



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

SDCL: Students Discrepancy-Informed Correction Learning for Semi-supervised Medical Image Segmentation

Bentao Song and Qingfeng Wang (✉)

School of Computer Science and Technology, Southwest University of Science and Technology, China
qfwangyy@mail.ustc.edu.cn

Abstract. Semi-supervised medical image segmentation (SSMIS) has been demonstrated the potential to mitigate the issue of limited medical labeled data. However, confirmation and cognitive biases may affect the prevalent teacher-student based SSMIS methods due to erroneous pseudo-labels. To tackle this challenge, we improve the mean teacher approach and propose the Students Discrepancy-Informed Correction Learning (SDCL) framework that includes two students and one non-trainable teacher, which utilizes the segmentation difference between the two students to guide the self-correcting learning. The essence of SDCL is to identify the areas of segmentation discrepancy as the potential bias areas, and then encourage the model to review the correct cognition and rectify their own biases in these areas. To facilitate the bias correction learning with continuous review and rectification, two correction loss functions are employed to minimize the correct segmentation voxel distance and maximize the erroneous segmentation voxel entropy. We conducted experiments on three public medical image datasets: two 3D datasets (CT and MRI) and one 2D dataset (MRI). The results show that our SDCL surpasses the current State-of-the-Art (SOTA) methods by 2.57%, 3.04%, and 2.34% in the Dice score on the Pancreas, LA, and ACDC datasets, respectively. In addition, the accuracy of our method is very close to the fully supervised method on the ACDC dataset, and even exceeds the fully supervised method on the Pancreas and LA dataset. (Code available at <https://github.com/pascalcpp/SDCL>).

Keywords: Semi-supervised learning · Medical image segmentation · Correction learning

1 Introduction

Deep learning techniques have revolutionized medical image processing, mainly due to the superiority of neural network algorithms over many conventional image processing techniques. Relative to the general computer vision field, medical image segmentation encounters the dual challenges of the scarcity of annotation data and the greater complexity of datasets. These data constraints suppress the accuracy of medical image segmentation, and semi-supervised learning

(SSL) methods are increasingly being recognized for their potential to tackle these challenges by leveraging both labeled and unlabeled data.

Current popular SSMIS methods focus on self-training, uncertainty estimation, consistency regularization, and distribution alignment. Self-training methods, such as self-training and co-training [2, 30], use current high-confidence pseudo-labels and ground truths for iteratively training. Uncertainty estimation uses measures like information entropy to assess unlabeled data and guide pseudo-label filtering or weighting. Mean Teacher (MT) [21] framework is a prevalent approach that enforces consistency between student and teacher models. UA-MT [28] refines consistency learning of MT with uncertainty, while CoraNet [20] applies different weights to teacher-generated pseudo-labels based on uncertainty. BCP [3] aims to reduce the empirical distribution gap by learning common semantics from both labeled and unlabeled data in the MT framework.

Despite the ongoing progress in SSMIS, confirmation and cognitive biases remains a critical limitation, especially in the widely used semi-supervised learning based on teacher-student framework. Since the framework often adds input perturbations and applies consistency regularization between teacher and student, a single model structure inevitably produces noisy or erroneous pseudo-labels [3, 22], resulting in model confirmation and cognitive biases [1, 22]. These biases limit the performance of the teacher-student framework and it is very difficult for the model to correct these biases on its own.

Recently, methods like multi-student [7, 15] and multi-teacher [12, 18] have emerged to provide diverse pseudo-labels to mitigate confirmation and cognitive biases. Multi-student approaches involve cross-consistency learning, [15] uses differences in learning paradigms of different network structures to achieve perturbation effects, followed by consistency learning through cross teaching. But they may encounter training instability without an Exponential Moving Average (EMA) teacher, and diversity and stability can be difficult to balance. On the other hand, the multi-teacher methods employ a single student to update multiple teachers with various strategies to promote diverse learning. However, these methods are constrained by the single model structure, limiting their ability to adequately address biases. The researchers have also employed correction learning methods, such as [29] incorporating dual-task network for bias correction and [23] utilizing complementary network for mapping to ground truth. MCF [22] proposes inter-subnet interaction as a means of bias corrections.

