A Virtual Domain Collaborative Learning Framework for Semi-supervised Microscopic Hyperspectral Image Segmentation

Geng Qin[1], Huan Liu[1], Wei Li[1]✉, Haihao Zhang[1], and Yuxing Guo[2]

[1] School of Information and Electronics, Beijing Institute of Technology, China
[2] the Department of Oral and Maxillofacial Surgery, Peking University School and Hospital of Stomatology, China
liwei089@ieee.org

Abstract. Microscopic hyperspectral image segmentation faces dual challenges of limited labeled data and insufficient utilization of unlabeled data. However, existing semi-supervised methods often isolate the training processes for labeled and unlabeled data, neglecting their potential synergistic effects. To address this, we propose a semi-supervised method based on **V**irtual **D**omain **C**ollaborative **L**earning (VDCL) to enhance the collaborative learning ability between labeled and unlabeled data and improve the quality of pseudo-labels. Specifically, by combining unlabeled background with labeled foreground and labeled background with unlabeled foreground to construct virtual domain data pairs, we established a collaborative learning bridge between labeled and unlabeled samples. Furthermore, we establish a repository of optimal models and employ an alternating co-training strategy. The current and historically optimal models jointly guide training, and this dynamic framework significantly improves pseudo-labels quality. We have verified the novel semi-supervised segmentation method on the widely-used public microscopic hyperspectral choledoch dataset from Kaggle and the oral squamous cell carcinoma dataset. On these datasets, our method has achieved the state-of-the-art performance. The code is available at https://github.com/Qugeryolo/Virual-Domain.

Keywords: Virtual domain · Alternate learning · Hyperspectral image · Semi-supervised segmentation.

1 Introduction

Microscopic hyperspectral imaging (MHSI) has emerged as a promising modality in medical diagnostics [1–3], combining the high spatial resolution of traditional microscopy with rich spectral information that captures subtle biochemical variations in tissues [4, 5]. This unique ability to provide detailed molecular and structural insights has positioned HSI as a transformative tool for applications such as disease diagnosis, tumor margin delineation, and pathological assessment [6–8]. Despite its potential, the adoption of MHSI in clinical practice faces significant challenges, particularly in the domain of MHSI segmentation.

A critical bottleneck in developing accurate and robust segmentation models lies in the requirement for large volumes of annotated data [9–11]. Manual labeling of MHSI is labor-intensive, time-consuming, and requires domain expertise, especially given the high dimensionality and complexity of HSI data [7, 8, 24]. Consequently, the scarcity of labeled datasets significantly limits the scalability and performance of fully supervised segmentation methods.

To alleviate the reliance on extensive labeled data, semi-supervised learning (SSL) has gained traction as a promising alternative. Existing research falls into three categories: (1) methods based on weak-to-strong consistency, where pseudo-labels are generated from weak perturbations and used to supervise predictions from strongly perturbed inputs [12–14]; (2) the Exponential Moving Average (EMA) teacher-based framework for semi-supervised segmentation [15, 16, 18], where a teacher network is derived from the EMA of the student model's weights; and (3) co-training paradigms [19, 20], which involve training multiple networks with different initializations in a mutually instructive manner. However, these methods still face two main limitations: (1) Existing enhancement methods, such as CutMix [26] and ClassMix [27], treat labeled and unlabeled data as isolated streams, but ignore the potential collaborative associations between labeled and unlabeled data. (2) The process of generating and refining pseudo-labels, which is crucial for effectively leveraging unlabeled data, frequently suffers from a lack of robustness and adaptability. This shortcoming ultimately results in suboptimal model performance.

To address these challenges, we propose a novel semi-supervised framework for microscopic hyperspectral image segmentation based on **V**irtual **D**omain **C**ollaborative **L**earning (VDCL). There are three aspects to our contribution: (1) We propose a virtual domain that bridges labeled and unlabeled data, enabling the seamless transfer of knowledge between the two. By creating virtual and real data pairs, we facilitate the collaborative learning between annotated and unannotated samples, enhancing their synergy. (2) We introduce an innovative dynamic co-training framework, which exploits the historical memory of optimal models. Two models iteratively refine pseudo-labels generation under the alternating guidance of current and historical optimal models, enhancing pseudo-labels quality and reliability. (3) Extensive experiments on two microscopic hyperspectral datasets demonstrates that our method outperforms state-of-the-art approaches.

