

Open-Set Semi-Supervised Medical Image Classification with Learnable Prototypes and Outlier Filter

Along He¹, Tao Li^{1,2}✉, Yitian Zhao³, Junyong Zhao⁴, and Huazhu Fu⁵

¹ Tianjin Key Laboratory of Network and Data Security Technology, College of Computer Science, Nankai University, Tianjin, China

litao@nankai.edu.cn

² Haihe Lab of ITAI, Tianjin, China,

³ Cixi Institute of Biomedical Engineering, Ningbo Institute of Industrial Technology, Chinese Academy of Sciences, China

⁴ College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing, China.

⁵ Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore, Singapore

Abstract. Semi-supervised learning (SSL) is an effective way to utilize unlabeled data due to the high annotation cost. Unfortunately, the collected unlabeled data will inevitably contain outliers not belonging to the labeled classes in many clinical practices, which is named open-set Semi-supervised learning (OSSL). Although existing methods have achieved decent performance on natural images, they ignore the fine-grained characteristics of medical images, and thus they are not suitable for medical image. In this paper, we propose a framework for the challenging **Open-set Semi-Supervised Classification** for medical image, named OpenSSC and it consists of three components. First, the learnable prototypes is proposed to learn the compact representation of the fine-grained seen classes. Then, we propose a multi-binary discriminator that integrate the closed-set output to distinguish seen and unseen classes. Based on the two components, a joint outlier filter is proposed to classify seen classes and identify unseen classes from unlabeled data. Our proposed method can handle well both the seen and unseen classes. We conduct extensive experiments to demonstrate the superiority of our method, and it outperforms other state-of-the-art OSSL methods on two medical image classification tasks.

Keywords: Semi-supervised learning · Open-set medical image classification

1 Introduction

Medical image classification has been evolving with the development of convolutional neural networks (CNN) [20, 5, 18]. The great success of these efforts has relied on sufficient labeled datasets, which is expensive in medical domain.

Clinically, there is a large amount of unlabeled data that can be exploited and semi-supervised learning (SSL) is a popular technique to utilize these unlabeled data [15, 16, 12]. SSL methods assume that the class space of labeled and unlabeled data is consistent. However, it is difficult to collect such a clean dataset in practice due to the high cost of manual screening. Therefore, it is inevitable that some samples that do not belong to the target classes in the unlabeled data. This setting is called open-set semi-supervised learning (OSSL) [17], which may cause problems of domain shift and class imbalance.

Although SSL is an effective learning approach that is able to utilize unlabeled data, its performances degrade when it faces unseen class data [11]. Therefore, traditional SSL methods [19, 21, 16] are limited in this real-world application scenarios. An intuitive solution is to detect unseen class samples (outliers) and filter out them [14, 17, 2] to reduce the harmful effect of outliers on SSL. However, these methods depend on the accuracy of outlier detection and they may classify some seen-class samples (inliers) as outliers and vice versa, especially for the confusing fine-grained medical image classification task. These methods discard outliers directly, and thus it is difficult to model the features of outliers. Therefore, it will cause low utilization of unlabeled data and has a negative impact on the ability to identify outliers. Thus, some works have been proposed [9, 6, 4, 7] to improve the outliers detection accuracy and utilize unseen unlabeled data effectively. They find that inliers classification performance can be largely improved by considering both high-confidence pseudo-labeled inliers and outliers from unlabeled data. However, they ignore the tiny visual differences among classes in fine-grained medical image classification task, and thus it is easy to confuse inliers and outliers in fine-grained medical OSSL scenarios, resulting in the degradation of performance [14].

In this work, we aim to train a discriminative feature extractor that facilitates the identification of fine-grained inliers and outliers. Addressing the aforementioned challenging problem involves the effective utilization of both inliers and outliers. Therefore, we propose an open-set semi-supervised classification (OpenSSC) framework for fine-grained medical images, and it contains three key components. The first components is the lightweight learnable prototypes to represent fine-grained inliers for each seen class. The learnable prototypes can be trained on labeled inliers and high-confidence pseudo inliers from unlabeled data. Then, the multi-binary discriminator is proposed to recognize whether the image belongs to its ground truth or not. In other words, the multi-binary discriminator treats the class of the image as inlier and the other classes as outlier, and each sub-classifier is only responsible for the current class, while the other classes are treated as negative classes. Finally, the outlier filter based on learnable prototypes and multi-binary discriminator is proposed to form a joint detector to identify inliers and outliers.