In this view, the teacher-student methods still lack a general approach for stably rectify own biases using diverse information, and the incorporation of bias correction can help to improve the performance of SSMIS. Therefore, we propose students discrepancy-informed correction learning (SDCL) based on the Mean Teacher (MT) framework, featuring one self-ensembling teacher with two trainable students. We ensure stability with an EMA teacher and promote diversity by using students with different structures. SDCL considers the discrepancy areas between the segmentations of the two students as the potential bias areas, and then conducts correction learning in these areas. The contributions of this study include: (1) Different from the traditional teacher-student framework, we

use two structurally different students and an EMA teacher to ensure the diversity and stability of the teacher-student framework. (2) We design a method to optimize bias correction learning that reviews correct cognition and rectifies error biases in the differences between the predictions of two students. (3) Our approach outperforms SOTA SSMIS methods on three datasets, and additionally, it performs comparably or surpasses the fully supervised method.

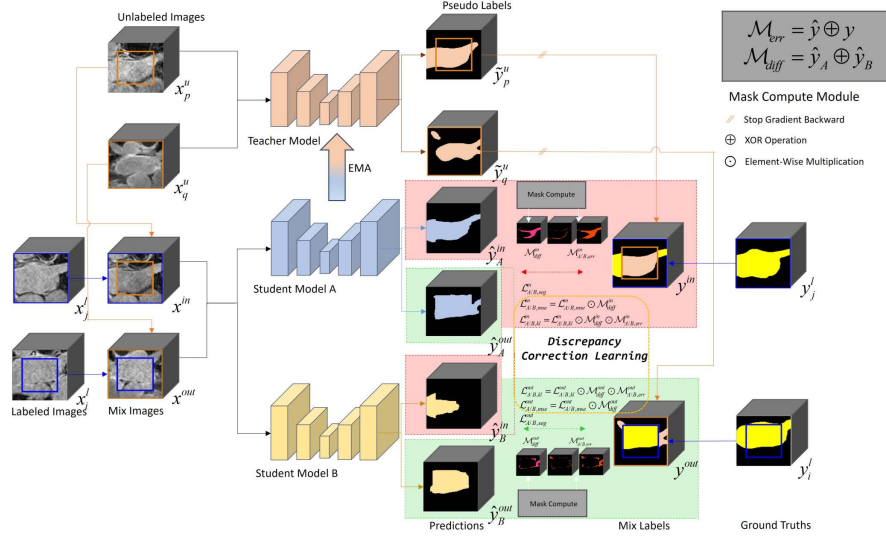


Fig. 1. Overview of the proposed students discrepancy-informed correction learning (SDCL) framework for semi-supervised medical image segmentation. The SDCL framework adopts a BCP strategy, merging two labeled and two unlabeled images to create mix images. The teacher uses unlabeled images to generate pseudo-labels. Ground truths and pseudo-labels are then mixed to produce mix labels. Students A and B process the mixed images separately to calculate segmentation loss. Finally, mask computation determines discrepancy and error masks, guiding the correction learning process.

2 Methodology

2.1 Problem Definition

Given a medical image dataset \mathcal{D} , it contains N labeled images \mathcal{D}^l and M unlabeled images \mathcal{D}^u ($N \ll M$), i.e. $\mathcal{D} = \{\mathcal{D}^l, \mathcal{D}^u\}$, where $\mathcal{D}^l = \{(x_i^l, y_i^l)\}_{i=1}^N$ and $\mathcal{D}^u = \{x_i^u\}_{i=N+1}^{N+M}$. Each 3D volume medical image $x_i \in R^{W \times H \times D}$ in \mathcal{D}^l have label $y_i^l \in \{0, 1, \dots, K-1\}^{W \times H \times D}$. The output prediction of the model is $\hat{y}_i \in \{0, 1, \dots, K-1\}^{W \times H \times D}$. SDCL includes two students and a self-ensembling teacher.