2 Methodology

2.1 Overview

Fig. 1 illustrates the proposed framework based on Virtual Domain Collaborative Learning (VDCL) for microscopic hyperspectral image segmentation. The training process of this framework consists of two core components: (1) introducing virtual domains for effective knowledge transfer and collaborative learning, bridging the gap between labeled and unlabeled data; (2) A dynamic alternative joint training framework is constructed to build a historical optimal model

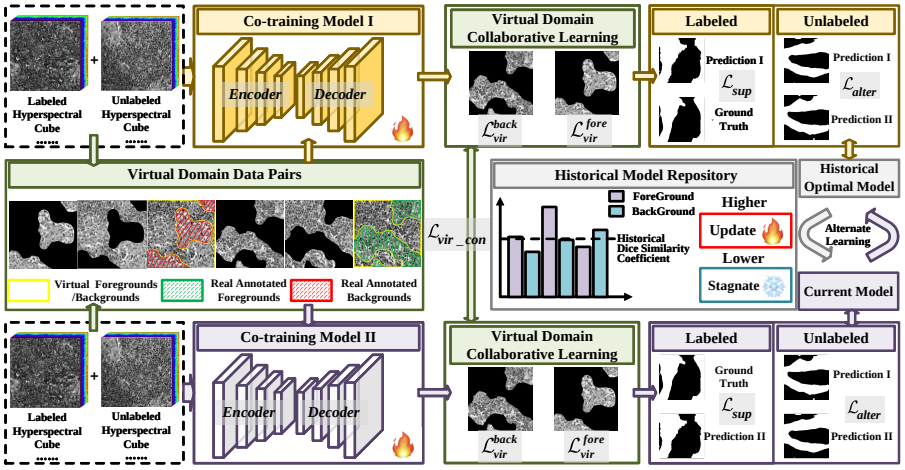

Fig. 1. The overview of our proposed VDCL.

repository, where the current model and the optimal model alternate to guide pseudo-labels generation and improve the quality of pseudo-labels. With U-Net [22] as the backbone for the co-training models, we detail the virtual domain collaborative learning and alternate co-training learning (ACTL) in the following sub-sections.

2.2 Virtual Domain Collaborative Learning

The virtual domain co-training method proposed in this paper aims to enhance the collaborative learning capability between labeled and unlabeled data. Specifically, this method bridges the gap between labeled and unlabeled samples by creating data pairs of virtual backgrounds and real annotated foregrounds $\mathcal{H}_{vir}^{back}$, as well as real backgrounds and virtual foregrounds $\mathcal{H}_{vir}^{fore}$.

$$\mathcal{H}_{vir}^{fore} = \left[\mathcal{H}_{labeled}^{back}, \mathcal{H}_{unlabeled}^{fore}\right], \mathcal{H}_{vir}^{back} = \left[\mathcal{H}_{labeled}^{fore}, \mathcal{H}_{unlabeled}^{back}\right] \tag{1}$$

Based on the construction of these virtual domain data pairs, we can define the corresponding loss functions to measure the performance of the model on the virtual domain. The following two loss functions, $\mathcal{L}_{vir}^{fore}$ and $\mathcal{L}_{vir}^{back}$, are used to guide the model to better learn the information from both labeled and unlabeled data.

$$\mathcal{L}_{vir}^{fore} = \log \frac{\exp\left(x_n, y_n\right)}{\sum_{c=1}^{C} \exp\left(x_{n,c}\right)} \cdot 1\left\{y_n \neq foreground\_index\right\} \tag{2}$$

$$\mathcal{L}_{vir}^{back} = \log \frac{\exp\left(x_n, y_n\right)}{\sum_{c=1}^{C} \exp\left(x_{n,c}\right)} \cdot 1\left\{y_n \neq background\_index\right\} \tag{3}$$

In this way, unlabeled data is not only used to generate pseudo-labels but also collaborates with labeled data for co-learning, significantly improving the utilization of unlabeled data. This approach strengthens the model's learning ability through the construction and mutual transformation of virtual data domains, greatly enhancing the contribution of unlabeled data, especially in semi-supervised learning tasks.

$$\mathcal{L}_{vir\_con} = \frac{1}{n}\sum_{i=1}^{n} \left\| y_1^i - y_2^i \right\|^2 \tag{4}$$

where y_1 and y_2 denote the predictions for the virtual domain data from M_1 and M_2, the two models employed in the co-training framework.

