Pattern Recognition 176 (2026) 113204



Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr





Amplitude-guided deep reinforcement learning for semi-supervised layer segmentation

Enting Gao [a,b], Zian Zha [a], Yonggang Li [c], Junhui Zhu [a], Yong Wang [a], Xinjian Chen [d], Naihui Zhou [e,*], Dehui Xiang [d,*]

[a] School of Electronic and Information Engineering, Suzhou University of Science and Technology, Suzhou, 215009, Jiangsu, China
[b] Jiangsu Key Laboratory of Biophotonics, Suzhou City University, Suzhou, 215009, Jiangsu, China
[c] Department of Radiology, The First Affiliated Hospital of Soochow University, Suzhou, 215006, Jiangsu, China
[d] School of Electronic and Information Engineering, Soochow University, Suzhou, 215006, Jiangsu, China
[e] Department of Dermatology,The First Affiliated Hospital of Soochow University, Suzhou, 215006, Jiangsu, China

ARTICLE INFO

Keywords:
Deep reinforcement learning
Semi-supervised layer segmentation
Phase component alignment
Cross-power spectrum correlation

ABSTRACT

Accurate segmentation of scalp tissue layers is essential for mechanistic studies and staging of androgenetic alopecia (AGA), a common form of hair loss that impacts quality of life and mental health. High-resolution magnetic resonance imaging (HR-MR) offers a promising assessment tool. However, accurate segmentation remains challenging due to the lack of large-scale annotated datasets, structural deformation, and low image quality. To address these issues, an Amplitude-guided Deep Reinforcement Learning (ADRL) framework is designed to decouple the data distribution of images and adaptively fuse into the distribution of unlabeled images. This enables effective feature learning of lamellar and asymmetrically thickened structures from both labeled and unlabeled data. Then, phase component alignment (PHA) is imposed to mitigate the adverse impacts of noise or artifacts. To further enhance the discriminative capability of this network, a Cross-Power Spectrum Correlation (CPSC) module is proposed to mitigate inaccurate segmentation of layer structures. Comprehensive experiments on a scalp HR-MR image dataset and a publicly available retinal OCT image dataset demonstrate that our method significantly outperforms state-of-the-art methods in semi-supervised layer segmentation.

1. Introduction

Androgenetic alopecia (AGA), the most common form of hair loss, has an overall incidence of 33% in European American men [1] and a lower incidence in women and Asians. In China, the prevalence is about 21.3% in men and 6.0% in women [2]. The pathogenesis of AGA involves multiple factors, including androgen level, genetics, gender, age and mental stress [3], which significantly impair the quality of life and mental health of affected individuals.

Recent studies [4] have demonstrated that High-Resolution Magnetic Resonance imaging (HR-MR) can enhance the understanding of the underlying pathogenesis and staging of AGA by providing clear visualization of scalp tissue structures. Although deep convolutional neural networks have achieved significant success in semantic segmentation, these methods strongly rely on having as much high quality labeled training data as possible to achieve considerable performance. However, accurately segmenting such layer structures remains a challenge in medical image segmentation tasks. There are three main reasons. **1) Manually Annotated Large-scale Images.** Manually labeling layers is extremely time-consuming, labour-intensive and expensive. As shown in Fig. 1, scalp or retina are extremely complex. Scalp in HR-MR images often can be segmented at least 6 layers, such as stratum corneum and stratum granulosum (SCG), stratum spinosum and stratum basale (SSB), papillary layer (PL), reticular layer (RL), subcutaneous fat layer containing hair follicles (SFL-HF) and subcutaneous fat layer without hair follicles (SFL-NHF). Retina in OCT images often can be segmented at least 8 layers, such as Retina nerve fiber layer (RNFL), Ganglion cell layer and inner plexiform layer (GCL + IPL), Inner nuclear layer (INL), Outer plexiform layer (OPL), Outer nuclear layer (ONL), Inner photoreceptor segments (IS), Outer photoreceptor segments (OS), and Retinal pigment epithelium (RPE). In addition, the numbers of HR-MR B-scan images and OCT B-scan images often be large. **2) Anatomical Variability.** Anatomical variability, both across individuals and within different regions of the same subject, introduces considerable differences in

* Corresponding authors.
E-mail addresses: entinggao@mail.usts.edu.cn, entinggao@foxmail.com (E. Gao), 2313041001@post.usts.edu.cn (Z. Zha), yonggang224@163.com (Y. Li), zhujh@mail.usts.edu.cn (J. Zhu), wangyong@mail.usts.edu.cn (Y. Wang), xjchen@suda.edu.cn (X. Chen), zhounaihui@163.com (N. Zhou), xiangdehui@suda.edu.cn (D. Xiang).

https://doi.org/10.1016/j.patcog.2026.113204
Received 8 June 2025; Received in revised form 18 December 2025; Accepted 28 January 2026
Available online 30 January 2026
0031-3203/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

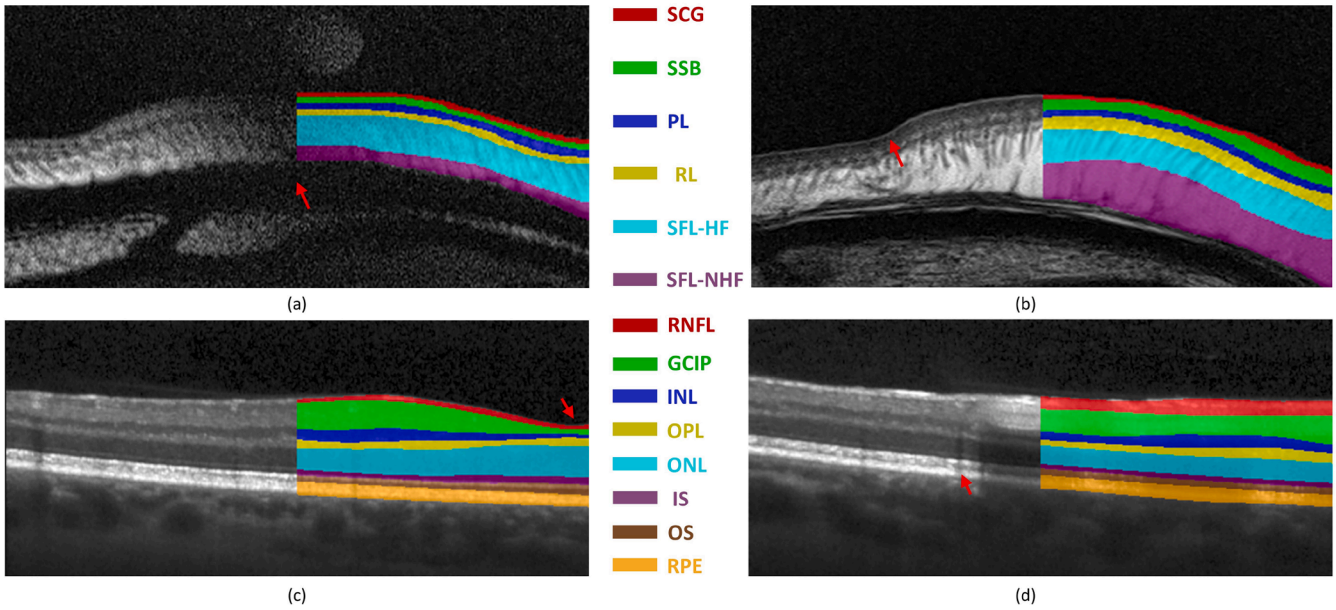


Fig. 1. Examples of scalp HR-MR images and retinal OCT images. The vertex (a) and occipital (b) scalps both exhibit six-layer anatomical structures in HR-MR images.

tissue layer thickness, shape, and arrangement, thereby increasing the complexity of the segmentation. For example, during HR-MR image acquisition of scalp, coils are used to stabilize the head, which may squeeze and deform the layer structures (as indicated by arrow in Fig. 1(b)). **3) Low Image Quality.** Noise or imaging artifacts often degrade image quality. Vascular (as indicated by arrow in Fig. 1(d)) and susceptibility artifacts can degrade image quality and speckle noise may obscure subtle tissue boundaries (as indicated by arrow in Fig. 1(c)). In addition, low inter-layer contrast–characterized by minimal signal differences between adjacent tissue layers–often results in indistinct boundaries (as indicated by the red arrow in Fig. 1(a)), posing significant challenges for accurate delineation. The extreme thinness and structural complexity of certain layers further hinder conventional segmentation methods from accurately capturing their contours (as indicated by arrow in Fig. 1(c)).

To effectively learn global-local dependency features, [5] presents a deep feature-fusion network that jointly models morphological, semantic, and local-global dependencies for accurate boundary-aware segmentation. Zhao et al. [6] combines hybrid global-local feature extraction and dynamic adaptive upsampling to achieve accurate and computationally efficient segmentation. Nevertheless, these fully supervised approaches require ample training data.To deal with these challenges, many semi-supervised methods have focused on utilization of unlabeled data. It can generally be divided into two categories: pseudo-label methods [7] and consistency regularization [8]. While most semi-supervised methods seek to take advantage of unlabeled data, they are often constrained by the absence of effective data augmentation strategies. Although the methods in [9] and [10] mitigate overfitting and improve data diversity to some extent by adopting CutMix [11], its stochastic and non-adaptive nature makes it difficult to ensure consistent effectiveness in complex scenarios and may even introduce undesirable side effects. For unlabeled data, pseudo-label methods usually generate pseudo-labels by pre-training a model with labeled data, but improvements are often limited. Consistency regularization can be applied to model [12], data [13] and task [14] to enhance robustness and reduce sensitivity to noise. However, both strong and weak perturbation consistency and the design of multiple encoders or decoders ignore that the distribution of random noise differs from real noise.

To address these issues, we introduce our **A**mplitude-guided **D**eep **R**einforcement **L**earning (ADRL) framework (as shown in Fig. 2). First, instead of relying on CutMix [11], which randomly exchanges image content, deep reinforcement learning (DRL) is proposed to further decouple the data distribution of images and adaptively transform it into the distribution of unlabeled images. Second, departing from conventional consistency strategies, the phase component enforces consistency and mitigates the adverse impacts of noise. Third, rather than employing pairwise similarities, the cross-power spectrum is proposed to impose constraints. Overall, the main contributions of this work are summarized as follows,

- An Amplitude-guided DRL Data Augmentation (ADDA) strategy is designed by integrating Fourier transform into a deep reinforcement learning framework. The proposed strategy aims to adaptively augment the labeled images with confident labels with different unlabeled images, such that the segmentation model can be improved by the labeled images with diverse styles from unlabeled images.
- A Phase Alignment (PHA) strategy is introduced to guide the model's attention toward structural components in the image, effectively mitigating the impact of noise and artifacts. The phases of unlabeled images are used to maintain the structural consistency among original, weakly and strongly augmented images, since phase components of features encode the structural components of the image.
- A Cross-Power Spectrum Correlation (CPSC) module is proposed to capture global constraints for mitigating inaccurate segmentation of layer structures. The cross-power spectrum is proposed to efficiently capture the spatial shift between the outputs of two different classifiers by identifying its peak. Correlation map is then constructed to modulate the high-level features for the alignment of the outputs of two different classifiers.
- Comprehensive experiments on a scalp HR-MR image dataset and a publicly available retinal OCT image dataset demonstrate that our method significantly outperforms state-of-the-art methods in semi-supervised layer segmentation.

