreas, U-Net for BraTS, with the input size and evaluation protocol remaining the same as prior work (Liu et al. 2024). The SGD optimizer and data loading configurations are also in line with prior work (Wu et al. 2022b; Liu et al. 2024). We employ the Dice and cross-entropy combination as the segmentation loss while using cross-entropy loss to evaluate student performance on labeled data. The likelihood  $\mathcal{P}$  is implemented as the cumulative product of per-voxel pseudo-labeling probability. For LA, we employ normalized gradients in Eq. 7 while Pancreas and BraTS do not, as we experimentally find it works better. Moreover, we find that using

likelihood of strong-augmented data can also be beneficial in some cases. Performance is evaluated by Dice (%) and 95% Hausdorff distance (95HD, voxel), displayed on the left and right sides of each table cell, respectively.

### Comparison with SOTAs

We first compare DualFete with the fully supervision baseline (only with  $\mathcal{D}_l$ ) and various state-of-the-art methods, including: UA-MT (Yu et al. 2019), SASSNet (Li et al. 2020), DTC (Luo et al. 2021), URPC (Luo et al. 2022b), PS-MT (Liu et al. 2022b), MC-Net+ (Wu et al. 2022a), SS-Net (Wu et al. 2022b), BCP (Bai et al. 2023), UniMatch (Yang et al. 2023), MutRel (Su et al. 2024), AD-MT (Zhao et al. 2024), and TraCoCo (Liu et al. 2024). The reported results of these methods are identical to those in Su et al. (2024), Zhao et al. (2024), and Liu et al. (2024) under the same settings. Our results are from the student, and we also report the fine-tuned performance (DualFete w.ft.).

While recent state-of-the-art methods [AD-MT (Zhao et al. 2024), TraCoCo (Liu et al. 2024)] have made significant progress in SSMIS, our DualFete pushes these boundaries further. As shown in Table 1, DualFete outperforms existing approaches across nearly all benchmarks and label settings, even without fine-tuning. Notably, on Pancreas with 10% labels, it achieves a +1.78% Dice improvement and 1.84-voxel 95HD reduction over previous best results [AD-MT (Zhao et al. 2024)]. This consistent outperformance suggests that our framework captures fundamental improvements beyond current teacher-student paradigms. The purely unlabeled training of the student model makes it possible to yield additional gains through fine-tuning on labeled data. These gains are substantial in the 20% labeled data setting and on the more challenging Pancreas dataset. However, fine-tuning on very limited data is susceptible to overfitting, e.g., on LA with 5% and 10% labels. Thus, DualFete has particular potential in label-scarce cases by using labeled data exclusively to guide pseudo-label updates toward correct di-

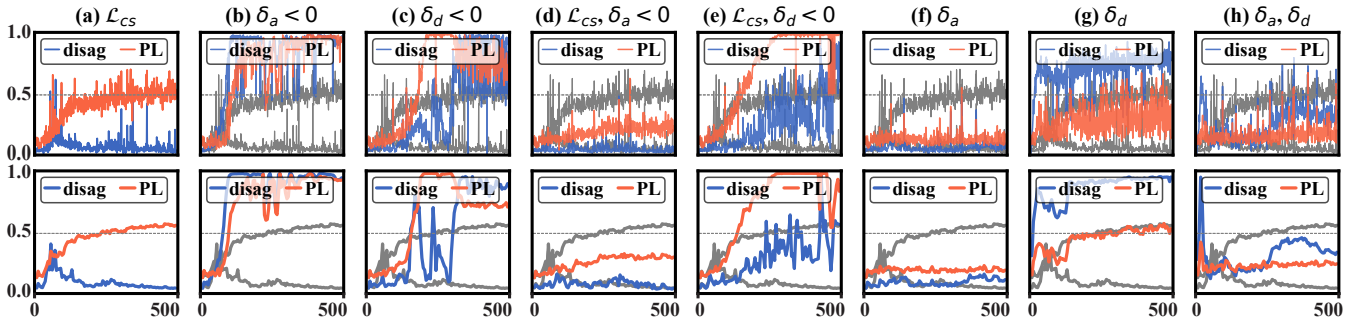


Figure 4: Disagreement between two teachers (disag) and the pseudo-label error (PL) [Eq. 6], measured by 1-Dice. We report results evaluated by training inputs (first row) and by the testing set (last row), respectively. (LA, 10% labels).