2.3 Alternate Co-training Learning

Generating and refining pseudo-labels is key to utilizing unlabeled data effectively. To achieve this, we construct a repository of optimal models M_2^{hist} based on the DSC (Dice Similarity Coefficient) metric. Models are selected for the repository based on their performance, ensuring that the best-performing models are used to guide the learning process. In this approach, M_2 alternates between being the current model M_2^{curr} and the historically optimal model M_2^{hist}.

$$M_2^{hist} = Repository \left\langle M_2 | DSC_{current} > Best \right\rangle_{update}, \tag{5}$$

$$M_2 = \begin{cases} M_2^{\text{hist}}, & \text{if epoch is odd} \\ M_2^{\text{curr}}, & \text{if epoch is even} \end{cases} \tag{6}$$

To prevent excessive coupling in model learning, we employ alternate co-training between two models, with each guiding the other to refine pseudo-labels and improve the overall learning process. This alternate loss is defined as:

$$\mathcal{L}_{alter} = \mathcal{L}_{ce}\left(y_1, y_2\right) + \mathcal{L}_{dice}\left(y_1, y_2\right) + \mathcal{L}_{ce}\left(y_2, y_1\right) + \mathcal{L}_{dice}\left(y_2, y_1\right) \tag{7}$$

where y_1 and y_2 are the pseudo-label predictions for the unlabeled data made by M_1 and M_2, respectively. $\mathcal{L}_{ce}$ denotes the cross-entropy loss, $\mathcal{L}_{dice}$ is the dice loss.

2.4 Overall Objective

For labeled data, we integrate the Cross-entropy loss and Dice loss to oversee the model training, which is denoted as $\mathcal{L}_{sup}$. Regarding virtual domain data, we calculate the virtual consistency loss $\mathcal{L}_{vir\_con}$ along with the virtual supervision losses $\mathcal{L}_{vir}^{fore}$ and $\mathcal{L}_{vir}^{back}$. In the case of unlabeled data, we determine the alternate learning loss $\mathcal{L}_{alter}$. The total loss function $\mathcal{L}_{total}$ is then given by the formula:

$$\alpha = e^{-0.5*\left(1-\frac{currunt\_epoch}{total\_epoch}\right)^2}, \tag{8}$$

$$\mathcal{L}_{total} = \mathcal{L}_{sup} + \alpha\left(\mathcal{L}_{vir}^{fore} + \mathcal{L}_{vir}^{back} + \mathcal{L}_{vir\_con} + \mathcal{L}_{alter}\right) \tag{9}$$

Table 1. Comparisons with other methods on the MHC Dataset.

Method	Publication	Labeled/unlabeled	DSC(%)↑	IoU(%)↑	HD(mm)↓	HD95(mm)↓
UAMT[16]	[MICCAI'19]		61.73	45.57	44.72	11.86
SSASNet[17]	[MICCAI'22]		63.33	47.18	41.19	10.20
Unimatch[14]	[CVPR'23]		61.29	45.38	44.92	10.94
CPS[19]	[CVPR'21]	20%/80%	62.97	47.54	42.88	11.64
DualTeacher[18]	[NeurIPS'23]		61.57	45.40	42.39	11.20
ABD[21]	[CVPR'24]		61.51	45.89	43.64	12.69
VDCL	-		**64.61**	**48.92**	**39.74**	**9.97**
UAMT[16]	[MICCAI'19]		60.31	44.02	43.48	12.74
SSASNet[17]	[MICCAI'22]		60.93	45.01	44.88	13.59
Unimatch[14]	[CVPR'23]		58.12	42.12	43.13	13.49
CPS[19]	[CVPR'21]	10%/90%	61.13	45.45	43.04	12.90
DualTeacher[18]	[NeurIPS'23]		58.09	41.69	43.85	13.32
ABD[21]	[CVPR'24]		60.50	44.46	43.05	12.41
VDCL	-		**62.20**	**46.33**	**42.82**	**12.22**
UAMT[16]	[MICCAI'19]		59.95	43.78	44.20	13.74
SSASNet[17]	[MICCAI'22]		58.91	42.58	43.75	13.83
Unimatch[14]	[CVPR'23]		57.36	41.12	43.29	13.33
CPS[19]	[CVPR'21]	5%/95%	57.30	42.81	44.13	13.49
DualTeacher[18]	[NeurIPS'23]		53.23	37.12	45.52	14.32
ABD[21]	[CVPR'24]		59.88	43.74	43.08	12.96
VDCL	-		**61.21**	**45.38**	**42.56**	**12.88**