2. Related works

2.1. Layer segmentation

Many layer segmentation methods have been proposed for medical images. Traditional energy-driven methods such as active contour models [15], Markov random field models [16], and graph models [17]

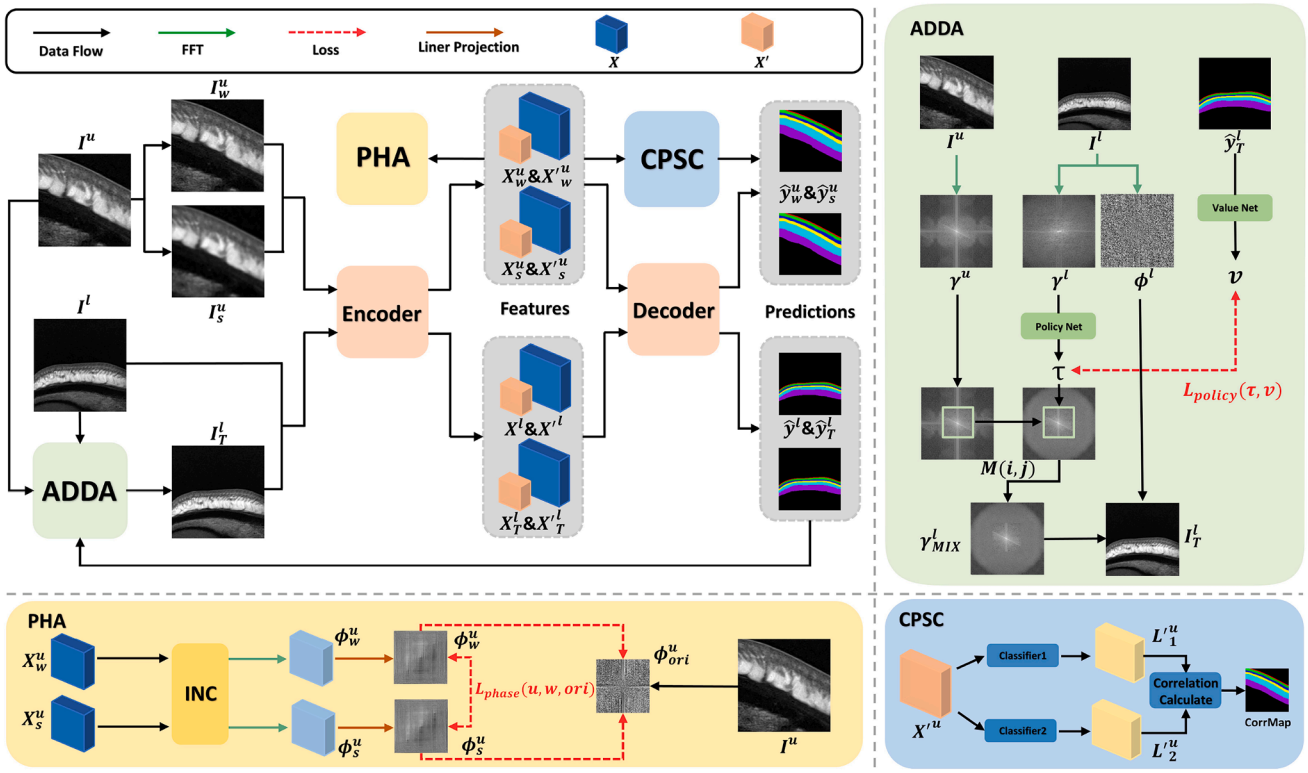


Fig. 2. The proposed ADRL framework. Weak and strong augmentation are applied to unlabeled image I^u to produce I^u_w and I^u_s, respectively. The encoder processes I^u_w and I^u_s to generate features X^u_w and $X^u_s \in \mathbb{R}^{C\times\frac{H}{4}\times\frac{W}{4}}$, while X'^u_w and $X'^u_s \in \mathbb{R}^{C\times\frac{H}{8}\times\frac{W}{8}}$, respectively. The PHA module and CPSC module are applied to the features of unlabeled images while the ADDA module is applied to I^l. The augmented image I^l_T shares the same encoder to get prediction $\hat{y}^l_T$ and completes a cycle of deep reinforcement learning.

exhibited show limited robustness when dealing with complex or low-quality medical images. Graph search-based methods represented each pixel as a node in a graph, where edge weights between neighboring nodes were computed based on the pixel intensities. For example, Xiang et al. [17] transformed the surface segmentation problem into a graph-theoretic framework. Xiang et al. [18] proposed a hybrid graph live wire algorithm to accurately detect surfaces between retinal layers in OCT images with fluids. However, graph-based layer segmentation in pathological retinas remains limited, primarily due to disease-induced alterations in tissue thickness, morphology, and spatial arrangement.

With the rise of deep learning, many researchers have turned to convolutional neural networks (CNNs) to overcome limitations. Early CNN-based methods extracted features of specific layer boundaries and trained classifiers for surface estimation [19]. To enhance feature representation, subsequent approaches integrated attention mechanisms or combined CNNs with Transformers to capture both local details and global context. For example, lightweight encoder-decoder networks employed multi-scale feature extraction and Transformer blocks to improve semantic understanding [20]. Despite these advances, fully supervised methods remain limited by the difficulty of obtaining large-scale manually annotated datasets.

2.2. Semi-supervised segmentation

Semi-supervised semantic segmentation reduced annotation costs by leveraging limited labeled data alongside abundant unlabeled data while maintaining robust performance. Existing approaches primarily followed two strategies: pseudo labeling and consistency learning. Pseudo labeling methods generated proxy labels for unlabeled images, often refining them iteratively or employing confidence-based selection to improve reliability. For example, ComWin [21] dynamically selected the most confident pseudo labels across multiple models, while DiffTCL [22] used a pretrained diffusion model for early pseudo-label generation.To further improve pseudo-label reliability, DOC-MLE [23] introduced a dynamic orthogonal constraint to enhance pixel diversity and a multi-level election strategy to generate robust pseudo-labels by jointly considering pixel-level and feature-level information. Nevertheless, pseudo labels may be noisy or misleading, especially in complex datasets.

To mitigate label noise, consistency learning encouraged models to produce stable predictions under various perturbations. Representative methods include MC-Net+ [24], which enforced output consistency across multiple decoders, MaskMatch [25], which introduces masked local consistency, and LeFeD [26], which leveraged feature-level discrepancy for semi-supervised learning. Some approaches, such as BE-SemiNet [27], further combined task-level and data-level consistency to enhance robustness.

Despite these advances, existing methods often failed to fully exploit global structural information, which was crucial for tasks like layer segmentation where features spanned the entire image. Methods such as DiffTCL [22] and BE-SemiNet [27] lacked explicit designs for global constraints, limiting performance under noise and unclear boundaries. To this end, we take advantage of the decoupling property of the Fourier transform and construct a consistent learning strategy in phase, which avoids the influence of extreme noise while reducing the amount of calculation. In addition, we adopt cross-power spectrum to construct global constraint consistency, which can effectively avoid outliers and unclear boundaries.

2.3. Deep reinforcement learning

Recently, deep reinforcement learning has achieved notable success across various fields [28], showing strong performance in both natural image segmentation [29] and medical image segmentation [30]. BS-IRIS [31] introduced a multi-agent reinforcement learning (MARL) approach for interactive 3D medical image segmentation, where each voxel acted as a collaborative agent guided by a shared policy. Xu et al. [32] developed a two-stage framework for pancreas segmentation in CT images by utilizing a deep Q network (DQN) to localize the pancreas and a deformable U-Net to adapt to adapt to its non-rigid geometry for accurate segmentation. AMP-DRL [33] formulated mask selection as a reinforcement learning task, using a dueling DQN agent to adaptively choose mask position and size based on reconstruction loss feedback and learn fine-grained, image-specific representations. Similarly, AHM [34] employed an A3C based strategy to enhance mask selection by choosing suitable and challenging masks.

The special strategy mechanism of deep reinforcement learning enables it to choose the best path according to different situations, which coincides with the difficulties faced by semi-supervised learning. Unlike all of the above methods, which require a static environment to model the true reward value, deep reinforcement learning is applied to a dynamic environment to build a more flexible policy. Unlike any traditional perturbation mechanism, our ADDA module is able to dynamically implement nonlinear transformations based on the model state, and we have demonstrated through a series of experiments that this method can significantly improve the model training.

3. Method

In semi-supervised segmentation task, a labeled dataset is denoted as $D^l = (I_s^l, Y_s^l)_{s=1}^{N_l}$ and an unlabeled dataset is denoted as $D^u = (I_s^u)_{s=1}^{M_u}$. $I_s^l \in \mathbb{R}^{H\times W}$ is the labeled image, and H and W denote the height and width of I_s^l. Y_s^l is the label of I_s^l. Similarly, $I_s^u \in \mathbb{R}^{H\times W}$ denotes an unlabeled image. Usually, the number of labeled images N_l is significantly smaller than the number of unlabeled images M_u, *i.e.*, $M_u > N_l$.

In the task of semi-supervised medical image segmentation, the utilization of labeled data is very important due to its scarcity. Therefore, data augmentation has become a very important step to improve the performance of the model. However, most methods focused on the consistency of segmentation results or constrains, which is not suitable for images with various noise and tissues. Therefore, we exploit a frequency-domain based deep reinforcement learning strategy. Fig. 2 illustrates the architecture of our proposed Amplitude-guided Deep Reinforcement Learning (ADRL) network, which integrates three key modules: Amplitude-guided DRL Data Augmentation (ADDA), Phase Alignment (PHA), and Cross-Power Spectrum Correlation (CPSC).

3.1. Amplitude-guided DRL data augmentation

CutMix [11] is a common data augmentation method, which makes the model more robust by mixing samples. Most methods often focus on the usage of unlabeled data and ignore the importance of data augmentation, just use the conventional CutMix method. However, the common CutMix method is not suitable for HR-MR images, especially for layer segmentation tasks. It tends to destroy the continuity of the layer structure, and the data after random exchange may not have a positive impact on the model, or even a side effect. In order to build an adaptive method that can choose the appropriate amount of data augmentation for different situations, we introduce our asynchronous advantage actor-critic (A3C) [35] algorithm. Unlike any traditional perturbation mechanism, our ADDA module is able to dynamically implement nonlinear transformations based on the model state for data augmentation.