We empirically demonstrate that our proposed method enhances the model’s robustness across both fine-grained seen and unseen classes. Our main contributions are: (1) We propose a simple and effective learnable prototype with a few trainable parameters to learn the compact seen-class features, providing a




Fig. 1. The diagram of the proposed framework. It contains three core components, one is the learnable prototypes to learn inliers features and the other is the multi-binary discriminator to recognize inliers and outliers. Based on these two components, we propose a outlier filter strategy to better utilize the seen-class and unseen-class unlabeled data.

precise feature representation for each seen class. It alleviates the feature conflict between inliers and outlier. (2) We introduce a multi-binary discriminator to differentiate between inliers and outliers. Subsequently, an outlier filter strategy is proposed for inliers and outlier detection. Our method can leverage both labeled and unlabeled data to learn feature representation, enhancing the utilization ratio of unlabeled data. (3) We conducted extensive experiments and achieve state-of-the-art (SOTA) results on two medical image classification tasks, which demonstrates the performance of fine-grained OSSL can be improved by properly utilizing inliers and outliers.

2 Method

2.1 Preliminaries and Overview.

The setup of the OSSL is defined as follows. For a C -class classification task, the labeled pairs and unlabeled images are denoted as $\mathcal{D}^l = \{(x_i, y_i)\}_{i=1}^M$ and $\mathcal{D}^u = \{x_j\}_{j=1}^N$, where i and j denote the index of labeled and unlabeled data, respectively. x_i and x_j are the input images, and $y_i \in \{1, \dots, C\}$. For unlabeled data, there exists a subset $\mathcal{D}_{out}^u \subset \mathcal{D}^u$, samples x_j^{out} in $\mathcal{D}_{out}^u$ does not belong to any of the C seen classes, and these unlabeled samples are called unseen-class (outliers) and other samples are called seen-class (inliers). The objective of OSSL is to accurately classify inliers and effectively detect outliers.

We propose a novel open-set semi-supervised classification (OpenSSC) framework for fine-grained medical images, and it is illustrated in Fig. 1. The mean

teacher (MT) [15] is employed as the baseline framework, which consists of a teacher F_ϕ and a student network F_θ , ϕ and θ are the parameters. For the student network, it contains a pre-trained ViT encoder F_e with adapter from AdaptFormer [1], a closed-set classifier F_h , a multi-binary discriminator F_b , a non-linear transformation layer F_g and learnable prototypes P . Except for pre-trained encoder F_e , other parameters are initialized with random weights. Closed-set classifier is to obtain seen-class classification results, and the multi-binary discriminator is used to classify whether the image belongs to specific seen-class or not. Non-linear layer used to transform the high-dimension encoder features to low-dimension prototype features. Formally, the outputs of teacher and student models are denoted as:

$$\hat{y}^s, \hat{y}_b^s, \hat{y}_p^s = F_\theta(x), \hat{y}^t, \hat{y}_b^t, \hat{y}_p^t = F_\phi(x) \quad (1)$$

where x is the input image, s and t denote the student and teacher networks, $\hat{y}$, $\hat{y}_b$ and $\hat{y}_p$ are the closed-set classification results, multi-binary classification results and prototypes classification results, respectively. Note that we do not update the pre-trained ViT encoder and only fine-tune other parameters. The parameters of teacher network are updated by the exponential moving average (EMA) as in MT [15].

2.2 Learnable Prototypes.

The encoder simultaneously learn features for both inliers and outliers, leading to inevitable feature conflicts arising from the shared backbone. Moreover, fine-grained medical image classification task often share significant similarities in features. It is crucial to discern distinctive features while preserving common ones during the learning process. Therefore, we propose the learnable prototypes to learn fine-grained features of inliers.