As a weight factor, α balances different loss terms. This allows the model to better utilize information from the virtual domain and unlabeled data, and adapting to training stages.

3 Experiments

3.1 Datasets & Implementation Details

Datasets and evaluation metrics. In this study, experiments are conducted on two datasets: the public Microscopic Hyperspectral Choledoch (MHC) Dataset [24] from Kaggle and the Oral Squamous Cell Carcinoma (OSCC) [7] Dataset. The MHC dataset has 325 scenes with 60 spectral bands in 550-1000 nm and 1280×1024 pixel images (260 for training, 65 for testing). The OSCC dataset includes data from 72 patients, with each scene at 696×520×60 resolution and 400-1000 nm wavelength range (58 for training, 14 for testing). To comprehensively evaluate performance, four metrics are used: Dice similarity coefficient (DSC), Intersection over Union (IoU), Hausdorff distance (HD), and 95% Hausdorff distance (HD95), widely applied in medical image segmentation. DSC is the main metric, with IoU, HD, and HD95 as supplementary ones.

Implementation details. The network in this study is implemented with Pytorch 1.13.0 and CUDA 11.7 on two NVIDIA GeForce RTX3090 24G GPUs. For MHC and OSCC datasets, each original image is divided into 256×256×60

Table 2. Comparisons with other methods on the OSCC Dataset.

Method	Publication	Labeled/unlabeled	DSC(%)↑	IoU(%)↑	HD(mm)↓	HD95(mm)↓
UAMT[16]	[MICCAI'19]		80.35	68.42	33.63	7.73
SSASNet[17]	[MICCAI'22]		80.24	68.60	32.82	6.09
Unimatch[14]	[CVPR'23]		79.69	67.59	33.69	7.23
CPS[19]	[CVPR'21]	20%/80%	81.36	69.58	32.52	6.55
DualTeacher[18]	[NeurIPS'23]		80.72	69.06	33.62	6.42
ABD[21]	[CVPR'24]		81.48	69.88	34.08	7.18
VDCL	-		**82.15**	**71.02**	**31.95**	**5.51**
UAMT[16]	[MICCAI'19]		78.38	65.50	33.65	6.95
SSASNet[17]	[MICCAI'22]		78.84	66.23	35.13	6.66
Unimatch[14]	[CVPR'23]		78.36	65.74	34.08	7.81
CPS[19]	[CVPR'21]	10%/90%	79.39	67.08	33.97	7.88
DualTeacher[18]	[NeurIPS'23]		79.30	67.26	34.37	6.79
ABD[21]	[CVPR'24]		79.84	67.26	34.26	6.58
VDCL	-		**81.38**	**70.01**	**33.63**	**5.66**
UAMT[16]	[MICCAI'19]		75.34	62.13	36.21	7.45
SSASNet[17]	[MICCAI'22]		76.65	63.67	34.94	7.41
Unimatch[14]	[CVPR'23]		74.82	61.57	35.33	7.08
CPS[19]	[CVPR'21]	5%/95%	76.84	64.23	35.65	8.44
DualTeacher[18]	[NeurIPS'23]		75.53	62.36	35.18	7.83
ABD[21]	[CVPR'24]		78.05	65.58	34.88	6.54
VDCL	–		**80.45**	**68.90**	**34.74**	**5.91**

image cubes. Training parameters are: SGD optimizer (momentum 0.99, weight decay 0.0005), batch size 4, 150 epochs, and initial learning rate 0.0001. To ensure fairness, following prior work [23, 25], we use 5%, 10%, and 20% labeled data for model training.