rections, thereby avoiding the need for fine-tuning and circumventing the overfitting issue. For the BraTS dataset, it is noteworthy that our DualFete shows some oscillations in terms of testing performance. We speculate that the validation set is prone to overfitting due to its relatively limited quantity (25 samples).

Figure 3 illustrates two challenging segmentation cases in the Pancreas and LA dataset. Our method shows the best alignments with the ground-truth masks, maintaining most spatial details, which suggests the improved generalizability.

## Ablation studies

**Qualitative analysis of DualFete.** To investigate the effectiveness of DualFete, we pretrain a cross-supervised model ( $\mathcal{L}_l + \mathcal{L}_{cs}$ ) for 6k steps and then experimentally impose different constraints with 0.5k extra steps. We evaluate the pseudo-label error and the disagreement between two teachers in the second stage.

Several findings can be summarized by Figure 4. (1) The cross-supervised loss  $\mathcal{L}_{cs}$  forces consensus between the teachers but amplifies errors, leading to consistently worse pseudo-labels [Figure 4(a)]. (2) The constraint of  $\delta_a < 0$  forces teachers to fully conflict with each other [Figure 4(b)], while  $\mathcal{L}_{cs}$  enables them to resolve these conflicts while reducing pseudo-labeling errors [Figure 4(d)]. (3) The  $\delta_d < 0$  constraint induces pseudo-labeling oscillation between the teachers. In extreme cases, this creates collapsing dynamics where object boundaries predicted by the teachers are alternatively eroded, resulting in background-only predictions ultimately. This collapse can be observed from the completely incorrect pseudo-labels yet relatively low disagreement [Figure 4(c,e)]. Nonetheless, the observation suggests that if the constraint is prevented from collapsing, it has the potential to avoid invariant pseudo-labels. (4)  $\delta_d > 0$  can foster disagreements by reinforcing pseudo-label confidence [Figure 4(g)]. Since the student is trained from scratch, it generates predominantly positive feedback, which continually reinforces one teacher’s confidence while relatively weakening the other. Ultimately, two teachers evolve complementary prediction behaviors. (5) The two types of feedback  $\delta_a, \delta_d$  collaboratively counteract degenerating cases found in (2)-(4) [Figure 4(h)].

These findings suggest that the two types of feedback op-

D.T.	Att.	Rec.	$\mathcal{A}$	LA 20%	Pancreas 20%
×	×	×	×	88.55, 8.58	77.18, 9.81
×	$\hat{y}^u$	$\mathcal{P}$	×	89.63, 7.92	79.27, 9.90
✓	$\hat{y}^u$	$\mathcal{P}$	×	89.83, 8.08	76.83, 12.2
✓	$\bar{y}^a$	$\mathcal{P}_l$	×	90.34, 6.43	79.56, 8.38
✓	$\bar{y}^d$	$\mathcal{P}_h$	×	<u>90.35, 6.19</u>	<u>80.77, 8.22</u>
✓	$\bar{y}^{a,d}$	$\mathcal{P}_{h,l}$	×	87.69, 8.73	78.06, 8.77
✓	$\bar{y}^{a,d}$	$\mathcal{P}_{l,h}$	×	<b>90.89, 6.11</b>	<b>81.12, 7.74</b>
×	$\hat{y}^u$	$\mathcal{P}$	✓	89.12, 8.68	78.59, 8.68
✓	$\bar{y}^{a,d}$	$\mathcal{P}_{l,h}$	✓	90.14, 7.30	80.77, 7.89