As Fig. 1 shows, labeled images are fed into the student model to learn the fine-grained seen-class prototypes. The output features of encoder are denoted as $f_i \in \mathbb{R}^{1 \times D}$, then, a fully connected layers followed by ReLU activation F_g is used to transform high dimensional features f_i to inliers specific features $f'_i \in \mathbb{R}^{1 \times d} = F_g(F_e(x_i))$. Then, it can be interacted with the learnable prototypes $P \in \mathbb{R}^{C \times d}$ to learn the inlier prototype output: $\hat{y}_{p,i} = f'_i P^T$, $\hat{y}_{p,i} \in \mathbb{R}^{1 \times C}$. Learnable prototypes are optimized by cross-entropy loss function:

$$\mathcal{L}_p(y_i, \hat{y}_{p,i}) = - \sum_{(x_i, y_i) \in \mathcal{D}_l} y_i \cdot \log(\hat{y}_{p,i}), \quad (2)$$

After training on the labeled data, seen-class prototypes P can learn the compact feature representations of each fine-grained seen class. It can mitigate the negative impact of shared backbone on feature learning of inliers and outlier. Therefore, inliers can produce high confidence scores, while outliers exhibit low confidence scores.

2.3 Multi-binary Discriminator.

The features f_i generated by the encoder can be used to obtain the seen-class output scores $\hat{y}$ through the closed-set classification head F_h . The scores $\hat{y}$ and the features f_i are concatenated, and then they are transformed through the multi-layer perceptron (MLP). Finally, inliers and outliers are discriminated by a multi-binary discriminator F_b as: $\hat{y}_b = F_b(MLP([f_i, \hat{y}]))$, $[,]$ denotes concatenation along the channel dimension, and $\hat{y}_b \in \mathbb{R}^{2 \times C}$. The multi-binary discriminator can be viewed as C binary discriminators. The output of each binary-discriminator is a probability to indicate how likely the image x_i is to be y_i or not. Then, the sub-discriminator is trained such that the data from class y_i are treated as inliers, and the data from the other classes are treated as outliers. It can be seen that binary-discriminator make decisions based on encoder features and outputs scores of the closed-set classifier, which is more conducive to the feature learning of closed-set classifier and backbone. Therefore, the model can more precisely decide whether the current sample belongs to a specific class or outlier. The loss is defined as:

$$\mathcal{L}_{mb}(y_i, \hat{y}_b) = - \sum_{(x_i, y_i) \in \mathcal{D}_l} \sum_{c=1}^C y_{i,c} \cdot \log(\hat{y}_{b,c}) + (1 - y_{i,c}) \cdot \log(1 - \hat{y}_{b,c}), \quad (3)$$

2.4 Outlier Filter.

Following training on labeled data, the closed-set classifier, multi-binary discriminator, and learnable prototypes have acquired discriminative features. For unlabeled data x_j , we can obtain closed-set output score $\hat{y}$, multi-binary discriminator output score $\hat{y}_b$, prototypes classification results $\hat{y}_p$. First, we apply softmax to $\hat{y}_p$ to get the seen-class probability. Then, we produces an outlier filter: $y' = \hat{y}_p \cdot \hat{y}_b$, final confidence score can be obtained by $score = \max(y')$. If $score > \sigma$, the unlabeled sample is inliers and the pseudo label of unlabeled sample is $\tilde{y} = \operatorname{argmax}(y')$, and it is used to train the closed-set classifier and learnable prototypes. Otherwise, the unlabeled sample is outlier, it can be used to train the multi-binary classifier to learn better decision boundary between inliers and outlier. σ is the inliers threshold and it is set to 0.4 (See Table 3).