3.2 Comparison Study

Performance on MHC Dataset. As shown in Table 1, we have compared our proposed VDCL method with other existing semi-supervised learning (SSL) methods. The methods are presented under different labeled/unlabeled data ratios, which mimic various realistic scenarios. It can be observed that our VDCL method outperforms the other methods in most cases. For example, when the labeled/unlabeled ratio is 20%/80%, VDCL achieves a DSC of 64.61%, which is higher than the best result (63.33% by SSASNet) among the other methods.

Performance on OSCC Dataset. Table 2 presents the comparative results of our proposed VDCL method and other semi-supervised learning (SSL) methods on the OSCC dataset. The methods are evaluated under different labeled/unlabeled data ratios. It can be seen that our VDCL method outperforms other methods. For example, when the labeled/unlabeled ratio is 20%/80%, VDCL reaches a DSC of 82.15%, IoU of 71.02%, HD of 31.95mm, and HD95 of 5.51mm, which are all better than the best results of other methods. Overall, our VDCL method demonstrates its effectiveness and superiority in handling the semi-supervised learning task on this dataset.

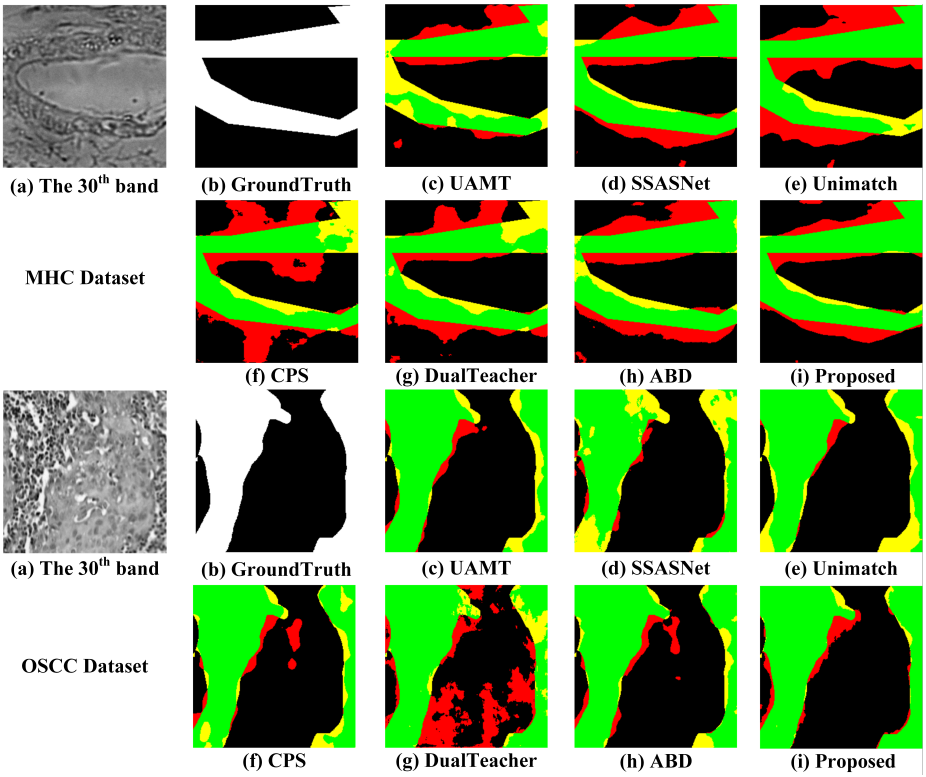

Fig. 2. Visual comparison of segmentation results on two datasets.

Visualization. In Fig. 2, green regions represent accurate cancer area predictions, red regions indicate false positives, and yellow regions signify missed cancer areas. Evidently, our proposed method yields more accurate segmentation masks for cancerous regions than other methods. The segmented boundaries obtained by our approach are sharper on the MHC and OSCC datasets, minimizing false positives and missed cancer areas.

3.3 Ablation Study

Ablation Study of VDAL and ACTL. In this section, we assess the effectiveness of the proposed VDAL and ACTL on the MHC and OSCC datasets. As shown in Table 3, on the MHC dataset with a 20%/80% labeled/unlabeled ratio, the DSC increases from 61.56% to 63.97%. The combined use of VDAL and ACTL results in the best performance, with a DSC of 64.61%, indicating that the two components complement each other. Similarly, on the OSCC dataset with a 20%/80% labeled/unlabeled ratio, VDAL alone boosts the DSC from 80.72% to 81.85%, and ACTL alone increases it to 81.06%. When used together, VDAL

Table 3. Ablation study of our proposed VDAL and ACTL.