Table 2: Quantitative analysis of DualFete.  $\hat{y}^u$  is either from Eq. 1 or Eq. 6 depended on whether the dual-teacher model is introduced (first column). In the second and third columns,  $\bar{y}^{a,d}$  and  $\mathcal{P}_{l,h}$  indicate that the update is triggered by agreement/disagreement pseudo-labels, and feedback is applied correspondingly to the lower/higher confidence side.  $\mathcal{A}$  indicates whether  $\mathcal{P}$  is from strong-augmented data.

erate in different manners, yet work synergistically in the dual-teacher model: they (1) generate productive prediction disagreements; (2) maintain pseudo-label accuracy while preventing error accumulation.

**Quantitative analysis of DualFete.** We further explore the proposed feedback mechanism quantitatively, and results are shown in Table 2. Compared to the baseline teacher-student model (first row), the feedback mechanism actively prompts the teacher to refine pseudo-labels, yielding consistent student performance gains (2, 4, 5, and 7-th rows). The single-teacher framework relies solely on naive feedback, while the dual-teacher extension introduces two distinct and collaborative feedback types ( $\delta_a$  and  $\delta_d$ ), each capable of operating independently while surpassing the baseline. This extension achieves significant improvements over the single-teacher protocol, demonstrating the value of coordinated yet independent feedback dynamics. However, performance degrades markedly when feedback attributors and receivers are mismatched (3 and 6-th rows), highlighting the criticality of proper component alignment. Strong-augmented likelihoods here do not result in performance gains, as only when weak-to-strong consistency is imposed can teacher models learn to handle such cases, which will be shown later.

loss	teacher	Ep.1 (Dice)	Ep.2 (Entropy $\times 10^4$ )
$+\mathcal{L}_{cs}^A$	$\phi$	$78.42 \pm 8.46$	$12.65 \pm 2.45$
	$\psi$	$80.51 \pm 6.23$	$13.13 \pm 2.82$
$+\mathcal{L}_{df}$	$\phi$	$57.40 \pm 24.63$	$14.30 \pm 2.77$
	$\psi$	$54.38 \pm 28.83$	$15.23 \pm 2.93$

Table 3: Experiments with either  $\mathcal{L}_{df}$  only or  $\mathcal{L}_{cs}^A$  only. Label setting: Pancreas 20%.

dataset	loss	10%	20%
LA	$\mathcal{L}_l$	86.82, 9.96	88.55, 8.58
	$+\mathcal{L}_{cs}$	88.37, 9.27	90.30, 6.80
	$+\mathcal{L}_{cs}+\mathcal{L}_{df}$	<u>90.77, 6.19</u>	<u>91.66, 6.12</u>
	$+\mathcal{L}_{cs}^A$	90.39, 7.03	91.10, 6.12
	$+\mathcal{L}_{cs}^A+\mathcal{L}_{df}$	<b>91.28, 5.51</b>	<b>91.89, 5.24</b>
Pancreas	$\mathcal{L}_l$	70.81, 14.3	77.18, 9.81
	$+\mathcal{L}_{cs}$	75.39, 10.6	79.62, 6.84
	$+\mathcal{L}_{cs}+\mathcal{L}_{df}$	78.13, 7.71	80.92, 7.89
	$+\mathcal{L}_{cs}^A$	80.30, 6.06	82.17, 5.43
	$+\mathcal{L}_{cs}^A+\mathcal{L}_{df}$	<b>81.99, 5.34</b>	<b>83.49, 4.76</b>

Table 4: Ablation study of  $\mathcal{L}_{df}$ ,  $\mathcal{L}_{cs}$ , and  $\mathcal{L}_{cs}^A$  in the dual-teacher framework. Results are evaluated by the student.