2.2 SDCL Framework

Our framework consists of two phases: initial pre-training with labeled data using Copy-Paste augmentation [3,9], followed by a semi-supervised learning (SSL) phase incorporating both labeled and unlabeled data. SSL begins by initializing students and the teacher with the pre-trained model. The process of SSL includes three parts: i) obtaining the basic SSL segmentation losses based on the BCP strategy, ii) obtaining DiffMask \mathcal{M}_{diff} from student segmentation discrepancies to guide the framework in reviewing correct cognition voxels, and iii) generating ErrMask \mathcal{M}_{err} from the difference between student segmentations and mix labels, then creating DiffErrMask $\mathcal{M}_{differr}$ by multiplying \mathcal{M}_{diff} and \mathcal{M}_{err} , which guides the repair of self-bias error voxels. To achieve diversity between the two students, for 3D tasks, we employ VNet as student A and ResVNet [22] as student B. Meanwhile, for 2D tasks, UNet is designated as student A, and ResUNet is assigned as student B. To accurately evaluate the impact of discrepancy correction learning, efforts are made to minimize interference from other factors. We use VNet/UNet (student A) for Exponential Moving Average (EMA) updates to the teacher, aligning with other methodologies. It’s important to emphasize that the performance gap between the two students is minimal. The framework is illustrated in Fig. 1.

2.3 Bidirectional Copy-Paste

We combine our framework with the current SSMIS SOTA method BCP [3]. In the SSL phase, we need to generate a zero-centered mask $\mathcal{M} \in (0, 1)^{W \times H \times D}$, where 0 represents foreground and 1 represents background, the size of 0 region is $\beta W \times \beta H \times \beta D$, $\beta \in (0, 1)$. Next, we use the mask to obtain mix images as the input of SSL as follows:

$$x^{in} = x_j^l \odot \mathcal{M} + x_p^u \odot (1 - \mathcal{M}), \quad x^{out} = x_q^u \odot \mathcal{M} + x_i^l \odot (1 - \mathcal{M}). \quad (1)$$

To obtain the pseudo labels, x_p^u and x_q^u are forwarded to the teacher to obtain pseudo-labels p_p^u and p_q^u . Because pseudo-labels contain a lot of noise [3, 10], which is very harmful to model training, the optimized pseudo labels \tilde{y}_p^u and \tilde{y}_q^u are obtained by selecting the largest connected component of raw pseudo labels. Then the mix labels are defined as follows:

$$y^{in} = y_j^l \odot \mathcal{M} + \tilde{y}_p^u \odot (1 - \mathcal{M}), \quad y^{out} = \tilde{y}_q^u \odot \mathcal{M} + y_i^l \odot (1 - \mathcal{M}), \quad (2)$$

where $i \neq j$ and $p \neq q$, and \odot denotes element-wise multiplication. Next, fed x^{in} and x^{out} into two students to obtain \hat{y}^{in} and \hat{y}^{out} for each student. Finally, we get BCP losses computed respectively by Eq. (3) and Eq. (4):

$$\mathcal{L}_{A/B,seg}^{in} = \mathcal{L}_{seg}(\hat{y}_{A/B}^{in}, y^{in}) \odot \mathcal{M} + \alpha \mathcal{L}_{seg}(\hat{y}_{A/B}^{in}, y^{in}) \odot (1 - \mathcal{M}), \quad (3)$$

$$\mathcal{L}_{A/B,seg}^{out} = \mathcal{L}_{seg}(\hat{y}_{A/B}^{out}, y^{out}) \odot (1 - \mathcal{M}) + \alpha \mathcal{L}_{seg}(\hat{y}_{A/B}^{out}, y^{out}) \odot \mathcal{M}, \quad (4)$$

\mathcal{L}_{seg} consists of Dice and Cross-entropy loss in equal proportions. Since ground truths are generally more accurate than pseudo labels, α is used to control the weight between them.