Dataset	labeled/unlabeled	VDAL	ACTL	DSC(%)↑	IoU(%)↑	HD(mm)↓	HD95(mm)↓
MHC	20%/80%	-	-	61.56	45.40	42.39	11.20
		✓	-	63.97	48.32	41.42	10.76
		-	✓	63.57	47.99	41.98	11.74
		✓	✓	**64.61**	**48.92**	**39.74**	**9.97**
OSCC	20%/80%	-	-	80.72	69.06	33.62	6.42
		✓	-	81.85	70.74	32.93	6.78
		-	✓	81.06	69.82	33.88	7.12
		✓	✓	**82.15**	**71.02**	**31.95**	**5.51**

Table 4. Ablation study of different paradigms.

Dataset	labeled/unlabeled	Paradigm	DSC(%)↑	IoU(%)↑	HD(mm)↓	HD95(mm)↓
MHC	20%/80%	EMA	63.12	47.66	42.50	11.68
		Co-training	**64.61**	**48.92**	**39.74**	**9.97**
OSCC	20%/80%	EMA	81.39	70.30	32.49	6.26
		Co-training	**82.15**	**71.02**	**31.95**	**5.51**

and ACTL achieve the highest DSC of 82.15%, along with improved IoU, HD, and HD95 values.

Ablation Study of Different Paradigms. We conduct an ablation study on the MHC and OSCC datasets to assess the EMA and Co-training paradigms. For the MHC dataset with a 20%/80% labeled/unlabeled ratio, Co-training achieves a DSC of 64.61% while EMA gets 63.12%. On the OSCC dataset at the same ratio, Co-training attains a DSC of 82.15% compared to EMA's 81.39%. The other metrics exhibit consistent trends, showing that Co-training is more effective in semi-supervised learning for these datasets.

4 Conclusion

In this paper, we propose a semi-supervised method based on Virtual Domain Collaborative Learning (VDCL). It leverages a historical memory of optimal supervision framework where two co-training models generate optimized pseudo-labels to enhance unlabeled data utilization. By creating specific data pairs of virtual and real backgrounds/foregrounds, we enable collaborative learning between labeled and unlabeled samples. Additionally, with a repository of optimal adversarial models and an alternating training approach, the current and historically optimal models jointly guide the training, improving pseudo-label quality. Extensive experiments on two datasets show that our VDCL-based method is superior to previous methods, indicating its great potential in this domain.

Acknowledgments. This work was supported by the National Key Research and Development Program of China under Grant 2024YFF0618204.