**Cross-teacher supervision.** We design two experiments to verify that the feedback loss  $\mathcal{L}_{df}$  operates distinctly from both consistency regularization and entropy minimization. Specifically, we train two dual-teacher frameworks under identical configurations but with different unsupervised constraints: one using only feedback loss  $\mathcal{L}_{df}$ , and the other employing cross-supervised loss  $\mathcal{L}_{cs}^A$ . In Experiment 1 (Ep.1), each test sample is evaluated six times with different strong augmentations, and we report the mean and standard deviation of Dice scores. Experiment 2 (Ep.2) similarly conducts six evaluations per test sample, yet with dropout perturbations. We measure the mean and standard deviation of per-image entropy summation. Table 3 suggests two key insights of using  $\mathcal{L}_{df}$  alone: (1) it fails to maintain robustness against input perturbations, evidenced by significant Dice degradation and high performance variance; (2) it increases prediction uncertainty, showing no reduction in output entropy.

For further validation, we conducted an ablation study examining the impact of relevant unsupervised constraints. As shown in Table 4, cross-supervision between the dual teachers produces superior student performance, indicating that the teachers generate higher-quality pseudo-labels through mutual learning [ $\mathcal{L}_{cs}$ ,  $\mathcal{L}_{cs}^A$ ]. This virtuous cycle is further amplified by student feedback [ $\mathcal{L}_{df}$ ], which continuously improves teacher pseudo-labeling accuracy, thereby enabling progressive student performance gains.

We experiment with varying confidence thresholds in cross-supervised loss  $\mathcal{L}_{cs}^A$  [Eq. 10] and analyze the effect of the likelihood type used in the feedback loss  $\mathcal{L}_{df}$  [Eq. 8]. Figure 5 demonstrates that a 0.7 threshold achieves optimal performance across most scenarios, validating the importance of filtering low-confidence predictions, which are

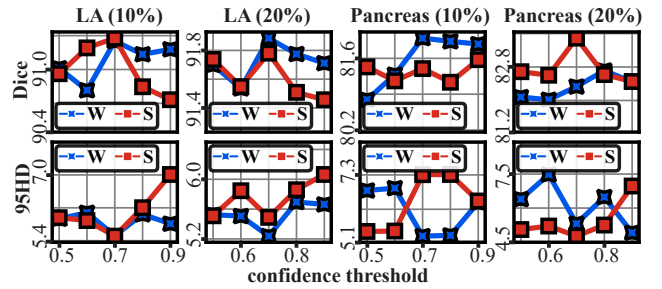


Figure 5: Tuning of confidence threshold in  $\mathcal{L}_{cs}^A$  and strong-augmentation likelihood in  $\mathcal{L}_{df}$  (W: weak, S: strong).

Method	Train (s/iter)	Mem. (GB)	Infer. (s/case)
FullySup	0.15	5.15	1.93
AD-MT	0.67	7.26	1.93
TraCoCo	2.39	21.93	1.81
DualFete	2.28	10.25	1.91

Table 5: Training (time, memory) and inference (time) cost on the LA dataset with V-Net as the backbone.

possibly incorrect. More importantly, preventing converging to low-confidence targets further ensures achieving expected performance using feedback loss, as it mostly modulates the likelihood of these easy-to-error targets, thereby circumventing error accumulation. Notably, strong-augmentation likelihood yields significant improvements on the Pancreas 20% benchmark and also remains effective with lower thresholds. We hypothesize this occurs because teachers benefit from additional low-confidence supervision to better align the likelihood from strong and weak augmentations.

**Efficiency.** The dual-teacher design is used only during training, while inference relies solely on the student. We consider the trade-off worthwhile between training cost and inference performance. As shown in Table 5, our method achieves comparable inference speed (about 1.9 s/case with TITAN RTX) with superior accuracy (see in Table 1).