2.4 Discrepancy Correction Learning

Minimize Discrepancy Correct Distance. From an intuitively straightforward perspective, during training, we can increase the weighting of learning in the correct regions to review correct voxels and avoid model biases. Here, we minimize Mean Squared Error (MSE) to enhance the model’s learning towards the correct voxels in discrepant regions. We first apply argmax to \hat{y}^{in} and \hat{y}^{out} , resulting in \tilde{y}^{in} and \tilde{y}^{out} . Then, \mathcal{M}_{diff} is then derived from the following formula: $\mathcal{M}_{diff}^{in/out} = \tilde{y}_A^{in/out} \oplus \tilde{y}_B^{in/out}$, where A and B respectively represent student A and B, \oplus denotes XOR operation. Intuitively, we can obtain losses using a method similar to BCP, and the MSE losses are computed through Eq. (5), Eq. (6), and Eq. (7):

$$\mathcal{L}_{A/B,mse}^{in} = \mathcal{L}_{mse}(\hat{y}_{A/B}^{in}, y^{in}) \odot \mathcal{M} + \alpha \mathcal{L}_{mse}(\hat{y}_{A/B}^{in}, y^{in}) \odot (1 - \mathcal{M}), \quad (5)$$

$$\mathcal{L}_{A/B,mse}^{out} = \mathcal{L}_{mse}(\hat{y}_{A/B}^{out}, y^{out}) \odot (1 - \mathcal{M}) + \alpha \mathcal{L}_{mse}(\hat{y}_{A/B}^{out}, y^{out}) \odot \mathcal{M}, \quad (6)$$

$$\mathcal{L}_{A/B,mse}^{in/out} = \mathcal{L}_{A/B,mse}^{in/out} \odot \mathcal{M}_{diff}^{in/out}. \quad (7)$$

Maximize Discrepancy Erroneous Entropy. In the teacher-student framework, inherent confirmation and cognitive biases emerge. We incorporate a loss function for penalization to enable the model to self-correct these biases. Maximizing the entropy of erroneous voxels may be an effective strategy, pulling misclassified voxels in uncertain regions back to an initial state and redirecting them toward the correct direction. The entropy of each voxel is defined as $\mathcal{H}(\hat{y}^{(x,y,z)}) = -\sum_{c=0}^{K-1} \hat{y}^{(x,y,z)}(c) \log \hat{y}^{(x,y,z)}(c)$, where (x, y, z) denotes the voxel’s position, and $\hat{y}^{(x,y,z)} \in \mathcal{R}^K$. The objective is to maximize $\mathcal{H}(\hat{y}^{(x,y,z)})$ for each erroneously classified voxel. We employ Kullback-Leibler (KL) divergence [4], a simple equivalent variant, to guide misclassified voxels in shifting their output distribution towards a uniform distribution. This minimizes the loss equation $\mathcal{L}_{kl}(\hat{y}, u) = \mathcal{D}_{KL}(u||\hat{y})$, where u represents the uniform distribution that all components are equal to $\frac{1}{K}$. Before deriving losses, we obtain \mathcal{M}_{err} using the following expression: $\mathcal{M}_{A/B,err}^{in/out} = \tilde{y}_{A/B}^{in/out} \oplus y^{in/out}$. Afterward, $\mathcal{M}_{A/B,differr}^{in/out} = \mathcal{M}_{diff}^{in/out} \odot \mathcal{M}_{A/B,err}^{in/out}$. The KL losses we obtain are as shown in the following equations Eq. (8), Eq. (9), and Eq. (10):

$$\mathcal{L}_{A/B,kl}^{in} = \mathcal{L}_{kl}(\hat{y}_{A/B}^{in}, u) \odot \mathcal{M} + \alpha \mathcal{L}_{kl}(\hat{y}_{A/B}^{in}, u) \odot (1 - \mathcal{M}), \quad (8)$$

