

Robust Source-Free Domain Adaptation for Fundus Image Segmentation

Lingrui Li^{1,2}, Yanfeng Zhou^{1,2}, Ge Yang^{1,2} *

¹Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

{lilingrui2021, zhouyanfeng2020, ge.yang}@ia.ac.cn

Abstract

Unsupervised Domain Adaptation (UDA) is a learning technique that transfers knowledge learned in the source domain from labelled training data to the target domain with only unlabelled data. It is of significant importance to medical image segmentation because of the usual lack of labelled training data. Although extensive efforts have been made to optimize UDA techniques to improve the accuracy of segmentation models in the target domain, few studies have addressed the robustness of these models under UDA. In this study, we propose a two-stage training strategy for robust domain adaptation. In the source training stage, we utilize adversarial sample augmentation to enhance the robustness and generalization capability of the source model. And in the target training stage, we propose a novel robust pseudo-label and pseudo-boundary (PLPB) method, which effectively utilizes unlabeled target data to generate pseudo labels and pseudo boundaries that enable model self-adaptation without requiring source data. Extensive experimental results on cross-domain fundus image segmentation confirm the effectiveness and versatility of our method. Source code of this study is openly accessible at <https://github.com/LinGray/PLPB>.

1. Introduction

Unsupervised domain adaptation (UDA) transfers the knowledge embedded in labeled source domain data to the target domain, mitigating the problem of performance degradation caused by domain shift and dependence on expensive pixel-wise annotations [12, 34, 36]. However, existing UDA methods [8, 28, 40] mainly consider a single target domain, resulting in limited applicability in the real world. Indeed, the target domain may feature multiple data distributions, while domain adaptation techniques may need to handle samples from unseen domains. Therefore it is more practical to study open compound domain adaptation

(OCDA) [22], where the target has multiple homogeneous unlabeled domains as well as open domains unseen before.

For OCDA, two practical issues need to be addressed. First, existing UDA methods typically focus on solving the adaptation problem, but often do not take robustness into consideration [11, 22, 34, 38, 40]. While image degradations and/or adversarial attacks in the natural image domain have been studied extensively [1, 4, 14, 18, 29], their study in the medical image domain remains limited. Vulnerability of medical deep learning systems to image degradations and/or adversarial attacks may lead to bias in downstream tasks such as patient diagnosis and treatment [3, 16, 24]. This raises safety concerns about the deployment of these systems in clinical settings.

Second, source data may not be available due to privacy or storage constraints. Recent studies have developed various source-free domain adaptation techniques [5, 13, 21, 39]. However, the performance of these techniques may drop substantially on open-domain datasets unseen before. Real-world applications require adaptation to both multiple target domains and unseen open domains. Together, this poses a more challenging problem, namely source-free open-compound domain adaptation (SF-OCDA) [43], which is the subject of this study.

In this study, we utilize adversarial sample augmentation in source training to enhance the robustness and generalization of the source model. And in the target model training stage, we propose a **pseudo-labeling** and **pseudo-boundary** (PLPB) method for SF-OCDA. Our PLPB method utilizes pseudo labelling, pseudo boundary modeling and entropy minimization to enable effective self-training in the target domain without source data. By adversarial sample augmentation in the source domain and explicit modeling of the boundaries in the target domain, our model generalizes better to the unseen target domain. In addition, our work uses adversarial samples in domain adaptation and leverages them to benchmark model performance. Experimental results show that our method substantially enhances model robustness while maintaining model adaptation performance on clean samples in target domains. Figure 1

*Corresponding author.

shows the overall performance of our method on clean and adversarial samples in comparison with competing methods. The contributions of our work are summarized as follows:

- We proposed a two-stage training strategy for robust domain adaptation for semantic segmentation of medical images without source data.
- We utilize a pseudo-boundary loss in the target adaptation stage and develop a new domain adaptation method PLPB. Our method models edge information and achieves good performance on Average Surface Distance (ASD) metric and obtains precise boundary prediction. Without requiring source data, our method achieves comparable and sometimes higher performance than state-of-the-art (SOTA) source-dependent UDA methods and other SF-DA methods.
- We evaluate the efficacy of PLPB on two target domain fundus datasets and one open domain dataset which are popular benchmarks for UDA tasks, demonstrating improvements on both clean and adversarial samples. By utilizing adversarial samples, our method also demonstrates good generalization capability in the open domain. Our method is flexible and can be combined with other existing adaptation techniques.

2. Related Works

2.1. Unsupervised Domain Adaptation (UDA)

Unsupervised domain adaptation (UDA) aims to tackle domain shift by adapting the training process of a model in an unsupervised manner. Leveraging unsupervised learning reduces the laborious and time-consuming data labeling work for the target domain. Therefore, UDA is a promising method to solve domain shift problems, especially in the medical image analysis field where data is diverse and requires expert data labeling.