ADDA consists of two networks: the policy network and the value network. The policy network encourages agents to behave better, and the value network scores more accurately based on state. I_i^l is the ith $(i = 1, \dots, N)$ pixel of the labeled image I_s^l with $N = H \times W$. Each pixel I_i^l in the labeled image is treated as an agent whose policy is denoted as $\pi_i(a_i^{(t)} \mid s_i^{(t)})$, t stands for iteration. For I_i^l at the iteration t, $a_i^{(t)} \in \mathcal{A}$ and $s_i^{(t)} \in \mathcal{S}$ denote the action and state, where $\mathcal{A}$ and $\mathcal{S}$ are the action space and the state space, respectively. The goal of ADDA is to learn the optimal policy $\pi = (\pi_1, \dots, \pi_N)$ that maximizes the mean of the expected total discounted rewards for all pixels as

$$\pi^* = \arg\max \; E_\pi(\bar{r}) \tag{1}$$

$$\bar{r} = \sum_{t=0}^{T} \frac{1}{N} \sum_{i=1}^{N} r_i^{(t)} \tag{2}$$

where T is the total number of iterations, $\bar{r}$ denotes the sum of all the mean rewards.

Through this DRL strategy, the model adaptively learns the nonlinear mapping from the labeled image to the actions. The actions of DRL are used to manipulate adaptive style transfer from unlabeled image to labeled image. Fast Fourier transform (FFT) can decouple the style and structure components of labeled images and unlabeled images into the amplitude and the phase. The amplitude contains the style component and the phase contains the structure component. The change of amplitude can enable the model to learn the different nonlinear mappings from the labeled image to the actions, and the change does not affect the structural component of the labeled image since the structural component corresponds to the label. Therefore, a amplitude-guided deep reinforcement learning strategy is proposed to adaptively fuse the style of labeled images with diverse styles of numerous unlabeled images. FFT is therefore applied to both the unlabeled image I^u and the labeled image I^l. The amplitude spectra γ^u and γ^l are derived from I^u and I^l, respectively, while the phase of the labeled image, denoted ϕ^l, is preserved for subsequent reconstruction. The style is typically represented in low-frequency components, while edge is often found in high-frequency components [36]. The labeled image with the style transferred from the unlabeled image can be obtained by exchanging the low-frequency components of the labeled and unlabeled images. Different from the way that MiDss [37] increases the swap ratio linearly, our A3C strategy is adopted to adaptively obtain the value $\tau \in (0, 1)$ to control the size of the swap region. We define the width $\mathcal{W}$ of the swap region $\mathcal{M}$ ($\mathcal{M}$ is a binary map, in which the region used for swapping is set to 0 and the remaining parts are set to 1) as

$$\mathcal{W} = \lfloor \tau \cdot H \rfloor \tag{3}$$

The policy network P is designed to generate the action a, and the value network V is designed to generate the value v. The adaptive scale factor τ at iteration t can be defined as

$$\tau^{(t)} = \sigma[\frac{1}{N} \sum_{i=1}^{N} a_i^{(t)}(\gamma_i^l)] \tag{4}$$

where σ is a softmax function. Based on Eq. (3), the synthetic image $\hat{I}^l$ can be formed through Inverse Fourier transform (IFFT) as

$$\hat{I}^l = real[\mathcal{F}^{-1}(\gamma^l \odot (1 - \mathcal{M}) + \gamma^u \odot \mathcal{M}, \phi^l)] \tag{5}$$

where $\odot$ is the element-wise multiplication. The prediction $\hat{y}^l$ of $\hat{I}^l$ is the output of model. Then, we need to design the reward in a suitable way to guide the agent towards the desired goal. Since the prediction is formed from the adaptive scale factor τ, it can be viewed as the dynamic state s that reflects the impact and changes that the policy brings to the environment. By using the cross entropy loss between ground truth Y^l and prediction $\hat{y}^l$, our model is driven to learn the potential non-linear relationship dynamically. To this end, the cross entropy loss is used to define the reward as

$$r = \frac{1}{N} \sum_{i=1}^{N} \mathcal{L}_{\text{trans}}(Y_i^l, \hat{y}_i{}^l) \tag{6}$$

where Y_i^l and $\hat{y}_i{}^l$ respectively denote ground truth and the prediction of model at pixel i, $\mathcal{L}_{\text{trans}}$ is cross entropy loss. As mentioned above, the

value network V is supposed to evaluate the true reward value as closely as possible from state s, which is the prediction $\hat{y}^l$. The parameters of the value network V is updated as

$$d\theta_v = \nabla_{\theta_v}(r - V(s)^2) \tag{7}$$

The policy network P is supposed to make actions with higher rewards, and the gradient for updating the parameters of the policy network θ_p is supposed to compute as

$$A(a,s) = r - V(s) \tag{8}$$

$$d\theta_p = -\nabla_{\theta_p} log\pi(a|s)A(a,s) \tag{9}$$

where $A(a,s)$ is the advantage function. It represents the advantage of taking an action in the current state relative to the average expectation. Since $\mathcal{L}_{\text{trans}}$ is adopted as our reward, unlike normal DRL problem, the 'reward' naturally declines, and we aim for a faster decline. Therefore, it is encouraged when $A > 0$ rather than $A < 0$. This turns our problem into building a model to guide the policy network updates under dynamic reward reduction, allowing $\mathcal{L}_{\text{trans}}$ to converge faster. Instead of the gradient update, the loss function is used to express our version of policy update function as

$$\mathcal{L}_{\text{policy}} = exp[\log(\tau) \cdot exp(A(a,s))] \tag{10}$$

Considering the characteristics of exponential function, $\mathcal{L}_{\text{policy}}$ can make the strategy network make better strategies to improve the accuracy of the model. Similarly, the value loss is defined as

$$\mathcal{L}_{\text{value}} = \left(\mathcal{L}_{\text{Trans}} - v\right)^2 \tag{11}$$

v represents the agent's expectation of the reward that the current action can obtain. Therefore, v should be aligned as closely as possible to the actual reward value. Then the total loss of ADDA module could be written as

$$\mathcal{L}_{\text{ADDA}} = \lambda_1\mathcal{L}_{\text{policy}} + \lambda_2\mathcal{L}_{\text{value}} + \lambda_3\mathcal{L}_{\text{Trans}} \tag{12}$$

where λ_1, λ_2 and λ_3 are the hyper parameters.

3.2. Phase alignment

In the task of layer segmentation, noise and artifacts significantly hinder performance, particularly for traditional consistency learning methods that primarily operated in the image domain. To overcome these obstacles, a phase consistency strategy is introduced by driving the model's focus on the structural components, to effectively mitigate the impact of noise and artifacts. Furthermore, multi-scale large-kernel convolutions [38] are employed to enhance the focus of features on global layer structures, such that global characteristics are preserved and fine details are maintained.

Low-level feature for unlabeled image with weak and strong augmentation $\mathbf{X}_w^u$ and $\mathbf{X}_s^u$ are fed into the multi-scale Inception-NeXt Convolution (INC) block [38] to get the filtered feature. FFT along with a liner projection is applied to get the phase ϕ_w^u and ϕ_s^u. The phase alignment loss can be defined as

$$\mathcal{L}_{\text{phase}}^{\mathbf{X}} = \frac{1}{N}\sum_{i=0}^{N-1}\left(\phi_{i_w}^u - \phi_{i_s}^u\right)^2 \tag{13}$$

The structure of high-level feature is also supposed to be similar to the original image, and therefore the phase alignment loss between weakly augmented image and the original image can be defined as

$$\mathcal{L}_{\text{phase}}^{\text{ori}} = \frac{1}{N}\left(\sum_{i=0}^{N-1}\left(\phi_{i_w}^u - \phi_i^{\text{ori}}\right)^2 + \sum_{i=0}^{N-1}\left(\phi_{i_s}^u - \phi_i^{\text{ori}}\right)^2\right) \tag{14}$$

where $\phi_{i\,ori}$ denotes the phase of the unlabeled image without augmentation while ϕ_w^u and ϕ_s^u denote the phase of the unlabeled image with weak and strong augmentation. The total loss of the phase alignment module can be written as

$$\mathcal{L}_{\text{phase}} = \lambda_4 \cdot \mathcal{L}_{\text{phase}}^{\mathbf{X}} + \lambda_5 \cdot \mathcal{L}_{\text{phase}}^{\text{ori}} \tag{15}$$

where λ_4 and λ_5 are the hyper parameters. These methods make sure the encoder focus on the phase with high frequency components for the boundary. It's benefit for the model to capture details of thin layers.

3.3. Cross-power spectrum correlation

Inspired by the correlation extraction method in CorrMatch [39], we have noticed the potential of embedding global pairwise similarities in features for layer segmentation tasks. Given that layer structures are continuous and span the entire image, global constraints are crucial for optimizing segmentation results and correcting misclassified pixels. To effectively capture these global constraints, we propose the use of correlation maps. Unlike pairwise similarity approach [39], our method employs the cross-power spectrum, a computationally efficient signal processing tool commonly used in image registration [40], to align predictions and construct a correlation map that prioritizes global consistency. This approach is particularly advantageous for layer segmentation, where maintaining the overall structure is more critical than local details. Therefore, a CPSC module is introduced to capture global constraints and enhance segmentation through prediction alignment via the cross-power spectrum.

Feature X'^u of an unlabeled image is extracted from the encoder (as shown in Fig. 2 and processed through two different classifiers to generate logits L'^u_1, L'^u_2. FFT is also applied to get the frequency $\mathcal{F}'^u_1$ and $\mathcal{F}'^u_2 \in \mathbb{R}^{C\times H\times W}$, where C is the number of classes. The normalized cross-power spectrum can be computed as

$$\mathbf{R} = \frac{\mathcal{F}'^u_1\mathcal{F}'^{u*}_2}{|\,\mathcal{F}'^u_1\mathcal{F}'^{u*}_2\,| + \varepsilon} \tag{16}$$

where $\mathcal{F}'^{u*}_2$ is the complex conjugate of $\mathcal{F}'^u_2$and ε is a small positive number. IFFT is applied to $\mathbf{R}$ to transform into the spatial domain as

$$\mathbf{R}_r = real(\mathcal{F}^{-1}(\mathbf{R})) \tag{17}$$

The real part is extracted to form the correlation map. The optimal spatial shift (d_x, d_y) is derived by identifying the peak in $\mathbf{R}_r$ as

$$(d_x, d_y) = argmax_{h,w}\,|\,\mathbf{R}_r(h,w)\,| \tag{18}$$

This step quantifies the misalignment between logits L'^u_1 and L'^u_2. The logits L'^u_2 are shifted to align with L'^u_1 as

$$L'^u_{2\,shifted} = \mathbf{S}(L'^u_1, d_x, d_y) \tag{19}$$

where $\mathbf{S}$ denotes a cyclic shift operation. The correlation map is constructed by averaging the softmax outputs of the aligned and original logits as

$$\mathbf{C} = mean(\sigma(L'^u_1) + \sigma(L'^u_{2\,shifted})) \tag{20}$$

where σ is the softmax function applied along the class dimension. The correlation map $\mathbf{C}$ modulates the interpolated high-level features out $L^u \in \mathbb{R}^{C\times H\times W}$ as

$$Corr_{out} = \mathbf{C} \odot L^u \tag{21}$$

Therefore, the correlation loss $\mathcal{L}^u_{corr}$ can be computed between the constrained prediction $y^u_c = \sigma(Corr_{out})$ and the original feature prediction $y^u = \sigma(L^u)$ as

$$\mathcal{L}^u_{corr} = -\frac{1}{N}\sum_{i=1}^{N} y_{i_c}{}^u \log y_i{}^u \tag{22}$$