To train the framework, training loss for labeled data is defined as:

$$\mathcal{L}_{sup} = -\frac{1}{|\mathcal{D}_l|} \sum_{(x_i, y_i) \in \mathcal{D}_l} \mathcal{L}_{ce}(y_i, \hat{y}_i^s) + \mathcal{L}_{mb}(y_i, \hat{y}_{b,i}^s) + \mathcal{L}_p(y_i, \hat{y}_{p,i}^s), \quad (4)$$

where $\mathcal{L}_{ce}$ is the cross-entropy loss. For unlabeled data, the training loss is defined as:

$$\mathcal{L}_{unsup} = \frac{1}{|\mathcal{D}_u|} \sum_{x_j \in \mathcal{D}_u} -y_j^s \cdot \log(y_j^s) + \mathcal{L}_{ce}(\tilde{y}_j, \hat{y}_j^s) + \mathcal{L}_p(\tilde{y}_j, \hat{y}_{p,j}^s) + \mathcal{L}_{mse}(\hat{y}_{p,j}^s, \hat{y}_{p,j}^t) \quad (5)$$

the first term is minimizing entropy, $\mathcal{L}_{mse}$ is the mean square error loss. The final training loss is $\mathcal{L} = \mathcal{L}_{sup} + \lambda \mathcal{L}_{unsup}$, λ is a time-dependent function, and it can be defined as $\lambda(i) = \lambda_{max} \cdot e^{-5(1-\frac{i}{i_{max}})^2}$, i is the current training step, i_{max} is the maximum training steps, and λ_{max} is empirically set to 1.0.

3 Experiments

3.1 Dataset and Evaluation Metrics.

Datasets. ISIC 2018 skin dataset [3] contains 10015, 194, 1512 dermoscopic images in training, validation and test sets for seven-class skin disease classification, including melanoma (MEL), basal cell carcinoma (BCC), actinic keratosis / bowen’s disease intraepithelial carcinoma (AKIEC), benign keratosis (BKL), dermatofibroma (DF), vascular lesion (VASC), melanocytic nevus (NV). The first five classes are considered inliers, and the remaining classes are outliers. DDR dataset [8] includes 13,673 fundus images for diabetic retinopathy (DR) grading, including 6,835 training images, 2,733 validation images and 4,105 test images. These images are graded into six classes: no DR, mild DR, moderate DR, severe DR, proliferative DR and ungradable. The first five classes are considered inliers, and the last class is outliers. Note that seen and unseen classes can be determined according to the target task.

Evaluation Metrics. The classification performance is evaluated on both seen and unseen classes. We regard all unseen classes as a new class, accuracy (Acc) and average recall scores [9] are used as evaluation metrics.

3.2 Implementation Details.

In the experiments, the optimizer is Adam and the weight decay is $1e-3$. The learning rates is 0.0002, which is reduced by the “Poly” strategy. We resized the input images to 320×320 , and the batch size is 46. we adopted horizontal flips, vertical flips and random crops as data augmentation. The pre-trained encoder is ViT-B/16 pre-trained on 400 million image-text pairs [13]. Our method keeps the pre-trained encoder frozen and only fine-tune the other parameters. The dimension d in learnable prototypes is 256. We randomly select 20% or 30% samples from the seen classes as the labeled data and use the remaining training data as the unlabeled data. Code will be available. Training epoch is set to 50 to ensure convergence. Our framework is implemented with Pytorch and all experiments are performed on an NVIDIA GeForce RTX 3090 GPU. Code will be available at <https://github.com/NKUhealong/OpenSSC>.

3.3 Comparisons with SOTA Methods.

To demonstrate the effectiveness of our proposed method, we compare our method against recent SOTA OSSL algorithms, including IOMatch [9], OpenMatch [14], Safe-student [6], SSB [4], T2T [7] and Safe OSSL [10], and the results are shown

Table 1. Comparisons of performance are conducted against other state-of-the-art methods under various seen/unseen class splits and labeled ratios. FS and MT denote fully-supervised and mean teacher.