$$\mathcal{L}_{A/B,kl}^{out} = \mathcal{L}_{kl}(\hat{y}_{A/B}^{out}, u) \odot (1 - \mathcal{M}) + \alpha \mathcal{L}_{kl}(\hat{y}_{A/B}^{out}, u) \odot \mathcal{M}, \quad (9)$$

$$\mathcal{L}_{A/B,kl}^{in/out} = \mathcal{L}_{A/B,kl}^{in/out} \odot \mathcal{M}_{A/B,differr}^{in/out}. \quad (10)$$

In the final step, we form the total loss by linearly combining \mathcal{L}_{seg} , \mathcal{L}_{mse} and \mathcal{L}_{kl} with specific weights, as shown in Eq. (11),

$$\mathcal{L}_{A/B} = \mathcal{L}_{A/B,seg}^{in} + \mathcal{L}_{A/B,seg}^{out} + \gamma(\mathcal{L}_{A/B,mse}^{in} + \mathcal{L}_{A/B,mse}^{out}) + \mu(\mathcal{L}_{A/B,kl}^{in} + \mathcal{L}_{A/B,kl}^{out}). \quad (11)$$

Table 1. Performance comparison on the Left Atrium and Pancreas datasets.

Method	Pancreas-CT					Left Atrium						
	Lb	Unlb	Dice	Jac	95HD ASD	Lb	Unlb	Dice	Jac	95HD ASD		
V-Net	12	0	70.59	56.77	14.19	2.25	8	0	79.87	67.60	26.65	7.94
ResV-Net	12	0	68.94	54.50	13.86	3.36	8	0	80.07	69.29	19.50	6.02
V-Net	62	0	82.60	70.81	5.61	1.33	80	0	91.47	84.36	5.48	1.51
ResV-Net	62	0	82.46	70.50	5.45	1.44	80	0	91.09	83.90	4.77	1.75
UA-MT [28](MICCAI'19)			77.26	63.82	11.90	3.06			87.79	78.39	8.68	2.12
DTC [14](AAAI'21)			78.27	64.75	8.36	2.25			87.51	78.17	8.23	2.36
CoraNet [20](TMI'21)			79.67	66.69	7.59	1.89			-	-	-	-
SS-Net [25](MICCAI'22)			-	-	-	-			88.55	79.62	7.49	1.90
MC-Net+ [24](MIA'22)			80.59	68.08	6.47	1.74			88.96	80.25	7.93	1.86
CAML [8](MICCAI'23)	12	50	-	-	-	-	8	72	89.62	81.28	8.76	2.02
DMD [26](MICCAI'23)			-	-	-	-			89.70	81.42	6.88	1.78
MCCauSSL [17](ICCV'23)			80.92	68.26	8.11	1.53			-	-	-	-
UPCoL [13](MICCAI'23)			81.78	69.66	3.78	0.63			-	-	-	-
BCP [3](CVPR'23)			82.91	70.97	6.43	2.25			89.62	81.31	6.81	1.76
SDCL(Ours)			85.04	74.22	5.22	1.48			92.35	85.83	4.22	1.44

Table 2. Performance on ACDC dataset

Method	ACDC					
	Lb	Unlb	Dice	Jac	95HD ASD	
U-Net	7	0	79.41	68.11	9.35	2.70
ResU-Net	7	0	80.04	68.73	7.83	1.94
U-Net	70	0	91.65	84.93	1.89	0.56
ResU-Net	70	0	90.44	82.95	1.77	0.47
UA-MT [28](MICCAI'19)			81.65	70.64	6.88	2.02
SASSNet [11](MICCAI'20)			84.50	74.34	5.42	1.86
DTC [14](AAAI'21)			84.29	73.92	12.81	4.01
SS-Net [25](MICCAI'22)			86.78	77.67	6.07	1.40
MC-Net+ [24](MIA'22)	7	63	87.10	78.06	6.68	2.00
DC-Net [6](MICCAI'23)			89.42	81.37	1.28	0.38
BCPCauSSL [17](ICCV'23)			89.66	81.79	3.67	0.93
BCP [3](CVPR'23)			88.84	80.62	3.98	1.17
SDCL(Ours)			90.92	83.83	1.29	0.34