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

References

1. Karim, S., Qadir, A., Farooq, U., Shakir, M., Laghari, A.A.: Hyperspectral Imaging: A Review and Trends towards Medical Imaging. Curr Med Imaging 19(5), 417-427 (2022)
2. Bahl, A., Horgan, C.C., Janatka, M., MacCormac, O.J., Noonan, P., Xie, Y., Qiu, J., Cavalcanti, N., F¨urnstahl, P., Ebner, M., et al.: Synthetic white balancing for intra-operative hyperspectral imaging. Journal of Medical Imaging 10(4), 046001–046001 (2023)
3. Calin, M.A., Manea, D., Savastru, R., Parasca, S.V.: Mapping the Distribution of Melanin Concentration in Different Fitzpatrick Skin Types Using Hyperspectral Imaging Technique. Photochem Photobiol 99(3), 1020-1027 (2023)
4. Halicek, M., Lu, G., Little, J.V., Wang, X., Patel, M., Griffith, C.C., El-Deiry, M.W., Chen, A.Y., Fei, B.: Deep convolutional neural networks for classifying head and neck cancer using hyperspectral imaging. Journal of biomedical optics 22(6), 060503–060503 (2017)
5. Q. Huang, W. Li, B. Zhang, Q. Li, R. Tao, and N. H. Lovell.: Blood cell classification based on hyperspectral imaging with modulated gabor and cnn. IEEE Journal of Biomedical and Health Informatics 24(1), 160–170 (2020)
6. W. Zeng, W. Li, M. Zhang, H. Wang, M. Lv, Y. Yang, and R. Tao.: Microscopic hyperspectral image classification based on fusion transformer with parallel cnn. IEEE Journal of Biomedical and Health Informatics 27(6), 2910–2921 (2023)
7. X. Zhang, Q. Li, W. Li, Y. Guo, J. Zhang, C. Guo, K. Chang, and N. H. Lovell.: FD-net: Feature distillation network for oral squamous cell carcinoma lymph node segmentation in hyperspectral imagery. IEEE Journal of Biomedical and Health Informatics 28(3), 1552–1563 (2024)
8. B. Yun, B. Lei, J. Chen, H. Wang, S. Qiu, W. Shen, Q. Li, and Y. Wang.: Spectr: Spectral transformer for microscopic hyperspectral pathology image segmentation. IEEE Transactions on Circuits and Systems for Video Technology 34(6), 4610–4624 (2024)
9. Zeng, L.L., Gao, K., Hu, D., Feng, Z., Hou, C., Rong, P., Wang, W.: Ss-tbn: A semi-supervised tri-branch network for covid-19 screening and lesion segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
10. Lyu, F., Ye, M., Carlsen, J. F., Erleben, K., Darkner, S., Yuen, P. C.: Pseudo-label guided image synthesis for semi-supervised covid-19 pneumonia infection segmentation. IEEE Transactions on Medical Imaging 42(3), 797-809 (2022)
11. Tang, C., Zeng, X., Zhou, L., Zhou, Q., Wang, P., Wu, X., Ren, H., Zhou, J., Wang, Y.: Semi-supervised medical image segmentation via hard positives oriented contrastive learning. Pattern Recognition 146, 110020. (2024)
12. K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. A. Raffel, E. D. Cubuk, A. Kurakin, and C.-L. Li.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. Advances in neural information processing systems, vol. 33, pp. 596–608 (2020)
13. Zheng, M., You, S., Huang, L., Wang, F., Qian, C., Xu, C.: Simmatch: Semi-supervised learning with similarity matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14471–14481 (2022)

14. L. Yang, L. Qi, L. Feng, W. Zhang, and Y. Shi.: Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7236–7246 (2023)
15. 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)
16. 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. pp. 605–613 (2019)
17. Wu, Y., Wu, Z., Wu, Q., Ge, Z., Cai, J.: Exploring Smoothness and Class-Separation for Semi-supervised Medical Image Segmentation. In: MICCAI. pp. 34–43 (2022)
18. J. Na, J.-W. Ha, H. J. Chang, D. Han, and W. Hwang.: Switching temporary teachers for semi-supervised semantic segmentation. Advances in Neural Information Processing Systems 36 (2023)
19. X. Chen, Y. Yuan, G. Zeng, and J. Wang.: Semi-supervised semantic segmentation with cross pseudo supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2613–2622 (2021)
20. Y. Li, X. Wang, L. Yang, L. Feng, W. Zhang and Y. Gao.: Diverse Cotraining Makes Strong Semi-Supervised Segmentor. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16009-16021 (2023)
21. H. Chi, J. Pang, B. Zhang, and W. Liu.: Adaptive Bidirectional Displacement for Semi-Supervised Medical Image Segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4070–4080 (2024)
22. Ronneberger O, Fischer P, Brox T, et al.: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds) MICCAI 2015, Part III 18, pp. 234-241. Springer, Cham (2015)
23. L. Li, S. Lian, Z. Luo, B. Wang, S. Li.: VCLIPSeg: Voxel-wise CLIP-Enhanced model for Semi-Supervised Medical Image Segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 692-701 (2024)
24. Q. Zhang, Q. Li, G. Yu, L. Sun, M. Zhou, and J. Chu.: A multidimensional choledoch database and benchmarks for cholangiocarcinoma diagnosis. IEEE Access 7, 149414–149421 (2019)
25. 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)
26. S. Yun, D. Han, S. Chun, S. J. Oh, Y. Yoo and J. Choe.: CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6022-6031 (2019)
27. V. Olsson, W. Tranheden, J. Pinto and L. Svensson.: ClassMix: Segmentation-Based Data Augmentation for Semi-Supervised Learning. In: 2021 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1368-1377 (2021)