Method	DDR				ISIC2018			
	5 seen/1 unseen				5 seen/2 unseen			
	30%		20%		30%		20%	
	Acc	recall	Acc	recall	Acc	recall	Acc	recall
Closed-set classification performance								
Openmatch	66.46	52.29	65.34	52.54	75.03	62.30	74.23	60.13
T2T	70.98	57.46	67.35	54.34	77.99	64.47	74.46	62.74
Safe-student	71.56	56.22	66.56	55.12	77.56	64.72	74.23	62.54
Safe OSSL	69.34	56.35	66.93	54.73	78.65	65.02	75.01	63.67
SSB	70.21	57.80	67.64	54.68	78.24	65.74	75.21	63.02
IOMatch	71.48	58.10	68.98	55.45	79.67	65.23	75.35	63.21
OpenSSC	73.13	60.47	70.67	57.84	81.61	67.58	78.57	64.91
Open-set classification performance								
Openmatch	46.34	37.55	44.29	35.24	68.23	50.36	64.32	44.72
T2T	49.76	40.23	46.89	37.88	71.12	53.24	67.46	49.56
Safe-student	48.74	39.57	46.12	37.21	69.23	52.58	67.89	49.64
Safe OSSL	49.24	41.12	46.34	37.13	70.12	52.34	67.41	49.35
SSB	48.39	40.34	47.73	37.76	70.87	52.98	68.34	50.32
IOMatch	50.39	41.78	48.67	38.21	71.41	53.46	68.23	50.63
OpenSSC	52.34	43.29	50.41	40.24	73.88	55.78	70.23	52.88

in Table 1. We compare with the SOTA method in two settings, one is open-set (the test samples contain unseen classes), and the other is closed-set (the test samples only contain seen classes). For closed-set classification, we can see that SOTA methods achieved competitive results and our method achieved the best results with 20% or 30% labeled data, outperforming the previous best results by more than 1.5% in Acc and recall.

For open-set classification, the performance degrades when encountering outliers, OpenMatch [9] achieved poor performance because it has difficulty in detecting outliers similar to inliers as discussed in their paper. The performance of other methods is unsatisfactory as they encounter significant challenges of the fine granularity in medical image classification tasks, often incorrectly identify inliers and outliers. Our method leverage learnable prototypes and multi-binary discriminator to represent each seen class and distinguish inliers and outliers, and then adopt outlier filter to determine whether the test sample is inliers and outliers, achieving the best results with different label ratios and seen/unseen classes splits. All these results demonstrate the effectiveness and robustness of our method in open-set and closed-set medical image classification problem.

Table 2. Ablation studies of our proposed method with 30% labeled data on the two datasets. A, B and C denote the multi-binary discriminator, learnable prototypes and outlier filter, respectively.

index	baseline	A	B	C	DDR		ISIC2018	
					Acc	recall	Acc	recall
E.1	✓				46.25	39.34	69.55	50.37
E.2	✓	✓			49.12	41.02	70.34	52.39
E.3	✓	✓	✓		50.78	41.69	71.45	53.41
E.4	✓	✓	✓	✓	52.34	43.29	73.88	55.78

3.4 Ablation Study.

In this section, we conduct ablation studies to investigate the effectiveness of our proposed method on the two datasets. The baseline is MT [15] and the results are reported in Table 2.

Effectiveness of multi-binary discriminator. For MT baseline, it does not take into account outliers in unlabeled data, and it will confuse inliers with outliers and thus affect the performance. From the results shown in Table 2, after considering multi-binary discriminator (E.2), the model can learn whether the current sample belongs to the specific seen class. Therefore, the model’s ability to recognize all the seen classes can be improved.

Analysis of learnable prototypes. If we only consider the results of multi-binary discriminator, it can not effectively learn the compact feature representation of seen classes. For medical images, they belong to fine-grained classification tasks, where a substantial similarity exists between the seen and unseen classes. It is necessary to learn compact seen-class prototypes. From the experimental results E.3, the proposed learnable prototypes can effectively learn seen-class distribution and distinguish the seen classes and outliers, which proves the effectiveness of learnable prototypes.

Analysis of outlier filter. From the results shown in Table 2, we can see that the single component is not sufficient to improve the performance, and thus we effectively integrate the two to form a joint outlier filter strategy, which can effectively reduce the risk of misclassification. The experimental results E.4 show that our method can effectively improve the Acc and recall. The ablation study reveals that each proposed component has a positive impact on our method.

Analysis of threshold σ . The threshold σ determines whether a sample belongs to inliers or not, and we analyses it in Table 3. We can see that this threshold is robust to different datasets, the model achieves the best results when σ is set to 0.4.