Recently, many deep learning-based domain adaptation methods [8, 12, 23] have been proposed and achieved encouraging results. Many methods tackle the domain shift issue by extracting invariant features. And a typical approach is adversarial learning [34, 36]. Another popular method is image-to-image translation [11, 40]. Noticeably, image-to-image translation usually introduces artifacts, which may be not a proper approach in the medical field. Furthermore, although adversarial learning can align the latent feature distribution of different domains, the results of multiple adversarial learning-based methods easily suffer from sub-optimal performance due to the difficulty of stabilizing the training process of multiple adversarial modules [19].

In this study, we use neither adversarial learning nor image-to-image translation in the adaptation stage. As

boundary is domain-invariant information in different domains, modeling the boundaries helps the model generalize better to unseen target domains and generate more precise and coherent predictions, we utilize the low-level boundary information in both source training and target training stages. The recent work [20] also considers low-level edge information by edge map. However, it does not consider the practical reality of source-free setting and open target domains.




Figure 1. Performance comparison of PLPB (OURS Both) with competing methods. (a) Comparison of Dice score, higher is better. (b) Comparison of Average Surface Distance (ASD), lower is better. These results of different methods are averaged over all domain adaptation tasks for multiple clean and adversarial target datasets. Our proposed methods show substantial improvement over the baselines. ST.: only the standard source model is used in the target adaptation stage. R.: only the robust source model is used. C denotes compound target domains, C+O denotes compound and open target domains.

2.2. Source Free Domain Adaptation (SF-DA)

Existing SF-DA methods can be mainly divided into two categories: (i) generating images bypassing the dependence on the source data [13, 15, 21]; (ii) self-supervision with target pseudo-labels [5, 17]. The generative approach is often difficult to scale up, as learning to generate the images is difficult. On the other hand, pseudo-label based methods are easy to handle and have recently provided very promising results. Thus, in this work, we propose to leverage the self-supervised method as in [5, 35].

In SF-DA, without access to the source data, only the source pre-trained model is provided in the target training stage. SHOT [17] maintains the source hypothesis by fixing the trained classifier and maximizes mutual information of target outputs for distribution alignment. Lately, Liu et al. [21] introduce source-free domain adaptation for semantic segmentation and utilize the batch normalization statistics of the source model to recover source-like samples.

Though recent SF-DA methods have made great progress in object recognition and semantic segmentation,

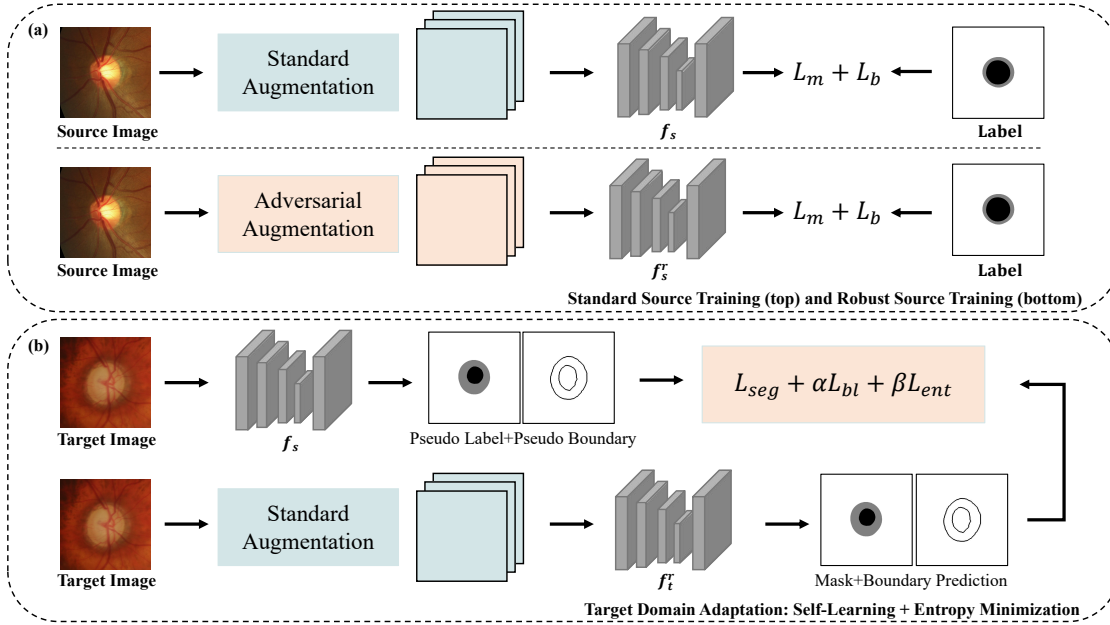


Figure 2. **Overview of our proposed method pseudo labeling and pseudo boundary (PLPB) for source-free open compound domain adaptation.** f_s denotes the standard source model. f_s^r denotes the robust source model. Standard Augmentation includes several typical image augmentations which are described in section 4.3, while Adversarial Augmentation means generating adversarial samples by Projected Gradient Descent (PGD) [26]. The training of the robust source model uses the source data, corresponding labels, and adversarial examples. The target model only uses the target data and the source models and combines three losses with pseudo-label and pseudo-boundary. Robust target model f_t^r is initialized from f_s^r .