## Conclusions

In this work, we revisit the teacher-student paradigm through the lens of feedback and propose a dual-teacher feedback model to address confirmation bias in semi-supervised medical image segmentation. Our framework makes two key contributions: (1) an error-correcting feedback mechanism that mitigates error propagation, and (2) a collaborative architecture that integrates the student’s individualized feedback with mutual refinement between teachers, thereby enhancing teaching effectiveness. The dual-teacher design strengthens feedback dynamics, enabling both productive disagreement between teachers and reduced error consistency, two essential factors for improving robustness under imperfect supervision. Extensive experiments on three benchmark datasets demonstrate consistent performance gains, with both quantitative metrics and qualitative analyses validating the effectiveness of our approach.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation for Distinguished Young Scholar under Grant No.62025601, in part by the National Natural Science Foundation Regional Innovation and Development Joint Fund under Grant No.U24A20341, in part by Transformation of Tianfu Jincheng Laboratory under Grant No.2025ZH013, and in part by Sichuan Province Innovative Talent Funding Project for Postdoctoral Fellows under Grant No.BX202512.

## References

- Arazo, E.; Ortego, D.; Albert, P.; O'Connor, N. E.; and McGuinness, K. 2020. Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE.
- Bai, Y.; Chen, D.; Li, Q.; Shen, W.; and Wang, Y. 2023. Bidirectional copy-paste for semi-supervised medical image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11514–11524.
- Chapelle, O.; et al. 2009. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3): 542–542.
- Chen, J.; Mei, J.; Li, X.; Lu, Y.; Yu, Q.; Wei, Q.; Luo, X.; Xie, Y.; Adeli, E.; Wang, Y.; et al. 2024. TransUNet: Rethinking the U-Net architecture design for medical image segmentation through the lens of transformers. *Medical Image Analysis*, 97: 103280.
- Chen, X.; Yuan, Y.; Zeng, G.; and Wang, J. 2021. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2613–2622.
- Chi, H.; Pang, J.; Zhang, B.; and Liu, W. 2024. Adaptive bidirectional displacement for semi-supervised medical image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4070–4080.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 702–703.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ghiasi, G.; Cui, Y.; Srinivas, A.; Qian, R.; Lin, T.-Y.; Cubuk, E. D.; Le, Q. V.; and Zoph, B. 2021. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2918–2928.
- Grandvalet, Y.; and Bengio, Y. 2004. Semi-supervised learning by entropy minimization. *Advances in neural information processing systems*, 17.
- Hu, Y.; Yang, J.; Chen, L.; Li, K.; Sima, C.; Zhu, X.; Chai, S.; Du, S.; Lin, T.; Wang, W.; et al. 2023. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17853–17862.
- Huang, W.; Zhang, L.; Wang, Z.; and Wang, L. 2024. Exploring inherent consistency for semi-supervised anatomical structure segmentation in medical imaging. *IEEE Transactions on Medical Imaging*.
- Huang, W.; Zhang, L.; Wang, Z.; and Wang, Y. 2025. GapMatch: Bridging Instance and Model Perturbations for Enhanced Semi-Supervised Medical Image Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 17458–17466.
- Ke, Z.; Wang, D.; Yan, Q.; Ren, J.; and Lau, R. W. 2019. Dual student: Breaking the limits of the teacher in semi-supervised learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6728–6736.
- Lan, T.; Chen, N.; Yi, Z.; Xu, X.; and Zhu, M. 2025a. Domain Generalization for Pulmonary Nodule Detection via Distributionally-Regularized Mamba. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 152–162. Springer.
- Lan, T.; Yi, Z.; Xu, X.; and Zhu, M. 2025b. LooBox: Loose-box-supervised 3D Tumor Segmentation with Self-correcting Bidirectional Learning. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 8077–8086.
- 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*.
- Li, S.; et al. 2020. Shape-aware semi-supervised 3D semantic segmentation for medical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 552–561. Springer.
- Liu, F.; Tian, Y.; Chen, Y.; Liu, Y.; Belagiannis, V.; and Carneiro, G. 2022a. Acpl: Anti-curriculum pseudo-labelling for semi-supervised medical image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20697–20706.
- Liu, Y.; Tian, Y.; Chen, Y.; Liu, F.; Belagiannis, V.; and Carneiro, G. 2022b. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4258–4267.
- Liu, Y.; Tian, Y.; Wang, C.; Chen, Y.; Liu, F.; Belagiannis, V.; and Carneiro, G. 2024. Translation consistent semi-supervised segmentation for 3d medical images. *IEEE Transactions on Medical Imaging*.
- Luo, X.; Chen, J.; Song, T.; and Wang, G. 2021. Semi-supervised medical image segmentation through dual-task consistency. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 8801–8809.
- Luo, X.; Hu, M.; Song, T.; Wang, G.; and Zhang, S. 2022a. Semi-supervised medical image segmentation via cross teaching between cnn and transformer. In *International conference on medical imaging with deep learning*, 820–833. PMLR.