4 Conclusion

In this work, we explored a framework for open-set semi-supervised classification for fine-grained medical images, and proposed learnable prototypes, multi-binary discriminator and outlier filter to identify inliers and detect outliers. Our method

Table 3. Ablation studies of σ with 30% labeled data on the two datasets.

dataset	DDR			ISIC2018		
σ	0.3	0.4	0.5	0.3	0.4	0.5
Acc	51.12	52.34	51.36	72.79	73.88	72.32
recall	42.23	43.29	42.45	55.01	55.78	54.97

can learn compact seen-class features in a end-to-end way, reducing the impact of unseen-class data on the performance and effectively improves the utilization of unlabeled data. Our proposed strategy can be readily combined with different model architectures like CNN and Transformer. The experimental results reveals that it is a promising method in OSSL scenario and outperforms current state-of-the-arts. Moreover, our method is versatile, allowing for application in both open-set and closed-set scenarios during testing and serving as a flexible and plug-and-play framework.

Acknowledgements This work is partially supported by the National Natural Science Foundation (62272248), the Huazhu Fu’s Agency for Science, Technology and Research (A*STAR) Career Development Fund (C222812010) and Central Research Fund (CRF) and the Natural Science Foundation of Tianjin (23JC-QNJ00010).

Disclosure of Interests The authors have no competing interests in this paper.

References

1. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems* **35**, 16664–16678 (2022)
2. Chen, Y., Zhu, X., Li, W., Gong, S.: Semi-supervised learning under class distribution mismatch. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34, pp. 3569–3576 (2020)
3. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., et al.: Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368* (2019)
4. Fan, Y., Kukleva, A., Dai, D., Schiele, B.: Ssb: Simple but strong baseline for boosting performance of open-set semi-supervised learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 16068–16078 (2023)
5. He, A., Li, T., Li, N., Wang, K., Fu, H.: Cabnet: Category attention block for imbalanced diabetic retinopathy grading. *IEEE Transactions on Medical Imaging* **40**(1), 143–153 (2020)
6. He, R., Han, Z., Lu, X., Yin, Y.: Safe-student for safe deep semi-supervised learning with unseen-class unlabeled data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14585–14594 (2022)

7. Huang, J., Fang, C., Chen, W., Chai, Z., Wei, X., Wei, P., Lin, L., Li, G.: Trash to treasure: Harvesting ood data with cross-modal matching for open-set semi-supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8310–8319 (2021)
8. Li, T., Gao, Y., Wang, K., Guo, S., Liu, H., Kang, H.: Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Information Sciences* **501**, 511–522 (2019)
9. Li, Z., Qi, L., Shi, Y., Gao, Y.: Iomatch: Simplifying open-set semi-supervised learning with joint inliers and outliers utilization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15870–15879 (2023)
10. Ma, Q., Gao, J., Zhan, B., Guo, Y., Zhou, J., Wang, Y.: Rethinking safe semi-supervised learning: Transferring the open-set problem to a close-set one. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16370–16379 (2023)
11. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems* **31** (2018)
12. Peng, Z., Tian, S., Yu, L., Zhang, D., Wu, W., Zhou, S.: Semi-supervised medical image classification with adaptive threshold pseudo-labeling and unreliable sample contrastive loss. *Biomedical Signal Processing and Control* **79**, 104142 (2023)
13. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
14. Saito, K., Kim, D., Saenko, K.: Openmatch: Open-set consistency regularization for semi-supervised learning with outliers. *arXiv preprint arXiv:2105.14148* (2021)
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. Yang, L., Zhuo, W., Qi, L., Shi, Y., Gao, Y.: St++: Make self-training work better for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4268–4277 (2022)
17. Yu, Q., Ikami, D., Irie, G., Aizawa, K.: Multi-task curriculum framework for open-set semi-supervised learning. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16. pp. 438–454. Springer (2020)
18. Yuan, L., Liu, X., Yu, J., Li, Y.: A full-set tooth segmentation model based on improved pointnet++. *Visual Intelligence* **1**(1), 21 (2023)
19. Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems* **34**, 18408–18419 (2021)
20. Zhang, J., Xie, Y., Xia, Y., Shen, C.: Attention residual learning for skin lesion classification. *IEEE transactions on medical imaging* **38**(9), 2092–2103 (2019)
21. 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)