SF-DA has not been well investigated in medical image segmentation. Moreover, they simply apply the source model as a pre-trained model to generate pseudo labels for unlabeled target data, ignoring useful information embedded in the boundaries of the objects which is important in medical images. In contrast, we utilize boundary information stored in the source model to generate pseudo boundaries, thus making better use of the source model.

In this work, we consider the source-free open compound domain adaptation (SF-OCDA) [43] for robust fundus semantic segmentation, extending SF-DA to a more realistic setting that includes an open target domain. SFDA-FSM [39] and DPL [5] also adopt two-stage training strategies and investigate domain adaptation in semantic segmentation. Nevertheless, the main difference between our method and the above two methods lies in the setting. First, both SFDA-FSM [39] and DPL [5] focus on the single or close-set target domain adaptation, where the performance on only one target domain is evaluated. Our SF-OCDA setting focuses on both the unlabeled compound domain and the unseen open domain, which is more challenging and practical than SFDA-FSM [39] and DPL [5]. Second, both methods do not take adversarial attacks into consideration.

2.3. Robust Training

A wide variety of attack mechanisms [14, 26, 29] has been proposed since the vulnerability of deep network was shown first by [10]. This also has led to the development of strategies to defend against such attacks, called defense mechanisms [1, 25]. Among them, adversarial training has stood out as the most reliable way to train robust models.

2.4. Robust Domain Adaptation

Our work is also inspired by recent studies on robust transfer learning [1, 30, 31]. A notable work [32] shows that robust source feature extractors can be effective in preserving robustness, while maintaining sufficiently high accuracy on clean samples. On the other hand, [1] shows that the robust pre-trained models not only perform well on targets without adversarial training, but also improve the accuracy on clean samples. Additionally, [9, 41] shows that a pre-trained model that is more robust to input perturbation provides a better initialization for generalization on downstream out-of-distribution data. These results support the hypothesis that robust models also transfer better. However, the existing methods are neither developed nor tested in the settings of SF-OCDA. Peshal et al. [1] studied the problem of SF-DA under the adversarial perturbations in image classification. However, pixel-level segmentation is more chal-

lenging and fundamentally different from the image-level classification which just associates the label with a whole image [21].

3. Methods

Given the source dataset $(x_s^i, y_s^i) \sim D_s$ from the source domain, we first train a source model f_s on D_s . Given f_s , our final goal is a target model f_t that can accurately segment target images x_t^i from target domain D_t , while being robust to adversarial examples x_t^{adv} simultaneously.

The entire procedure consists of two steps, shown in Figure 2. In step (i), as shown in Figure 2(a), we train source models on the source domain in a supervised way. We train two models in the source domain, i.e., a standard model and a robust model, following standard and robust training processes, respectively, where the robust training data is augmented with adversarial examples. In step (ii), we initialize the target models with the aforementioned source models, and then adapt the source model to the target domain without source data, using pseudo-labels and pseudo-boundaries generated from the source model, as shown in Figure 2(b). The pseudo-labels generated from the standard and robust source models are referred to as non-robust labels and robust labels, respectively, for simplicity. The same goes for the definitions of non-robust boundaries and robust boundaries. Figure 3 illustrates our framework of target model training for fundus image segmentation via denoised pseudo-labeling and pseudo-boundary (PLPB) with entropy minimization in SF-OCDA. In the following, We first describe the supervised training in the source domain in Section 3.1, then we propose the pseudo-boundary strategy and generating pseudo-labels on the target adaptation stage in Section 3.2.

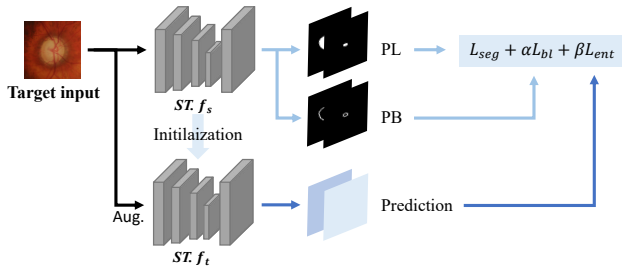


Figure 3. **Overview of standard flow on target adaptation stage.** Aug. denotes augmentation we used in target training which is described detailedly in section 4.3. PB denotes pseudo-boundary. PL denotes pseudo-label. This is the case of giving only the standard source model, i.e. standard target adaptation. The only difference of robust flow is that the robust source model is given instead of the standard model.