3 Experiments

Datasets. Our method is evaluated on three datasets: the Pancreas-NIH [19] with 82 CT volumes (12 labeled (20%), 50 unlabeled), the LA (Left Atrium) [27] dataset with 100 3D GE-MRIs (8 labeled (10%), 72 unlabeled for training, 20 for testing), and the ACDC [5] dataset with 100 cardiac MRI scans (70 for training, 10 for validation, 20 for testing, with 7 labeled(10%) and 63 unlabeled). For fair comparison, we use the same experimental setup as prior works like CoraNet and SS-Net [16, 20, 25], normalizing images and applying standard data augmentation. We follow the dataset splits used in previous studies. The ACDC result represents the average performance of four-class segmentation on the test set.




Fig. 2. 2D and 3D Visualization of segmentation results on Left Atrium dataset.

Implementation Details. We used PyTorch 2.1.0 and an NVIDIA RTX 4090 GPU to experiment, averaging segmentation results from students A and B for evaluation. Pre-training had iterations of 3k (Pancreas), 3k (LA), and 11k (ACDC), while SSL had iterations of 8k (Pancreas), 10k (LA), and 45k (ACDC). We employed the Adam optimizer with a learning rate of 0.001, a batch size of 8, 8, 48 (half for labeled and half for unlabeled data), and set hyperparameters $\alpha = 0.5$, $\beta = 2/3$. Specific values for γ and μ were 0.3 and 0.1 (Pancreas), 0.5 and 0.05 (LA and ACDC). Input patch sizes were $96 \times 96 \times 96$ (Pancreas), $112 \times 112 \times 80$ (LA), and 256×256 (ACDC). Testing used strides of $16 \times 16 \times 4$ (Pancreas) and $18 \times 18 \times 4$ (LA) [16, 20, 25], with no post-processing during evaluation. Performance metrics included Dice Score (Dice), Jaccard Score (Jac), 95% Hausdorff Distance (95HD), and Average Surface Distance (ASD).

Results on the LA and Pancreas Datasets. Table 1 provides a detailed comparison of our approach with the current SOTA SSMIS methods and presents the performance upper and lower bounds of our two backbones using only labeled data in a fully supervised manner. The table demonstrates significant improvements in our method across four metrics compared to the baseline (BCP) and other SOTA methods. As depicted in Fig. 2, our approach closely aligns with the Ground Truth, especially in regions prone to errors at boundaries and connections. This underscores the efficacy of our discrepancy correction learning, significantly contributing to enhancing the model’s edge and shape segmentation ability.

Results on the ACDC Dataset. Table 2 presents a comparative analysis of our results against current methods and the upper and lower bounds of fully supervised methods. Similarly, our method exhibits a substantial improvement over the baseline, particularly evident in a significant enhancement of the ASD

metric. Our approach demonstrated a 39% reduction on the metric of ASD relative to the fully supervised upper bound of U-Net. The presence of more 2D slices in the dataset likely contributes to the enhancement of geometric segmentation integrity through discrepancy correction learning.

Ablation Study. Table. 3, we conducted an ablation study on various components of our framework, comparing their impact on Dice scores against the baseline. Omitting \mathcal{M}_{diff} , \mathcal{L}_{mse} and \mathcal{L}_{kl} individually resulted in improvements of 0.43% and 0.38%, while the simultaneous use of both losses showed an enhancement of over 0.53%. Introducing \mathcal{M}_{diff} increased the improvements to 1.17% and 1.26%, and utilizing all components achieved a significant improvement of 2.16%. These findings highlight the positive effect of incorporating correct cognition review and self-bias error correction in SDCL on model performance. Notably, we observed the most favorable outcomes when simultaneously employing both losses and \mathcal{M}_{diff} , affirming the efficacy of our proposed correction learning based on students’ discrepancy. For further insights, the Supplementary Materials include the variation of biased error voxels during the training process, detailed experiment results, and hyperparameter ablation studies in this paper.