- Luo, X.; Wang, G.; Liao, W.; Chen, J.; Song, T.; Chen, Y.; Zhang, S.; Metaxas, D. N.; and Zhang, S. 2022b. Semi-supervised medical image segmentation via uncertainty rectified pyramid consistency. *Medical Image Analysis*, 80: 102517.
- Menze, B. H.; Jakab, A.; Bauer, S.; et al. 2014. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE transactions on medical imaging*, 34(10): 1993–2024.
- Nickerson, R. S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2): 175–220.
- Pham, H.; Dai, Z.; Xie, Q.; and Le, Q. V. 2021. Meta pseudo labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11557–11568.
- Qiao, S.; Shen, W.; Zhang, Z.; Wang, B.; and Yuille, A. 2018. Deep co-training for semi-supervised image recognition. In *Proceedings of the european conference on computer vision (eccv)*, 135–152.
- Rollwage, M.; Loosen, A.; Hauser, T. U.; Moran, R.; Dolan, R. J.; and Fleming, S. M. 2020. Confidence drives a neural confirmation bias. *Nature communications*, 11(1): 2634.
- Roth, H. R.; Lu, L.; Farag, A.; et al. 2015. Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, 556–564. Springer.
- Shen, Z.; Cao, P.; Yang, H.; et al. 2023. Co-training with high-confidence pseudo labels for semi-supervised medical image segmentation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 4199–4207.
- 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.
- 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. *Advances in neural information processing systems*, 30.
- Wang, Y.; Xiao, B.; Bi, X.; Li, W.; and Gao, X. 2023. Mcf: Mutual correction framework for semi-supervised medical image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15651–15660.
- Wu, H.; Li, X.; Lin, Y.; and Cheng, K.-T. 2023. Compete to win: Enhancing pseudo labels for barely-supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(11): 3244–3255.
- Wu, Y.; Ge, Z.; Zhang, D.; Xu, M.; Zhang, L.; Xia, Y.; and Cai, J. 2022a. Mutual consistency learning for semi-supervised medical image segmentation. *Medical Image Analysis*, 81: 102530.
- Wu, Y.; Wu, Z.; Wu, Q.; Ge, Z.; and Cai, J. 2022b. Exploring Smoothness and Class-Separation for Semi-supervised Medical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 13435, 34–43. Springer, Cham.
- Xiong, Z.; Xia, Q.; Hu, Z.; Huang, N.; 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, Z.; Wang, Y.; Lu, D.; Luo, X.; Yan, J.; Zheng, Y.; and Tong, R. K.-y. 2023. Ambiguity-selective consistency regularization for mean-teacher semi-supervised medical image segmentation. *Medical Image Analysis*, 88: 102880.
- Yang, L.; Qi, L.; Feng, L.; Zhang, W.; and Shi, Y. 2023. Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7236–7246.
- Yi, L.; Zhang, L.; Zhao, K.; and Xu, X. 2025. Learning from certain regions of interest in medical images via probabilistic positive-unlabeled networks. *Medical Image Analysis*, 103745.
- 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 *Medical image computing and computer assisted intervention—MICCAI 2019*, 605–613. Springer.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*.
- Zhao, Z.; Wang, Z.; Wang, L.; Yu, D.; Yuan, Y.; and Zhou, L. 2024. Alternate diverse teaching for semi-supervised medical image segmentation. In *European Conference on Computer Vision*, 227–243. Springer.