3.1. Source Domain Model Training

Similar to [38], our segmentation network includes a boundary prediction branch to regress the boundary, and a mask prediction branch to produce mask prediction for each input image. Previous methods often perform unsatisfactorily in the accurate prediction of boundaries in the target domain, particularly in areas with indistinct boundaries between different structures, as a result of domain shift between source and target domain, combined with low-intensity contrast between the various structures. Incorporating boundary loss has been demonstrated to be an effective approach in mitigating the performance degradation caused by inaccurate prediction of boundaries in the target domain [38].

In the source domain, we train a deep neural network f_s by minimizing the standard binary cross-entropy loss and boundary loss. Given a labeled image from the source domain, we denote \hat{y}_m^i and \hat{y}_b^i as its mask prediction and boundary prediction respectively from the source model, and y_m^i and y_b^i as the ground truth of the mask and boundary, respectively. To optimize the network, we calculate the mask prediction loss L_m and the boundary regression loss L_b on the source domain images. The total loss of the segmentation network in the source domain is:

$$L = L_m + L_b,$$

$$L_m = L_{BCE}(y_m^i, \hat{y}_m^i),$$

$$\text{and } L_b = \frac{1}{N} \sum_i (y_b^i - \hat{y}_b^i)^2. \quad (1)$$

Besides a standard source model f_s , we also train a robust source model f_s^r . Specifically, the only difference of f_s^r is that we augment the training data with adversarial examples generated by Projected Gradient Descent (PGD) [26] for robust training to improve robustness.

3.2. Target Domain Model Training

In the target domain, the only available components are source models and target data $x_t^i \sim D_t$. We adapt the standard source model f_s and robust source model f_s^r to the target domain separately so that we obtain two target models f_t and f_t^r . Standard and robust target models are trained in the same way, only with different initializations. In this section, we only describe details of standard target model training for simplicity.

Segmentation Loss with Pseudo Labels. As we do not have the ground truth of target data, we generate the pseudo labels y_{PL} and pseudo boundaries y_{PB} from the aforementioned source model. Given an unlabeled image from the target domain, we denote p_v as the mask prediction probability on v -th pixel obtained from the source model, then use \hat{p}_v and \hat{b}_v to denote the mask

Table 1. Quantitative comparison of different methods on the target domain datasets, including both adversarial and clean images. (Note: D denotes Dice, R. stands for robust, i.e. adversarial metrics that are tested on the adversarial samples, - means the result is not reported by that method, S-F denotes source-free.) All our methods have higher adversarial values compared to the baselines. And the performance of our methods on clean samples is comparable and mostly higher than the other methods. The best value is presented in bold and the second best is underlined. (D \uparrow : %, ASD \downarrow : pixel.)