3.4. Loss function

The supervised loss $\mathcal{L}_s$ is computed between the prediction of I^l and label y using cross-entropy loss. The semi-supervised loss $\mathcal{L}_{semi}$ is computed between the prediction of unlabeled images with two different

augmentations I_w^u and I_s^u using cross-entropy loss. Kullback-Leibler Divergence loss [41] is also used to measure the consistency between the logits of weakly and strongly augmented images as

$$\mathcal{L}_{\text{kl}} = \frac{1}{N}\sum_{i=0}^{N-1} \sigma(y_{i_w}^u) \log\left(\frac{\sigma(y_{i_w}^u)}{\sigma(y_{i_s}^u)}\right) \tag{23}$$

where $y_{i_w}^u$ is the prediction of unlabeled images with weak augmentation and $y_{i_s}^u$ is the prediction of unlabeled images with strong augmentation. The total loss $\mathcal{L}_{total}$ is defined as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_s + \mathcal{L}_{\text{ADDA}} + \mathcal{L}_{\text{phase}} + \lambda_6\mathcal{L}_{\text{corr}}^u + \lambda_7\mathcal{L}_{\text{semi}} + \lambda_8\mathcal{L}_{\text{kl}} \tag{24}$$

where $[\lambda_6, \lambda_7, \lambda_8]$ are the hyper parameters.

4. Experiments

4.1. Evaluation metric and datasets

Two layer segmentation datasets were used in this study to evaluate our method. There was no overlap of patients among the training, validation and test sets. Dice score (%) was used the evaluation metric.

Scalp HR-MR dataset was acquired from the First Affiliated Hospital of Soochow University. Approvals of all ethical and experimental procedures and protocols were granted by the Ethics Committee of First Affiliated Hospital of Soochow University under Application No.(2023)517. Scalp HR-MR images were scanned from 54 cases. Each case contained two scans which had 18 slices with the size of 384×384, annotated for 6 layers: stratum corneum and stratum granulosum(SCG), stratum spinosum and stratum basale(SSB), papillary layer(PL), reticular layer(RL), subcutaneous fat layer containing hair follicles(SFL-HF) and subcutaneous fat layer without hair follicles (SFL-NHF). The two scans were acquired from the vertex and occipital region of each subject. After removing invalid slices, a total of 1240 slices were obtained. Among them, 37 cases (833 slices) were randomly selected for training, including 12 cases (245 slices) as 30% labeled data, 4 cases (83 slices) as 10% labeled data, and 2 of these cases (41 slices) as 5% labeled data. Additionally, 4 cases (100 slices) were used for validation, and the remaining 11 cases (307 slices) were used for testing. This dataset is challenging due to noise, inter-layer artifacts, and ambiguous layer boundaries.

Retinal OCT dataset [42] comprised retinal OCT images from 35 subjects, including 14 healthy controls (HC) and 21 multiple sclerosis (MS) patients. Each retinal OCT encompassed 49 B-scans with a resolution of 496×1024 pixels, annotated for 8 retinal layers: RNFL, GCIP, INL, OPL, ONL, IS, OS, and RPE. Following [22,27], each B-scan was divided into two non-overlapping samples, and then center-cropped to a height of 384. This process resulted in 3430 retinal OCT scans, each 384×512 pixels in size, randomly divided into 2940 training, 196 validation, and 294 test images. For the training, in addition to the experiment that only used 5 labeled data (0.2%), we also set up two additional annotation ratios: 1% and 10%.

4.2. Implementation details

All models were implemented using PyTorch and run on an NVIDIA GeForce GTX 4090 GPU. U-Net was used for fully supervised training. Following previous work [43], DeepLabV3+ based on ResNet101 was used as our encoder. Details of the hyper-parameter settings are provided in Section 5.3. The Adam optimizer was used with a learning rate set to 1e-4, the learning rate dropped by a factor of 0.9 every 25 epoch. The batch size was set to 6. The number of epochs was set to 200 on scalp HR-MR dataset and 100 on retinal OCT dataset.

4.3. Comparison with state-of-the-art methods

In order to comprehensively verify the proposed method, comparative experiments have been performed on scalp HR-MR dataset and retinal OCT dataset with several recent state-of-the-art (SOTA) methods, such as Crcfp [44], BCP [45], SDCL [10], CorrMatch [39], ABD(BCP) [46], UniMatch [43], BE-SemiNet [27], DIFFTCL [22]. Specifically, since BE-SemiNet [27] and DIFFTCL [22] did not open their source code, we only adopted their results in retinal OCT dataset [42].

4.3.1. Scalp HR-MR

As reported in Table 1, our ADRL method achieved the highest average Dice scores. ("Mean" values indicated the average of layer-wise metrics across all test cases.) With 5% annotation on the Scalp HR-MR dataset, it outperformed the second-best method by 0.35, 4.14, 1.96, 3.09 and 0.86 in Dice score on the SCG, SSB, PL, SFL-HF, and SFL-NHF layers, respectively, while slightly underperforming on RL. Under the 10% annotation, ADRL improved Dice scores by 1.84, 4.04, 2.72, 1.22, 2.88 and 1.55 over the second-best method across all layers. When the annotation ratio increased to 30%, ADRL still consistently outperformed the second-best methods by 1.36, 1.29, 2.61, 1.22 and 1.62 on the first five layers in terms of Dice score. A paired two-tailed *t*-test yielded $p \ll 0.001$, indicating that the performance improvements of the ADRL were statistically significant.

As shown in Fig. 3, the visual comparisons among all these methods with 10% labeled data. The existing methods were unable to accurately segment thin layers, where the boundaries were blurred and the contrasts between neighboring layers were low. Crcfp led to over-segmentation of layers, while BCP, ABD(BCP) and SDCL tended to produce under-segmentation in layers with low contrast. UniMatch missed several thin layers and CorrMatch failed to capture the detailed boundary of the thinnest layer. In contrast, ADRL preserves the structure of layer boundary and retains the details of thin layers.

In addition, we compared the probability heat maps produced by ADRL and related methods. As shown in Fig. 5, for the SCG layer, the intensities in the probability heat maps of most methods, such as Crcfp, BCP, SDCL and ABD(BCP), were low, and the structures were not fully connected due to the low contrast under the influence of artifacts. The probability heat maps of U-Net, UniMatch and CorrMatch also showed low intensities in many segments of SCG and/or SFL-NHF. In contrast, ADRL maintained high confidence and the continuity of layers. For thick layer such as SFL-NHF, methods with global constraints such as CorrMatch performed well, while the rest of the methods produced weak heat maps. Nevertheless, ADRL still achieved good performance, with continuous and complete segmentation boundary.

4.3.2. Retinal OCT

Table 2 reported the performance of the proposed method and recent SOTA methods for retinal layer segmentation on retinal OCT dataset with 0.2%, 1%, and 10% labeled data settings. When only 0.2% were available, the proposed method achieved stable and competitive performance, outperforming the second-best method by 1.94, 2.41, 2.92, 1.43, 0.08, and 2.05 in Dice score on the RNFL, GCIP, INL, IS, OS, and RPE layers, respectively, and by 1.78 in terms of average Dice score. Although most methods can accurately segment thin layers such as IS and OS, the proposed method demonstrated consistently superior performance on these layers and even outperformed fully supervised methods on the OS layer with only 0.2% labeled data. A paired two-tailed *t*-test yielded $p \ll 0.001$, indicating that the performance improvements were statistically significant.

Fig. 4 shows the visual comparison of ADRL and related methods. Segmenting macula layers is challenging. Crcfp produced over- and under-segmentation in many layers, while BCP, SDCL, and ABD(BCP) tended to under-segment multiple layers. UniMatch and CorrMatch yielded relatively complete results, however, for fine layers such as RNFL, both methods showed poor segmentation and loss of details. In contrast, ADRL clearly delineated the structure of each layer, mitigating over-segmentation and under-segmentation. As shown in Fig. 5, on the retinal OCT dataset, most methods generated low-intensity heat maps and disconnected OS layers in the presence of artifacts, whereas ADRL maintained continuity and integrity. For thicker layers such as GCIP,

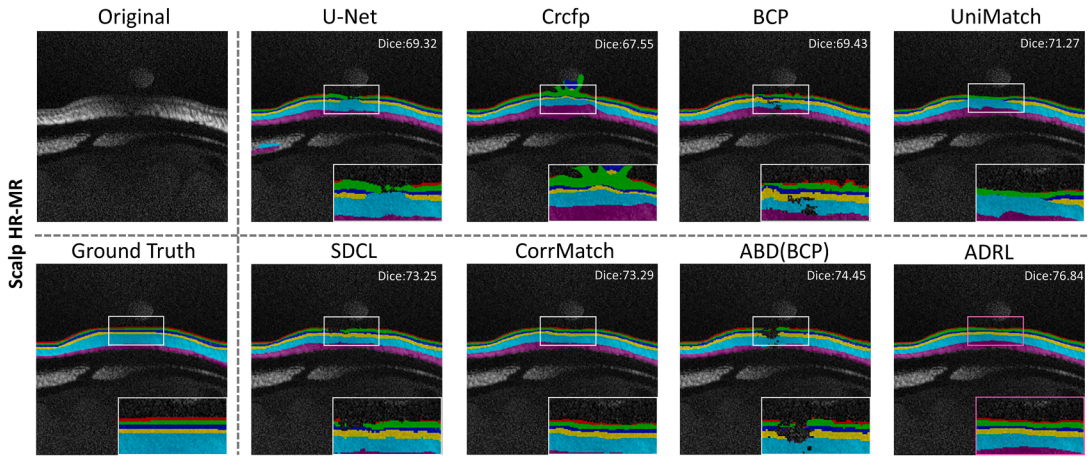


Fig. 3. Visualized segmentation results of ADRL and the state-of-the-art methods on the scalp HR-MR dataset using 10% labeled data. Boxes indicate regions with notable differences across methods, including boundary inconsistencies, shape deviations, and segmentation errors in anatomically complex regions.