Table 3. Ablation study results on the pancreas dataset.

Scans used		Components				Metrics			
Lb	Unlb	\mathcal{L}_{seg}	\mathcal{L}_{mse}	\mathcal{L}_{kl}	\mathcal{M}_{diff}	Dice	Jac	95HD	ASD
		✓				83.23	71.57	8.53	2.49
		✓	✓			83.59	72.18	7.20	2.30
		✓	✓		✓	84.20	73.01	6.25	2.03
12	50	✓		✓		83.55	72.05	7.36	2.09
		✓		✓	✓	84.28	73.12	6.31	1.97
		✓	✓	✓		83.67	72.20	9.12	2.80
		✓	✓	✓	✓	85.04	74.23	5.22	1.48

4 Conclusion

We propose a novel SSMIS framework, extending the Mean-Teacher with an additional student for correction learning based on student discrepancies. The method aims to review correct voxels and repair self-bias error voxels in discrepancies and is compatible with other teacher-student models. It outperforms existing methods in 2D and 3D tasks. Future research will consider using students’ information to refine teacher and further improve the performance of semi-supervised learning in medical image segmentation.

Acknowledgments. This work was supported by the Natural Science Foundation of SiChuan, China (No.2022NSFSC0940).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arazo, E., Ortego, D., Albert, P., O'Connor, N.E., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. In: 2020 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE (2020)
2. Bai, W., Oktay, O., Sinclair, M., Suzuki, H., Rajchl, M., Tarroni, G., Glocker, B., King, A., Matthews, P.M., Rueckert, D.: Semi-supervised learning for network-based cardiac mr image segmentation. In: MICCAI 2017. pp. 253–260. Springer (2017)
3. Bai, Y., Chen, D., Li, Q., Shen, W., Wang, Y.: Bidirectional copy-paste for semi-supervised medical image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11514–11524 (2023)
4. Belharbi, S., Rony, J., Dolz, J., Ayed, I.B., McCaffrey, L., Granger, E.: Deep interpretable classification and weakly-supervised segmentation of histology images via max-min uncertainty. *IEEE Transactions on Medical Imaging* **41**(3), 702–714 (2021)
5. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
6. Chen, F., Fei, J., Chen, Y., Huang, C.: Decoupled consistency for semi-supervised medical image segmentation. In: MICCAI 2023. pp. 551–561. Springer (2023)
7. Chen, X., Yuan, Y., Zeng, G., Wang, J.: Semi-supervised semantic segmentation with cross pseudo supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2613–2622 (2021)
8. Gao, S., Zhang, Z., Ma, J., Li, Z., Zhang, S.: Correlation-aware mutual learning for semi-supervised medical image segmentation. In: MICCAI 2023. pp. 98–108. Springer (2023)
9. Ghiasi, G., Cui, Y., Srinivas, A., Qian, R., Lin, T.Y., Cubuk, E.D., Le, Q.V., Zoph, B.: 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. pp. 2918–2928 (2021)
10. Han, K., Liu, L., Song, Y., Liu, Y., Qiu, C., Tang, Y., Teng, Q., Liu, Z.: An effective semi-supervised approach for liver ct image segmentation. *IEEE Journal of Biomedical and Health Informatics* **26**(8), 3999–4007 (2022)
11. Li, S., Zhang, C., He, X.: Shape-aware semi-supervised 3d semantic segmentation for medical images. In: MICCAI 2020. pp. 552–561. Springer (2020)
12. Liu, Y., Tian, Y., Chen, Y., Liu, F., Belagiannis, V., Carneiro, G.: Perturbed and strict mean teachers for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4258–4267 (2022)
13. Lu, W., Lei, J., Qiu, P., Sheng, R., Zhou, J., Lu, X., Yang, Y.: Upcol: Uncertainty-informed prototype consistency learning for semi-supervised medical image segmentation. In: MICCAI 2023. pp. 662–672. Springer (2023)
14. Luo, X., Chen, J., Song, T., Wang, G.: Semi-supervised medical image segmentation through dual-task consistency. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 8801–8809 (2021)