Dataset	Method	S-F	Optic disc segmentation				Optic cup segmentation			
			D[%]	ASD	R.D[%]	R.ASD	D[%]	ASD	R.D[%]	R.ASD
	W/o adaptation		83.02	23.36	-	-	73.10	13.87	-	-
	Oracle		95.74	6.05	-	-	83.97	5.38	-	-
RIM-ONE-r3	AdvEnt [36]	✗	86.76	12.36	71.30	35.48	74.48	11.43	46.01	43.07
	BEAL [38]	✗	87.54	19.96	49.77	27.04	74.62	12.88	38.32	<u>29.95</u>
	OCDA [22]	✗	86.47	16.76	62.86	49.17	76.74	10.94	52.11	53.76
	DPL [5]	✓	90.13	9.43	72.78	35.43	76.78	9.01	52.29	40.85
	TT-SFUda [35]	✓	85.00	17.05	62.84	61.22	76.62	10.31	43.60	61.21
	TENT [37]	✓	82.92	23.63	61.72	50.47	72.95	14.00	45.70	52.92
	SFDA-FSM [39]	✓	82.98	23.69	73.60	37.59	73.56	14.51	54.27	30.65
	Ours (Standard Source)	✓	<u>92.43</u>	<u>6.96</u>	74.35	35.58	<u>77.82</u>	<u>10.03</u>	55.91	40.85
	Ours (Robust Source)	✓	91.39	8.07	77.22	<u>23.63</u>	75.59	11.13	60.34	50.11
	Ours (Both)	✓	92.89	6.52	<u>76.20</u>	22.92	77.94	10.07	<u>60.13</u>	26.83
	W/o adaptation		94.04	7.47	-	-	80.22	13.47	-	-
	Oracle		97.40	3.58	-	-	90.10	9.50	-	-
Drishti-GS	AdvEnt [36]	✗	95.85	4.65	91.18	10.81	78.27	14.71	65.52	25.29
	BEAL [38]	✗	95.74	8.32	72.34	20.57	86.45	16.65	57.47	18.47
	OCDA [22]	✗	95.41	5.56	91.63	11.43	79.95	13.68	67.56	22.98
	DPL [5]	✓	<u>96.39</u>	4.88	93.02	9.58	83.53	11.39	75.99	16.35
	TT-SFUda [35]	✓	95.22	6.00	91.64	12.27	80.67	13.00	64.01	24.01
	TENT [37]	✓	94.06	7.56	91.32	13.32	80.12	13.52	72.17	18.77
	SFDA-FSM [39]	✓	93.83	7.76	90.26	17.22	83.19	11.95	70.21	16.09
	Ours (Standard Source)	✓	96.01	<u>4.70</u>	92.05	11.25	<u>83.71</u>	10.91	73.92	18.15
	Ours (Robust Source)	✓	95.67	5.09	95.55	5.13	82.87	11.55	<u>78.08</u>	<u>15.11</u>
	Ours (Both)	✓	96.51	4.01	<u>95.29</u>	<u>5.25</u>	83.56	<u>11.11</u>	80.02	13.69
	W/o adaptation		70.36	52.75	-	-	71.79	26.09	-	-
	Oracle		95.47	6.14	-	-	88.82	4.20	-	-
REFUGE val (Open)	AdvEnt [36]	✗	67.73	35.59	41.25	53.65	60.93	26.10	16.47	85.11
	BEAL [38]	✗	72.21	52.94	24.47	58.67	64.64	47.34	5.35	42.79
	OCDA [22]	✗	85.38	29.70	48.71	58.98	79.14	14.20	37.44	63.59
	DPL [5]	✓	85.48	8.23	52.74	54.88	72.14	15.11	41.33	59.12
	TT-SFUda [35]	✓	82.60	33.00	50.80	51.73	77.86	12.56	41.61	56.53
	TENT [37]	✓	67.73	35.59	41.25	53.65	60.93	26.10	16.47	85.11
	SFDA-FSM [39]	✓	80.64	8.33	77.27	35.39	78.19	8.91	40.03	39.80
	Ours (Standard Source)	✓	91.54	6.87	88.51	17.68	79.78	7.44	46.46	37.46
	Ours (Robust Source)	✓	<u>91.86</u>	<u>6.77</u>	<u>91.11</u>	<u>8.40</u>	80.40	<u>7.14</u>	72.31	<u>25.71</u>
	Ours (Both)	✓	92.53	6.54	91.35	8.11	80.31	7.12	<u>71.05</u>	20.87

probability and boundary prediction obtained from the target model. The pseudo label can be generated as: $y_{PL_v} = 1[p_v \geq t]$, where $1(\cdot)$ is the indicator function, $t \in (0, 1)$ is a probability threshold to generate binary pseudo labels for the segmentation task. And we alleviate the noise of the pseudo labels by using the label selection mask m_v [5], in which a pseudo label is selected when the network’s uncertainty on prediction is low and the encoded feature lies closer to the object prototype than the background prototype:

$$L_{seg} = - \sum_v m_v * [y_{PL_v} \cdot \log(\hat{p}_v) + (1 - y_{PL_v}) \cdot \log(1 - \hat{p}_v)]. \quad (2)$$

Pseudo-Boundary Loss. In semantic segmentation, accurate boundary prediction between objects is critical.

To encourage the model to predict more accurately, we use pseudo-boundaries obtained from a source model during the training. Explicitly modeling the boundaries helps the model generalize better to unseen target domains and generate more precise and coherent predictions, which can further improve the accuracy. Additionally, boundaries of foreground objects are domain-invariant information in different domains.

The pseudo boundary segmentation loss is defined as:

$$L_{bl} = \frac{1}{N} \sum_v (y_{PB_v} - \hat{b}_v)^2. \quad (3)$$

Note that using robust pseudo labels and pseudo boundaries from a robust model can result in less accuracy, as the clean accuracy of the standard model is generally higher than that of the robust model. On the one hand, we hope to transfer the source robustness using a robust source model. On

Table 2. Quantitative comprehensive comparisons of performance on **adversarial and clean** images on the test data averaged over all domain and over all OC and OD segmentation. C+O denotes overall results averaged on the compound and open target domains.