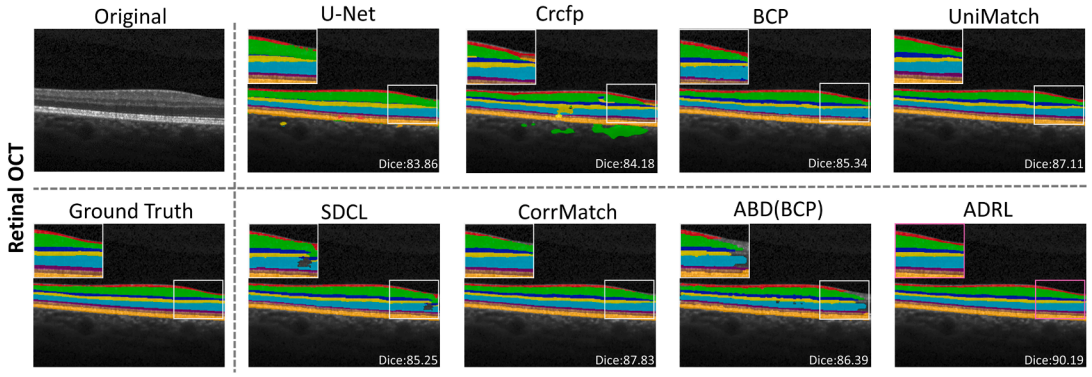


Fig. 4. Visualized segmentation results of ADRL and the state-of-the-art methods on retinal OCT dataset with 0.2% labeled data. Boxes indicate regions with notable differences across methods, including boundary inconsistencies, shape deviations, and segmentation errors in anatomically complex regions.

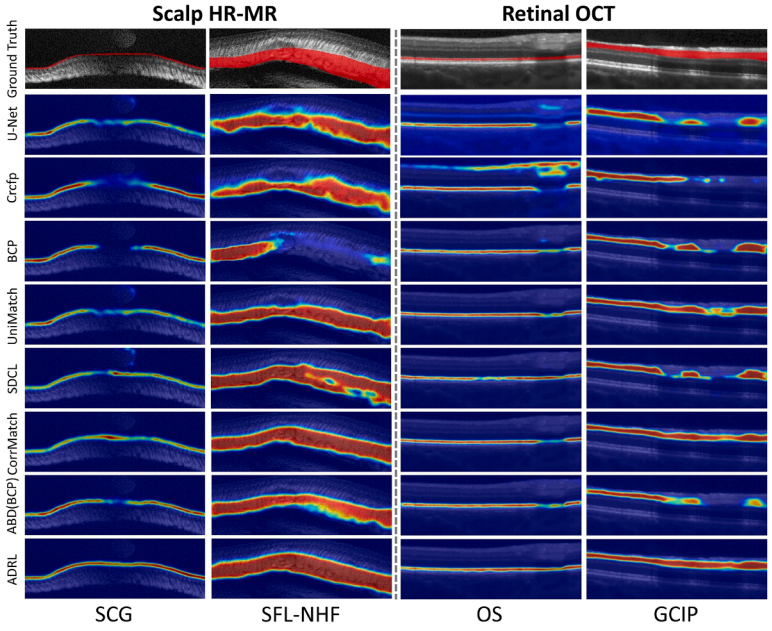


Fig. 5. Visualized comparisons of our method and the state-of-the-art methods by probability heat map of detailed layers in scalp HR-MR dataset(10% labeled) and retinal OCT (0.2% labeled) dataset.

Table 1
Results of Our Method and the State-of-the-Art Methods on Scalp HR-MR Dataset with 5%, 10%, and 30% labeled data. The best and second-best Dice score (%) are shown in bold and underline, respectively.

Methods		Scalp HR-MR						
		SCG	SSB	PL	RL	SFL-HF	SFL-NHF	Mean
5%	U-Net(2015MICCAI)	58.82	70.55	62.79	63.13	70.46	78.79	67.42[a]
	Crcfp(2024MIA)	53.25	63.60	56.64	60.18	72.38	77.36	63.90[a]
	BCP(2023CVPR)	59.17	72.91	68.49	59.47	61.79	62.72	64.09[a]
	UniMatch(2023CVPR)	56.14	71.80	64.19	67.57	79.95	81.52	70.20[a]
	SDCL(2024MICCAI)	60.48	74.58	71.12	70.16	75.88	80.40	72.10[a]
	CorrMatch(2024CVPR)	58.38	73.90	67.08	68.95	80.90	84.30	72.25[a]
	ABD(BCP)(2024CVPR)	61.41	73.37	69.34	**72.58**	80.98	78.49	72.70[a]
	ADRL(Ours)	**61.76**	**78.72**	**73.08**	68.70	**84.07**	**85.16**	**75.25**
10%	U-Net(2015MICCAI)	59.81	72.88	64.71	65.07	73.62	79.84	69.32[a]
	Crcfp(2024MIA)	57.86	68.69	62.87	63.82	72.30	79.77	67.55[a]
	BCP(2023CVPR)	60.68	73.75	68.58	67.11	73.04	73.42	69.43[a]
	UniMatch(2023CVPR)	58.02	73.93	66.42	67.71	79.89	81.66	71.27[a]
	SDCL(2024MICCAI)	62.08	74.63	71.72	70.60	78.51	80.94	73.25[a]
	CorrMatch(2024CVPR)	59.02	74.54	69.22	70.41	82.67	83.87	73.29[a]
	ABD(BCP)(2024CVPR)	62.53	75.35	70.97	72.54	82.48	82.82	74.45[a]
	ADRL(Ours)	**64.37**	**79.39**	**74.44**	**73.76**	**85.55**	**85.42**	**77.15**
30%	U-Net(2015MICCAI)	61.82	75.54	70.63	69.45	78.42	83.05	73.15[a]
	Crcfp(2024MIA)	58.55	68.85	69.34	63.97	75.03	79.81	69.26[a]
	BCP(2023CVPR)	62.14	76.63	71.02	70.82	81.96	85.15	74.62[a]
	UniMatch(2023CVPR)	62.20	75.60	67.76	69.20	81.45	83.46	73.28[a]
	SDCL(2024MICCAI)	62.92	77.05	71.74	72.52	82.92	85.82	75.50[a]
	CorrMatch(2024CVPR)	60.02	78.64	70.11	70.48	84.22	**87.50**	75.16[a]
	ABD(BCP)(2024CVPR)	64.81	77.01	73.43	74.43	84.22	85.58	76.58[a]
	ADRL(Ours)	**66.17**	**79.93**	**76.04**	**75.65**	**85.84**	85.48	**78.19**
100%	U-Net(2015MICCAI)	65.66	79.26	74.47	74.66	84.16	84.65	77.31

[a] p of paired two-tailed t-test $\ll 0.001$. The best is in bold font, and the second best is underlined.

Table 2
Results of Our Method and State-of-the-Art Methods on Retinal OCT Dataset with 0.2%, 1%, and 10% labeled data. The best and second-best Dice scores(%) are shown in bold and underline, respectively.

Methods		Retinal OCT								
		RNFL	GCIP	INL	OPL	ONL	IS	OS	RPE	Mean
0.2%	U-Net(2015MICCAI)	85.80	86.25	77.10	81.58	91.62	78.73	81.42	88.42	83.86[a]
	Crcfp(2024MIA)	89.32	88.74	80.47	81.02	91.56	73.91	79.14	89.25	84.18[a]
	SDCL(2024MICCAI)	89.07	92.60	82.87	81.93	87.04	77.21	82.10	89.14	85.25[a]
	BCP(2023CVPR)	86.96	90.58	81.61	84.98	89.24	81.06	82.94	85.32	85.34[a]
	ABD(BCP)(2024CVPR)	87.20	91.59	79.38	85.32	93.47	83.95	83.48	86.77	86.39[a]
	UniMatch(2023CVPR)	88.18	89.86	82.12	86.72	**93.93**	83.88	82.73	89.46	87.11[a]
	CorrMatch(2024CVPR)	88.64	91.65	84.65	87.23	93.57	83.85	83.81	89.22	87.83[a]
	BE-SemiNet[b](2023CMIG)	89.62	91.12	84.69	87.56	92.78	83.82	86.09	87.34	87.88
	DiffTCL[b](2024MICCAI)	90.69	91.14	84.98	**88.63**	92.98	84.60	86.09	87.42	88.32
	ADRL(Ours)	**92.63**	**95.01**	**87.90**	87.69	93.88	**86.03**	**86.17**	**91.51**	**90.10**
1%	U-Net(2015MICCAI)	89.46	92.90	79.32	85.03	93.84	85.56	84.11	86.70	87.12[a]
	Crcfp(2024MIA)	88.54	92.37	85.45	86.95	93.45	84.63	83.52	88.96	87.98[a]
	SDCL(2024MICCAI)	89.94	93.25	83.28	87.23	93.14	77.83	83.81	89.81	87.29[a]
	BCP(2023CVPR)	89.87	93.76	85.38	86.95	93.74	82.89	83.28	88.69	88.07[a]
	ABD(BCP)(2024CVPR)	90.27	94.01	85.62	87.09	93.15	82.87	83.21	89.32	88.19[a]
	UniMatch(2023CVPR)	90.45	93.37	86.97	88.65	93.89	84.76	85.34	90.45	89.24[a]
	CorrMatch(2024CVPR)	92.77	94.55	85.29	88.49	94.46	84.56	84.74	90.14	89.38[a]
	ADRL(Ours)	**93.15**	**94.59**	**87.10**	**89.33**	**94.81**	**86.89**	**85.68**	**91.20**	**90.34**
10%	U-Net(2015MICCAI)	90.51	94.69	86.60	87.81	93.54	84.99	85.18	89.26	89.07[a]
	Crcfp(2024MIA)	90.78	93.54	85.58	88.17	94.42	85.47	85.80	90.35	89.26[a]
	SDCL(2024MICCAI)	90.19	94.21	84.31	86.48	92.64	83.24	84.59	91.28	88.38[a]
	BCP(2023CVPR)	90.43	94.36	85.43	87.68	93.82	83.38	83.65	90.12	88.61[a]
	ABD(BCP)(2024CVPR)	90.71	94.43	86.01	87.48	93.48	83.50	83.93	90.18	88.72[a]
	UniMatch(2023CVPR)	92.89	94.83	86.31	88.98	84.59	85.67	86.04	**91.29**	90.07[a]
	CorrMatch(2024CVPR)	93.14	95.03	87.03	88.68	94.56	86.46	85.56	90.94	90.18[a]
	ADRL(Ours)	**93.26**	**95.10**	**87.73**	**89.23**	**95.16**	**87.87**	**86.29**	91.16	**90.73**
100%	U-Net(2015MICCAI)	92.79	95.59	87.91	89.87	95.17	87.20	86.15	91.67	90.79

[a] p of paired two-tailed t-test $\ll 0.001$.
[b] The numbers are borrowed from previous work. The best is in bold font, and the second best is underlined.

Table 3
Ablation Study on Scalp HR-MR Dataset with 10% Labeled Data.