15. Luo, X., Hu, M., Song, T., Wang, G., Zhang, S.: Semi-supervised medical image segmentation via cross teaching between cnn and transformer. In: International Conference on Medical Imaging with Deep Learning. pp. 820–833. PMLR (2022)
16. Luo, X., Wang, G., Liao, W., Chen, J., Song, T., Chen, Y., Zhang, S., Metaxas, D.N., Zhang, S.: Semi-supervised medical image segmentation via uncertainty rectified pyramid consistency. *Medical Image Analysis* **80**, 102517 (2022)
17. Miao, J., Chen, C., Liu, F., Wei, H., Heng, P.A.: Caussl: Causality-inspired semi-supervised learning for medical image segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21426–21437 (2023)
18. Na, J., Ha, J.W., Chang, H.J., Han, D., Hwang, W.: Switching temporary teachers for semi-supervised semantic segmentation. *Advances in Neural Information Processing Systems* **36** (2024)
19. Roth, H.R., Lu, L., Farag, A., Shin, H.C., Liu, J., Turkbey, E.B., Summers, R.M.: Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In: MICCAI 2015. pp. 556–564. Springer (2015)
20. Shi, Y., Zhang, J., Ling, T., Lu, J., Zheng, Y., Yu, Q., Qi, L., Gao, Y.: Inconsistency-aware uncertainty estimation for semi-supervised medical image segmentation. *IEEE transactions on medical imaging* **41**(3), 608–620 (2021)
21. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* **30** (2017)
22. Wang, Y., Xiao, B., Bi, X., Li, W., Gao, X.: Mcf: Mutual correction framework for semi-supervised medical image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15651–15660 (2023)
23. Wu, S., Li, J., Liu, C., Yu, Z., Wong, H.S.: Mutual learning of complementary networks via residual correction for improving semi-supervised classification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6500–6509 (2019)
24. Wu, Y., Ge, Z., Zhang, D., Xu, M., Zhang, L., Xia, Y., Cai, J.: Mutual consistency learning for semi-supervised medical image segmentation. *Medical Image Analysis* **81**, 102530 (2022)
25. Wu, Y., Wu, Z., Wu, Q., Ge, Z., Cai, J.: Exploring smoothness and class-separation for semi-supervised medical image segmentation. In: MICCAI 2022. pp. 34–43. Springer (2022)
26. Xie, Y., Yin, Y., Li, Q., Wang, Y.: Deep mutual distillation for semi-supervised medical image segmentation. In: MICCAI 2023. pp. 540–550. Springer (2023)
27. Xiong, Z., Xia, Q., Hu, Z., Huang, N., Bian, C., Zheng, Y., Vesal, S., Ravikumar, N., Maier, A., Yang, X., et al.: A global benchmark of algorithms for segmenting the left atrium from late gadolinium-enhanced cardiac magnetic resonance imaging. *Medical image analysis* **67**, 101832 (2021)
28. Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In: MICCAI 2019. pp. 605–613. Springer (2019)
29. Zhang, R., Liu, S., Yu, Y., Li, G.: Self-supervised correction learning for semi-supervised biomedical image segmentation. In: MICCAI 2021. pp. 134–144. Springer (2021)
30. Zhou, Y., Wang, Y., Tang, P., Bai, S., Shen, W., Fishman, E., Yuille, A.: Semi-supervised 3d abdominal multi-organ segmentation via deep multi-planar co-training. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 121–140. IEEE (2019)