Method	Compound(C)				Open(O)		Avg.			
	RIM-ONE-r3		Drishti-GS		REFUGE val		C		C + O	
	D[%]	ASD	D[%]	ASD	D[%]	ASD	D[%]	ASD	D[%]	ASD
AdvEnt [36]	69.64	25.58	82.71	13.86	46.60	50.11	76.17	19.72	66.31	29.85
BEAL [38]	62.56	22.46	78.00	16.00	41.67	50.44	70.28	15.47	60.74	27.13
OCDA [22]	69.55	32.66	83.64	13.41	62.67	41.62	76.59	23.03	71.95	29.23
DPL [5]	73.00	23.68	87.23	10.35	62.92	34.33	80.11	17.02	74.38	22.79
TT-SFUDA [35]	67.02	37.45	82.89	13.82	63.22	38.45	74.95	25.64	71.04	29.91
TENT [37]	65.82	35.26	84.42	13.29	67.62	30.58	75.12	24.27	72.62	26.38
SFDA-FSM [39]	71.10	26.61	84.87	13.25	69.03	23.11	77.99	19.93	75.00	20.99
Ours (Standard Source)	75.13	24.49	86.42	11.25	76.57	17.36	80.78	17.87	79.37	17.70
Ours (Robust Source)	<u>76.14</u>	<u>23.24</u>	<u>88.04</u>	<u>9.22</u>	83.92	<u>12.01</u>	<u>82.09</u>	<u>16.23</u>	<u>82.70</u>	<u>14.82</u>
Ours (Both)	76.79	16.59	88.85	8.52	83.81	10.66	82.82	12.55	83.15	11.92

the other hand, we require better pseudo-labels to generate adversarial examples. Therefore, to balance the trade-off between clean and robust accuracy, we use both robust and standard source models and transfer them to the target domain. In Ours (*Both*) method we obtain the required non-robust pseudo labels and pseudo boundaries using the standard source model f_s , then adapt the robust source model f_s^r to the target domain. We use the non-robust parts from f_s and prediction from f_t^r to calculate the loss, as shown in in Figure 2.(b). And only using standard source model f_s (Ours (*Standard Source*)) is shown as Standard Flow in Figure 3, while only using robust source model f_s^r (Ours (*Robust Source*)) is not shown but the only difference with standard flow is that only robust source model is available.

Entropy Minimization. Entropy minimization is a useful technique for UDA [2, 37], and minimizing entropy increases the confidence of the network’s output. We apply this technique in our method to further improve the adaptation performance and encourage the target model to output confident predictions on unlabeled data. The Shannon entropy for a prediction probability is defined as:

$$L_{ent} = - \sum_v \hat{p}_v \cdot \log \hat{p}_v. \quad (4)$$

\hat{p}_v denotes the mask probability obtained from the target model. To optimize the target segmentation model f_t , the overall loss function is:

$$L = L_{seg} + \alpha L_{bl} + \beta L_{ent}, \quad (5)$$

where α and β are hyper-parameters for balancing the effect between the pseudo-boundary loss and entropy minimization loss.

4. Experiments

4.1. Datasets

We utilize four public optic disc (OD) and optic cup (OC) segmentation datasets. Specifically, we use the training set of the REFUGE challenge [27] as the source domain,

RIM-ONE-r3 [7] and Drishti-GS [33] as the compound target domains, and the testing set of REFUGE val [27] as the open domain. The source domain consists of 400 annotated training images, and two compound target domain data are split to 99/60 and 50/51 images for training/testing respectively, following the same setup in DPL [5] and BEAL [38]. The open domain consists of 80 images.

4.2. Performance Metrics

We employ two metrics to evaluate the segmentation performance: Dice coefficient and Average Surface Distance (ASD). The Dice coefficient measures pixel-wise segmentation accuracy, while ASD measures segmentation boundary accuracy. Higher Dice coefficients and lower ASD indicate better segmentation performance.

4.3. Implementation Details

We use standard data augmentation, including Gaussian noise, contrast adjustment, and random erasing, same with [5] to slightly disturb the inputs when training target models, in order to make the predictions deviate from pseudo labels. And we also use these augmentation strategies in the standard source model training stage. For all methods, we use a MobileNetV2 adapted DeepLabv3+ [6] as the backbone, similar as in [5]. The threshold t is 0.75, same with BEAL [38]. We train the source and target model both with batch size 8, but 200 epochs and 2 epochs respectively. The model is implemented with PyTorch 1.12.1 and trained on one NVIDIA GeForce RTX 2080 Ti GPU. For generating adversarial images, we set the number of PGD steps to 20 and ϵ to 4/255. The loss components weights are set as $\alpha = 1.0$, $\beta = 0.4$.