Basic	CPSC	PHA	ADDA	SCG	SSB	PL	RL	SFL-HF	SFL-NHF	Mean
✓				60.17	75.89	70.90	71.57	82.07	84.11	74.12
✓	✓			59.69	73.44	72.42	73.55	84.03	**86.27**	74.90
✓		✓		59.16	75.95	72.04	74.20	84.17	85.36	75.15
✓			✓	63.87	76.87	72.30	**74.72**	83.81	85.92	76.25
✓	✓	✓		63.19	77.37	74.29	72.82	84.26	82.83	75.79
✓	✓		✓	62.68	79.11	73.33	73.53	83.29	85.38	76.22
✓		✓	✓	63.78	78.01	72.81	73.64	85.18	85.15	76.43
✓	✓	✓	✓	**64.37**	**79.39**	**74.44**	73.76	**85.55**	85.42	**77.15**

Table 4
Ablation Study on the Retinal OCT Dataset with 0.2% Labels Data.

Basic	CPSC	PHA	ADDA	RNFL	GCIP	INL	OPL	ONL	IS	OS	RPE	Mean
✓				90.92	93.81	86.21	85.83	93.27	83.93	80.98	87.88	87.79
✓	✓			91.10	94.02	84.90	85.57	93.10	85.41	83.98	90.60	88.58
✓		✓		92.09	93.69	85.99	87.59	93.58	83.74	85.44	90.39	89.06
✓			✓	91.74	94.73	85.82	86.42	93.81	85.96	85.95	91.42	89.48
✓	✓	✓		92.61	94.95	87.18	**88.82**	**94.14**	83.65	83.30	89.38	89.25
✓	✓		✓	92.32	94.74	86.29	88.34	93.42	83.63	85.19	90.17	89.26
✓		✓	✓	92.49	94.17	85.64	88.68	94.06	85.68	85.83	90.49	89.63
✓	✓	✓	✓	**92.63**	**95.01**	**87.90**	87.69	93.88	**86.03**	**86.17**	**91.51**	**90.10**

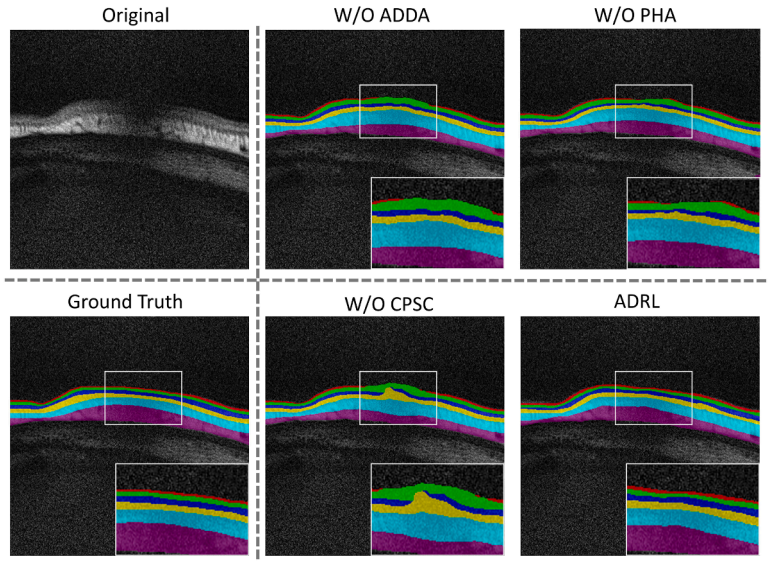


Fig. 6. Comparison of pseudo labels with different strategies on scalp HR-MR dataset with 10% labeled data.

CorrMatch [39] performed comparably to results on the scalp HR-MR dataset, while other methods exhibited low-intensity heat maps and disconnected layers. ADRL consistently produced high-intensity heat maps and well-connected layers.

4.4. Ablation studies

4.4.1. Effects of each module in ADRL

Tables 3 and 4 demonstrated the effectiveness of each module in ADRL with different comparisons in two datasets.

The integration of all three modules–ADDA, PHA, and CPSC–yielded the highest average segmentation performance both the two datasets. On the retinal OCT dataset, the RNFL, GCIP, INL, IS, OS and RPE layer segmentation exhibited the highest performance. This demonstrated the complementary nature of the modules in addressing diverse segmentation challenges.

ADDA: As shown in Table 3, on scalp HR-MR dataset, when the ADDA was incorporated independently, it particularly enhanced the segmentation performance of the first four thin layers, with the SCG layer segmentation reaching a Dice score of 63.87, which represented the second-best result achieved. When CPSC and ADDA are used together, the segmentation accuracy for the SSB and PL layers was improved, with the SSB layer reaching 79.11, the second-best performance. When PHA and ADDA were combined, the performance for the SFL-HF layer and the average performance both achieved the second-best results. When CPSC, PHA, and ADDA modules were integrated, the fourth layer achieved the second-highest Dice score.

Similarly, on the MS dataset (Table 4), when the ADDA module was used alone, the performance for the IS and OS layers reached 85.96 and 85.95, respectively, both ranking the second best, representing improvements of 2.03 and 4.97. Additionally, the segmentation accuracy for the RPE layer reached 91.42, also the second-best. When PHA and ADDA were jointly applied, the performance for the OPL and ONL layers attained the second-best level, while the IS and OS layers maintained significant improvements, and the average performance ranked the second best. Across both datasets, experiments incorporating the ADDA

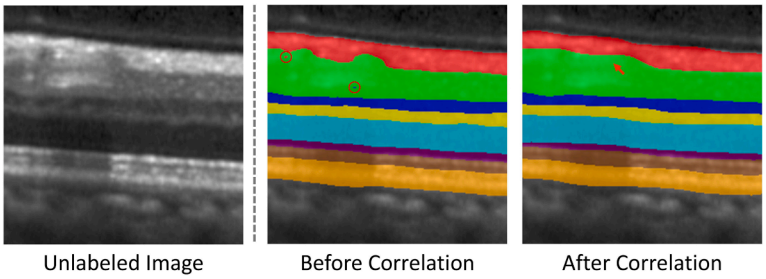


Fig. 7. Comparison of pseudo labels before and after CPSC correlation on retinal OCT dataset with 0.2% labeled data.

module consistently yielded substantial performance gains for thin layer segmentation. These results confirmed ADDA's effectiveness in enhancing thin layer segmentation across datasets.

Fig. 6 presented the segmentation results of various ablation experiments on the scalp HR-MR dataset. Under strong noise, the absence of the ADDA module led to discontinuities in SCG layer segmentation and indistinct PL boundaries, consistent with our hypothesis. Moreover, as shown in Fig. 9, during the early training phase (0–40 epochs), the value of τ rapidly decreased from a larger value and stabilizes, approaching 0 around 100 epochs,indicating it no longer significantly influenced training. Notably, the absence of the ADDA resulted in a slower decline in loss compared to scenarios where the ADDA was included, and the convergence was less optimal. As can be seen according to the mean Intersection over Union (mIoU) on the validation set, a significant performance gap emerged around 75 epochs, with the disparity was stabilized in the later stages of training. This experiment demonstrated the ADDA's ability to adaptively enhance model performance during training.

PHA: The PHA module played a compelling role, as evidenced by the Table 3. On the scalp HR-MR dataset, adding the PHA module alone notably improved the segmentation of the latter four layers. In particular, the RL layer achieved a Dice score of 74.20, the second-best result. Combining CPSC and PHA yielded the second-best performance for PL, while combining PHA and ADDA achieved the second-best results for SFL-HF. These findings, along with the prior observation that ADDA excelled in thin layer segmentation, suggested that PHA acts as both an amplifier and harmonizer. On the retinal OCT dataset, Table 4 shows that combining CPSC and PHA improved OPL and ONL segmentation. Although PHA with ADDA did not achieve top performance for any single layer across both datasets, the average performance consistently ranked the second best, demonstrating PHA's role as a harmonizer and amplifier between CPSC and ADDA.

Moreover, Fig. 8 showed the probability heat map comparison with and without the PHA module on the scalp HR-MR dataset. As shown in Fig. 6, without PHA, the model exhibited mild discontinuities in SCG and blurred boundaries in PL segmentation. Fig. 8 further investigated the role of the PHA module, without PHA, the model could segment the general contour of thin layers like SCG under artifacts, but produced low probabilities. For thicker layers such as SFL-NHF, incorrect segmentation occurred due to interference from shadow of scalp tissue e.g., blood vessels. The PHA significantly improved overall segmentation, effectively mitigating disturbances and enhancing performance in challenging regions by aligning and adjusting model features in the phase domain.

CPSC: The CPSC module captured global constraints to guide model segmentation. As shown in Table 3, addition the CPSC alone significantly improved the segmentation of thicker layers, such as SFL-HF and SFL-NHF, with Dice score improvements of 1.96 and 2.16, respectively. Combination all three modules (CPSC, PHA, and ADDA) achieved the best performance, while removing CPSC reduced SFL-HF and SFL-NHF segmentation. This suggested that global constraints were crucial for thicker layers. Table 4 further confirmed this on the retinal OCT dataset, where omitting CPSC decreased accuracy for GCIP, INL and RPE layers by 0.84, 2.26, and 1.02, respectively. These results highlighted the

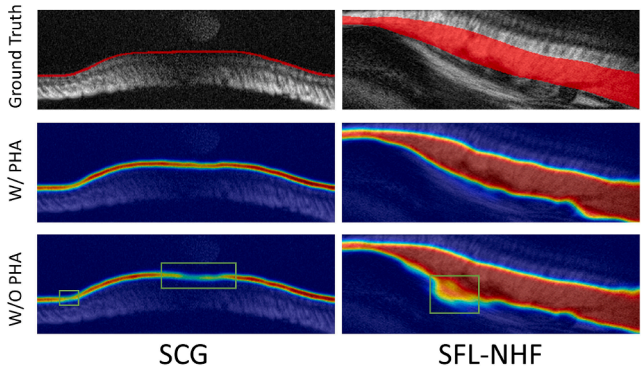


Fig. 8. Heat map comparison with and without PHA on the scalp HR-MR dataset.

CPSC's effectiveness in leveraging global constraints to enhance segmentation accuracy for thicker layers.

In Fig. 7, we showed a visual comparison between before and after applying the CPSC. The prediction obviously produced incorrect segmentation and inaccurate segmentation of layer boundaries without CPSC. After applying CPSC, the incorrect segmentation was mitigated, and the edge of the thicker layer was also corrected and supplemented. The combined results from Tables 3 and 4 clearly demonstrated that the global constraints captured by the CPSC significantly enhanced model segmentation performance. These constraints effectively mitigated pixel misclassification and optimized the segmented layer structures, aligning them more closely with the true anatomical structures.

These ablation experiments showed that ADDA significantly improves thin layer segmentation. The PHA acts as regulator between ADDA and CPSC, so that the model can achieve very good results for various types of layer structures. The CPSC enhanced thicker layer segmentation and effectively mitigates pixel misclassification.

5. Discussion

Segmentation of scalp tissue layers has significant clinical value. It enables quantification of layer thickness, providing a precise measure of disease severity. This is particularly valuable in clinical trials and research requiring objective outcomes. By employing image analysis algorithms, we can facilitate accurate treatment monitoring, assisting in the evaluation of treatment efficacy.