4.4. Comparisons with the State-of-the-Art

We conducted extensive experiments to compare our method with recent UDA methods, including BEAL [38], DPL [5], AdvEnt [36], TT-SFUDA [35], TENT [37] and OCDA [22] on fundus image datasets. BEAL [38] is an UDA method with adversarial learning between

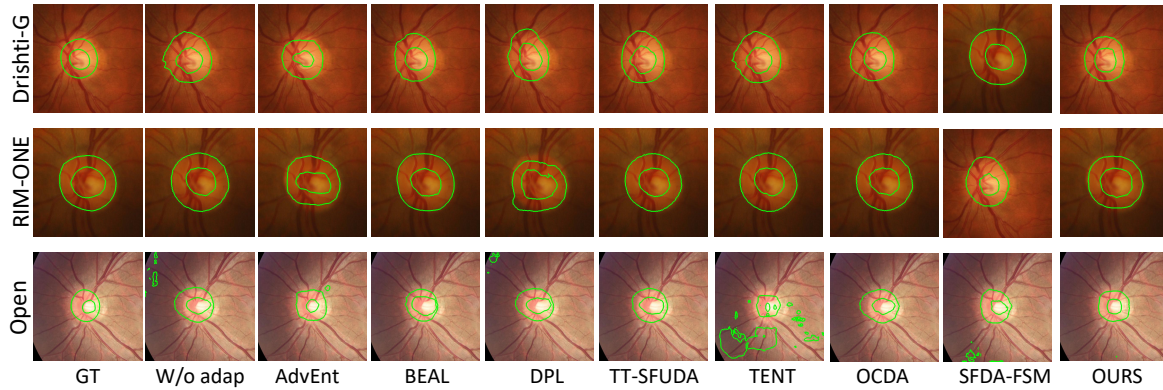


Figure 4. Comparison of adaptation performance of different methods on **clean** samples.




Figure 5. Comparison of adaptation performance of different methods on **adversarial** samples.

source and target data with boundary prediction on cross-domain fundus image segmentation; DPL [5] is a state-of-the-art SFUDA method with denoised pseudo labeling, which is the best-reported SFUDA model on cross-domain fundus image segmentation; AdvEnt [36] is a popular UDA benchmark approach that encourages entropy consistency between the source and target domains; TT-SFUDA [35] is a state-of-the-art two-stage approach for source-free domain adaptive image segmentation, including target-specific adaptation and task-specific adaptation; TENT [37] is a popular test-time adaptation method which optimizes the model by minimizing entropy of its prediction; OCDA [22] is a memory-based curriculum learning framework that improves generalization on the compound and open domains. 1) a curriculum domain adaptation strategy to bootstrap generalization across domain distinction in a data-driven self-organizing fashion and 2) a memory module to increase the model’s agility towards novel domains. SFDA-FSM [39] is a two-stage state-of-the-art two-stage approach for source-free domain adaptive image segmentation including a Fourier Style Mining (FSM) generation stage and an adaptation stage which has a Contrastive Domain Distillation (CDD) module to achieve feature-level

adaptation. The detailed comparison results are presented in Table 1, where we also include “**W/o adaptation**” lower bound and the supervised upper bound “**Oracle**” in the target domain. While SOTA DPL performs well on clean samples, our two best approaches (Robust source and Both) significantly improve adversarial performance while maintaining performance on clean data. Notably, without source data during adaptation, our method (Both) achieves the highest Dice for OD segmentation on three target domains.

The quantitative comparisons of different methods averaged on clean and adversarial images are shown in Table 2 and the Dice and ASD are averaged over OC and OD segmentation. Ours (*Both*) outperforms all SF-DA methods on clean samples, especially the state-of-the-art source-free DPL by improving C+O Dice by 3.22%. Our method (*Both*) also outperforms all UDA methods on adversarial samples, especially the baseline source-dependent BEAL by improving C Dice by 23.44% and C+O Dice by 37.72%. Moreover, when the gap between the source and target domain is trivial, as Drishti-GS, using non-robust (*Both*) or robust pseudo parts (*Robust source*) does not cause a big difference. Otherwise, non-robust counterparts are suggested.

Apart from closed-set UDA, we also demonstrate the ef-

fectiveness of our method for the open domain. As shown in Table 2, the proposed PLPB achieves the best accuracy and even outperforms the standard UDA method [22]. The improvements for open-set tasks are significant, as low-level boundary information of the source model is utilized sufficiently to address the distribution shift and the robust model generalizes well on out-of-distribution data.

We again conduct experiments of all the baseline approaches with adversarial training [26] in the source domain, as shown in Table S1 in Supplementary. Ours (*Both*) still have the best overall performance compared with robust methods. And it is important to note that the clean metrics of other robust methods drop considerably if they are directly trained robustly, which is in line with the general observation that robust models tend to hurt the performance on clean samples [42].

Figure 4 shows the qualitative comparison of adaptation performance on clean samples. Other methods hardly predict accurate boundaries in ambiguous regions and our method produces more accurate boundaries and mask predictions. Comparison of adversarial samples is shown in Figure 5. Other methods have collapsed segmentation results on adversarial samples, while our PLPB remains well shapes thanks to robust augmentation and boundary information from the robust source model.