In this section, we further discuss the clinical relevance of the strategies and parameters in the proposed framework on model segmentation performance.

5.1. Masking strategy in ADDA module

In semi-supervised medical image segmentation, confident pseudo-labels can enrich training data diversity, thereby improving the robustness of the model, particularly for delicate anatomical structures. We

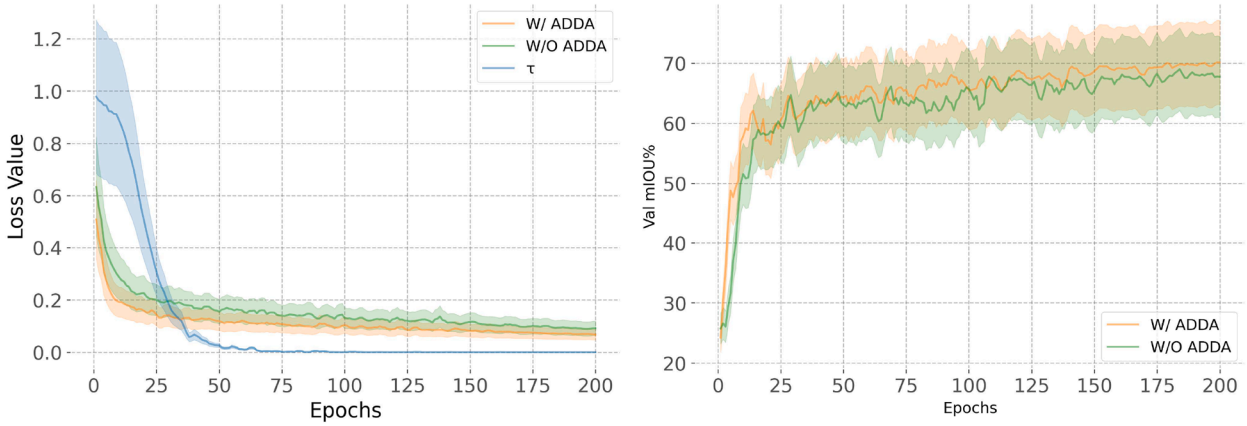


Fig. 9. Effectiveness of ADDA.

Table 5
Comparison Between Adaptive and Fixed Masking Strategies on Scalp HR-MR Dataset.

	Adaptive	Fixed($\tau = 0.04$)
SCG	64.37	**61.53(-2.84↓)**
SSB	79.39	**76.53(-2.86↓)**
PL	74.44	**69.31(-5.13↓)**
RL	73.76	**70.13(-3.63↓)**
SFL-HF	85.55	82.51(-3.04↓)
SFL-NHF	85.42	85.34(-0.08↓)
Mean	77.15	**74.22(-2.93↓)**

Table 6
Comparison Between Adaptive and Fixed Masking Strategies on Retinal OCT Dataset.

Layer	Adaptive	Fixed($\tau = 0.04$)
RNFL	92.63	91.14(-1.49↓)
GCIP	95.01	94.16(-0.85↓)
INL	87.90	**86.00(-1.90↓)**
OPL	87.69	87.83(-0.14↓)
ONL	93.88	93.81(-0.07↓)
IS	86.03	**81.50(-4.53↓)**
OS	86.17	**81.14(-5.03↓)**
RPE	91.51	88.79(-2.72↓)
Mean	90.10	**88.05(-2.05↓)**

Table 7
Parameter Selection Experiments on Retinal OCT Dataset, the best is shown in bold.

λ_1	DSC	λ_2	DSC	λ_3	DSC	λ_4	DSC
1.25	88.29	1.25	88.70	0.75	89.17	1.5	89.33
1	**89.12**	**1**	**89.12**	**0.5**	**89.32**	**1.25**	**89.42**
0.75	88.96	0.75	88.92	0.25	89.03	1	89.32
λ_5	DSC	λ_6	DSC	λ_7	DSC	λ_8	DSC
0.15	88.94	0.15	89.12	0.75	89.69	0.75	89.93
0.1	**89.42**	**0.1**	**89.42**	**0.5**	**89.78**	**0.5**	**90.10**
0.05	89.34	0.05	89.36	0.25	89.52	0.25	89.71

further studied the effects of the fixed mask strategy versus our adaptive masking strategy. As shown in Table 5, using the fixed mask strategy ($\tau = 0.04$ [37]) reduced Dice score by 2.84, 2.86, 5.13, 3.63 on SCG, SSB, PL, RL, and 3.04, 0.08 on SFL-HF and SFL-NHF, indicating notable performance decline, especially in thinner layers. Similarly, Table 6 showed Dice score reductions of 1.49, 0.85, 1.90, 0.14, 0.07, 4.53, 5.03 on PNFL, GCIP, INL, OPL, ONL, IS, OS layers, with IS and OS showing the largest decline. These results demonstrated the advantage of the ADDA module in significantly improving thin layer segmentation.

5.2. Contribution of PHA and CPSC

Accurate tissue boundary delineation was critical in clinical practice, yet frequently challenged by noise, artifacts, and patient variability. The proposed PHA and CPSC modules enhanced segmentation reliability through global and local feature fusion.

The PHA module emphasized high-frequency structural information and provided robustness to noise and contrast variations, enabling reliable detection of subtle or low-contrast tissue boundaries. For example, in cases where scalp layer boundaries were blurred, the PHA module can effectively distinguish them, enhancing segmentation accuracy and capturing fine anatomical details important for diagnosis. This capability ensured precise thickness measurements and early detection of subtle pathological changes, while stabilizing pseudo-labels and supporting semi-supervised learning across diverse patient populations.

The CPSC module complemented this by enforcing global structural consistency through cross-power spectrum correlation, suppressing non-structural noise and improved boundary delineation in complex or ambiguous regions. Together, PHA secured boundary accuracy, while CPSC preserved global integrity. Global-local fusion yielded more robust, consistent, and clinically interpretable segmentation results, supporting precise measurement, disease monitoring, and treatment assessment in applications.

5.3. Parameter selection

The influence of different weights was also explored. In Section 3, parameters λ_1 and λ_2 controlled the weights of the policy network and value network, while λ_3 regulated the influence of the augmentation module on the encoder. λ_4 and λ_5 determined the weight of the phase consistency losses between perturbed image pairs and between perturbed images and the original image, regulating the impact of the structural consistency loss on the encoder's updates. λ_6 controlled the influence of the CPSC module on the encoder, indicating the degree of constraint imposed by the cross-power spectrum on the feature maps. λ_7 and λ_8 controlled the weights of the Kullback-Leibler divergence loss and the unsupervised consistency loss, balancing the strength of distribution alignment and prediction consistency. To prevent any single term from dominating the optimization process and degrading feature extraction capability, we empirically set the initial values of λ_1 and λ_2 to 1, and λ_4 to 1.25, with an adjustment step of 0.25. The initial values of λ_3,λ_7, λ_8 were set to 0.5 with an adjustment step of 0.25, while λ_5 and λ_6 were initialized to 0.1 with an adjustment step of 0.05. Each parameter was tuned individually while keeping the others fixed to identify the optimal configuration. As shown in Table 7, our model exhibited

Table 8
Model Complexity.

	Baseline	ADDA	PHA	CPSC
Params	59.463M	0.177K	1.345K	4.626K
FLOPs	51.330G	0.005G	0.013G	0.231G

optimal performance when setting $\lambda_1 = 1, \lambda_2 = 1, \lambda_3 = 0.5, \lambda_4 = 1.25, \lambda_5 = 0.1, \lambda_6 = 0.1, \lambda_7 = 0.5, \lambda_8 = 0.5$.

5.4. Model complexity

As shown in Table 8, the baseline DeepLabV3 + based on ResNet101 had 59.463M parameters and required 51.33G FLOPs, while our ADDA, PHA, and CPSC modules contained only 0.177K, 1.345K and 4.626K parameters, and only required 0.005G, 0.013G and 0.231G FLOPs, respectively. Despite being extremely lightweight, these modules contributed to a significant performance improvement, achieving strong results on both datasets.

5.5. Limitations and future work

Although the proposed ADRL framework demonstrates promising performance across both datasets, several limitations remain. In cases with highly complex noise or extensive shadowing, the model may still struggle to accurately capture fine edge details of tissue layers. Furthermore, the interpretability of our method is relatively limited. In addition, our method is designed for 2D slice-based segmentation, whereas layered anatomical structures are inherently three-dimensional. In future work, we plan to extend the framework to 3D semi-supervised segmentation and further improve interpretability by incorporating causal inference theory.

6. Conclusion

In this work, to address the challenge of segmenting thin and complex anatomical layers with limited labeled data, we propose a novel amplitude-guided and deep reinforcement learning semi-supervised method for scalp layer segmentation. The ADDA leverages a dynamic, adaptive, reward-driven deep reinforcement learning paradigm to model the distribution of unlabeled images within the frequency domain for robust data augmentation. Additionally, we incorporate the PHA, which uses a phase consistency strategy to improve the ability to capture fine details and boundaries. The CPSC construct a global consistency constraint through cross-power spectrum, mitigating the impact of outliers and clarify boundaries of layers. Extensive experiments on scalp HR-MR and retinal OCT datasets demonstrate the robustness and generalization of the proposed framework under limited annotation settings, showing potential for scalable and annotation-efficient segmentation in real world. Beyond the performance gains, the framework also offers useful insights for designing deep reinforcement learning semi-supervised segmentation strategies in the frequency domain, which may inspire further advances in the field. Nevertheless, the method still faces challenges in accurately capturing subtle boundaries and extending to 3D data. Additionally, it remains largely heuristic and lacks causal interpretability, limiting insights into how features influence optimization and representation learning. Future work will focus on incorporating causal inference, adapting the framework for 3D segmentation tasks, and designing more computationally efficient variants to further improve its generalization and facilitate clinical adoption.

CRediT authorship contribution statement

Enting Gao: Validation, Software, Methodology, Conceptualization; **Zian Zha:** Writing – original draft, Visualization, Methodology; **Yonggang Li:** Data curation; **Junhui Zhu:** Writing – review & editing, Formal analysis; **Yong Wang:** Writing – review & editing, Formal analysis; **Xinjian Chen:** Writing – review & editing, Visualization, Resources, Funding acquisition, Formal analysis; **Naihui Zhou:** Writing – review & editing, Supervision, Resources, Conceptualization; **Dehui Xiang:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Data availability

The data that has been used is confidential.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This study was supported in part by National Nature Science Foundation of China (Grant Nos. 62371328, 62273247, 61971298, 61771326, 81871352), and in part by Jiangsu Key Laboratory of Biophotonics Open Projects Fund (Grant No. 2025BPT004), and in part by Key Laboratory of Advanced Robotics of Jiangsu Province (Grant No. KJS2446), and in part by Beijing Natural Science Foundation (Grant No. 7264253), and in part by National Key R&D Program Disruptive Technology Innovation Project of China (Grant No. 2023YFF1500804). Dehui Xiang and Naihui Zhou are corresponding authors.

References

[1] T. Rhodes, C.J. Girman, Savin, Prevalence of male pattern hair loss in 18–49 year old men, Dermatologic Surg. 24 (12) (1998) 1330–1332.
[2] T.L. Wang, C. Zhou, Y.W. Shen, X.Y. Wang, X.L. Ding, S. Tian, Y. Liu, G.H. Peng, S.Q. Xue, J.E. Zhou, Prevalence of androgenetic alopecia in China: a community-based study in six cities, Br. J. Dermatol. 162 (4) (2010) 843–847.
[3] T. Simakou, J.P. Butcher, S. Reid, F.L. Henriquez, Alopecia areata: a multifactorial autoimmune condition, J. Autoimmun. 98 (2019) 74–85.
[4] Y. Ye, Y. Wang, J. Zhu, R. Huang, Q. Yu, J. Zhang, X. Chen, Z. Wei, Y. Han, N. Zhou, et al., Diagnosis and differential diagnosis of tertiary androgenetic alopecia with severe alopecia areata based on high-resolution MRI, Skin Res. Technol. 29 (7) (2023) e13393.
[5] J. Xue, Q. Li, G. Gong, X. Liu, C. Li, L. Cheng, S. Chang, S. Liang, D. Li, MF2N: multi-view feature fusion network for pancreatic cancer segmentation, Pattern Recognit. 167 (2025) 111738.
[6] Y. Zhao, C. Wang, Y. Hao, L. Li, T. Liao, DyGLNet: hybrid global-local feature fusion with dynamic upsampling for medical image segmentation, Pattern Recognit. 173 (2026) 112792.
[7] B.H. Thompson, G. Di Caterina, J.P. Voisey, Pseudo-label refinement using superpixels for semi-Supervised brain tumour segmentation, in: ISBI, 2022, pp. 1–5.
[8] L. Yu, S. Wang, X. Li, C.-W. Fu, P.-A. Heng, Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation, in: MICCAI, 2019, pp. 605–613.
[9] W. Li, R. Bian, W. Zhao, W. Xu, H. Yang, Diversity matters: cross-head mutual mean-teaching for semi-supervised medical image segmentation, Med. Image Anal. 97 (2024) 103302.
[10] B. Song, Q. Wang, SDCL: students discrepancy-informed correction learning for semi-supervised medical image segmentation, in: MICCAI, LNCS 15008, 2024.
[11] S. Yun, D. Han, S. Chun, S.J. Oh, Y. Yoo, J. Choe, CutMix: regularization strategy to train strong classifiers with localizable features, in: ICCV, 2019, pp. 6023–6032.
[12] A. Tarvainen, H. Valpola, Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results, Neural Inf. Process. Syst. 30 (2017) 1195–1204.
[13] X. Xu, T. Sanford, B. Turkbey, S. Xu, B.J. Wood, P. Yan, Shadow-consistent semi-supervised learning for prostate ultrasound segmentation, IEEE Trans. Med. Imag. 41 (6) (2022) 1331–1345.
[14] A.R. Zamir, A. Sax, N. Cheerla, R. Suri, Z. Cao, J. Malik, L.J. Guibas, Robust learning through cross-task consistency, in: CVPR, 2020, pp. 11197–11206.
[15] A. Yazdanpanah, G. Hamarneh, B.R. Smith, M.V. Sarunic, Segmentation of intra-retinal layers from optical coherence tomography images using an active contour approach, IEEE Trans. Med. Imag. 30 (2) (2010) 484–496.
[16] D. Koozekanani, K. Boyer, C. Roberts, Retinal thickness measurements from optical coherence tomography using a Markov boundary model, IEEE Trans. Med. Imag. 20 (9) (2002) 900–916.
[17] D. Xiang, H. Tian, X. Yang, F. Shi, W. Zhu, H. Chen, X. Chen, Automatic segmentation of retinal layer in OCT images with choroidal neovascularization, IEEE Trans. Image Process. 27 (12) (2018) 5880–5891.

[18] D. Xiang, G. Chen, F. Shi, W. Zhu, Q. Liu, S. Yuan, X. Chen, Automatic retinal layer segmentation of OCT images with central serous retinopathy, IEEE J. Biomed. Health Inform. 23 (1) (2019) 283–295.
[19] A. Shah, L. Zhou, M.D. Abrámoff, X. Wu, Multiple surface segmentation using convolution neural nets: application to retinal layer segmentation in OCT images, Biomed. Opt. Express 9 (9) (2018) 4509–4526.
[20] X. He, W. Song, Y. Wang, F. Poiesi, J. Yi, M. Desai, Q. Xu, K. Yang, Y. Wan, Lightweight retinal layer segmentation with global reasoning, IEEE Trans. Instrum. Meas. 73 (2024) 1–14.
[21] H. Wu, X. Li, Y. Lin, K.-T. Cheng, Compete to win: enhancing pseudo labels for barely-supervised medical image segmentation, IEEE Trans. Med. Imag. 42 (11) (2023) 3244–3255.
[22] X. Li, H. Fang, M. Liu, Y. Xu, L. Duan, Diffusion-enhanced transformation consistency learning for retinal image segmentation, in: MICCAI, 2024, pp. 221–231.
[23] W. Zuo, H. Liu, H. Wang, L. Zeng, N. Tang, F. Shang, L. Wan, J. Deng, Enhancing the impact of model performance gains for semi-supervised medical image segmentation, Pattern Recognit. 168 (2025) 112889.
[24] Wu, Z. Ge, D. Zhang, M. Xu, L. Zhang, Y. Xia, J. Cai, Mutual consistency learning for semi-supervised medical image segmentation, Med. Image Anal. 81 (2021) 102530.
[25] W. Pan, Z. Xu, J. Yan, Z. Wu, R.K.-Y. Tong, X. Li, J. Yao, Semi-supervised semantic segmentation meets masked modeling : fine-grained locality learning matters in consistency regularization, Pattern Recognit. 172 (2026) 112586.
[26] Q. Zeng, Y. Xie, Z. Lu, M. Lu, J. Zhang, Y. Zhou, Y. Xia, Consistency-guided differential decoding for enhancing semi-supervised medical image segmentation, IEEE Trans. Med. Imag. 44 (2024) 44–56.
[27] Y. Lu, Y. Shen, X. Xing, C. Ye, M.Q.-H. Meng, Boundary-enhanced semi-supervised retinal layer segmentation in optical coherence tomography images using fewer labels, CMIG 105 (2023) 102199.
[28] Y. Song, J. Wang, T. Lukasiewicz, Z. Xu, S. Zhang, M. Xu, Mega-reward: achieving human-level play without extrinsic rewards, in: AAAI, abs/1905.04640, 2019, pp. 5826–5833.
[29] W. Zhang, X. Ding, Y. Liu, B. Qiao, Slide deep reinforcement learning networks: application for left ventricle segmentation, Pattern Recognit. 141 (2023) 109667.
[30] D. Yuan, Y. Liu, Z. Xu, Y. Zhan, J. Chen, T. Lukasiewicz, Painless and accurate medical image analysis using deep reinforcement learning with task-oriented homogenized automatic pre-processing, Comput. Biol. Med. 153 (2023) 106487.
[31] C. Ma, Q. Xu, X. Wang, B. Jin, X. Zhang, Y. Wang, Y. Zhang, Boundary-aware supervoxel-level iteratively refined interactive 3D image segmentation with multi-agent reinforcement learning, IEEE Trans. Med. Imag. 40 (10) (2020) 2563–2574.
[32] C. Xu, T. Zhang, D. Zhang, D. Zhang, J. Han, Deep generative adversarial reinforcement learning for semi-supervised segmentation of low-contrast and small objects in medical images, IEEE Trans. Med. Imag. 43 (9) (2024) 3072–3084.
[33] G. Xu, S. Wang, T. Lukasiewicz, Z. Xu, Adaptive-masking policy with deep reinforcement learning for self-supervised medical image segmentation, in: ICME, 2023, pp. 2285–2290.
[34] Z. Xu, Y. Liu, G. Xu, T. Lukasiewicz, Self-supervised medical image segmentation using deep reinforced adaptive masking, IEEE Trans. Med. Imag. 44 (2024) 180–193.
[35] V. Mnih, A.P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, K. Kavukcuoglu, Asynchronous methods for deep reinforcement learning, in: M.F. Balcan, K.Q. Weinberger (Eds.), ICML, 48, 2016, pp. 1928–1937.
[36] H. Zhou, S. Liu, S. Hu, Multi-scale dual-attention frequency fusion for joint segmentation and deformable medical image registration, Inf. Fusion 123 (2025) 103293.
[37] Q. Ma, J. Zhang, L. Qi, Q. Yu, Y. Shi, Y. Gao, Constructing and exploring intermediate domains in mixed domain semi-supervised medical image segmentation, in: CVPR., 2024, pp. 11642–11651.
[38] W. Yu, P. Zhou, S. Yan, X. Wang, InceptionNeXt: when inception meets convnext, in: CVPR, 2024, pp. 5672–5683.
[39] B. Sun, Y. Yang, L. Zhang, M.-M. Cheng, Q. Hou, CorrMatch: label propagation via correlation matching for semi-supervised semantic segmentation, in: CVPR, 2023, pp. 3097–3107.
[40] R.E. Rivadeneira, H.O. Velesaca, A. Sappa, Cross-spectral image registration: a comparative study and a new benchmark dataset, in: Lecture Notes in Networks and Systems, Singapore, 2024, pp. 1–12.
[41] S. Kullback, R.A. Leibler, On information and sufficiency, Ann. Math. Stat. 22 (1951) 79–86.
[42] Y. He, A. Carass, S.D. Solomon, S. Saidha, P.A. Calabresi, J.L. Prince, Retinal layer parcellation of optical coherence tomography images: data resource for multiple sclerosis and healthy controls, Data Brief 22 (2019) 601–604.
[43] L. Yang, L. Qi, L. Feng, W. Zhang, Y. Shi, Revisiting weak-to-strong consistency in semi-supervised semantic segmentation, in: CVPR, 2023, pp. 7236–7246.
[44] R.M.S. Bashir, T. Qaiser, S.E.A. Raza, N.M. Rajpoot, Consistency regularisation in varying contexts and feature perturbations for semi-supervised semantic segmentation of histology images, Med. Image Anal. 91 (102997) (2024) 102997.
[45] Y. Bai, D. Chen, Q. Li, W. Shen, Y. Wang, Bidirectional copy-paste for semi-supervised medical image segmentation, in: CVPR, 2023, pp. 11514–11524.
[46] H. Chi, J. Pang, B. Zhang, W. Liu, Adaptive Bidirectional Displacement for semi-supervised medical image segmentation, in: CVPR., 2024, pp. 4070–4080.