These results demonstrate the superiority of our method to adapt the model even without the source data. This also indicates that SF-DA would not necessarily underperform the vanilla UDA. One possible reason is that vanilla UDA methods assume that the source and target domains share common features, which is not always the case in real-world scenarios. Moreover, vanilla UDA methods aim to find an invariant latent space between the source and target distribution which could be challenging. Our method directly adapts the model to the target domain, allowing it to capture more discriminative representations from target distribution and be more robust to domain shifts.

4.5. Ablation Studies

We study the effectiveness of key components in PLPB, with results shown in Table 3 and Table 4 for clean and adversarial samples respectively. Adding boundary loss significantly improves the source-only model. The entropy loss also contributes to improvements in terms of Dice score and ASD value. The comparison with Ours (Standard) and Ours (Robust) in Table 1 shows the superiority of using both standard and robust models. The results of ablation studies confirm the effectiveness of the components in PLPB.

4.6. Discussion

Based on our experimental evaluations, we find that consistent with previous works [1, 32], robust source feature extractors are effective in preserving robustness, while main-

Table 3. Ablation results with different losses on **clean** samples.

Method	Compound(C)				Open(O)		Avg.			
	RIM-ONE-r3		Drishiti-GS		REFUGE val		C		C + O	
	D[%]	ASD	D[%]	ASD	D[%]	ASD	D[%]	ASD	D[%]	ASD
L_{seg}	83.85	9.88	89.30	8.12	85.14	7.72	86.57	9.00	86.09	8.57
$L_{seg} + \alpha L_{bl}$	84.12	9.27	89.76	7.79	85.76	7.26	86.94	8.53	86.55	8.10
$L_{seg} + \beta L_{ent}$	84.08	8.75	89.93	7.84	85.30	8.33	87.00	8.30	86.44	8.31
$L_{seg} + \alpha L_{bl} + \beta L_{ent}$	85.42	8.30	90.04	7.56	86.42	6.83	87.73	7.93	87.29	7.56

Table 4. Ablation results with different losses on **adversarial** samples.

Method	Compound(C)				Open(O)		Avg.			
	RIM-ONE-r3		Drishiti-GS		REFUGE val		C		C + O	
	D[%]	ASD	D[%]	ASD	D[%]	ASD	D[%]	ASD	D[%]	ASD
L_{seg}	67.98	29.84	86.45	10.20	81.08	17.45	77.22	20.02	78.50	19.66
$L_{seg} + \alpha L_{bl}$	68.03	26.60	86.98	9.97	81.09	18.33	77.50	18.29	78.70	18.30
$L_{seg} + \beta L_{ent}$	68.57	28.72	86.38	11.61	80.93	18.86	77.47	20.16	78.62	19.73
$L_{seg} + \alpha L_{bl} + \beta L_{ent}$	68.17	24.88	87.66	9.47	81.20	14.49	77.91	17.17	79.01	16.28

taining sufficiently high accuracy on clean samples. And robust pre-trained models not only perform well on targets without adversarial training but also improve the accuracy on clean samples. Furthermore, adversarial training on source training enhances the ability of adaptation by improving the low-level diversity and generalization on the open dataset. We assume that adversarial attacks are based on the gradient of the training model, which can be considered as a type of low-level diversity so that adversarial samples can also serve as augmentation for diversity. It is consistent with the previous work that a pre-trained model that is more robust to input perturbation provides a better initialization for generalization on downstream out-of-distribution data [41]. After disabling the pseudo-boundary loss, the performance of the target model has dropped noticeably. By explicitly modeling the boundaries in the target domain, our target model achieves better generalization and generates more precise and coherent predictions.

5. Conclusion

In this study, we develop a method for robust source-free domain adaptation for the segmentation of fundus images. Our method PLPB outperforms competing state-of-the-art methods, achieving SOTA performance on adversarial and clean samples. A limitation of our method is that we need to train two source models, which increases training costs. In future studies, we will address this limitation by using a single source model. We will also explore methods to handle more severe domain shifts.

6. Acknowledgement

The work was supported in part by the National Natural Science Foundation of China (grants 31971289 & 91954201 to G.Y.) and the Strategic Priority Research Program of the Chinese Academy of Sciences (grant XDB37040402 to G.Y.).

References

- [1] Peshal Agarwal, Danda Pani Paudel, Jan-Nico Zaech, and Luc Van Gool. Unsupervised robust domain adaptation without source data. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 2009–2018, 2022. 1, 3, 8
- [2] Mathilde Bateson, Hoel Kervadec, Jose Dolz, Hervé Lombaert, and Ismail Ben Ayed. Source-relaxed domain adaptation for image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 490–499. Springer, 2020. 6
- [3] Gerda Bortsova, Cristina González-Gonzalo, Suzanne C Wetstein, Florian Dubost, Ioannis Katramados, Laurens Hogeweg, Bart Liefers, Bram van Ginneken, Josien PW Pluim, Mitko Veta, et al. Adversarial attack vulnerability of medical image analysis systems: Unexplored factors. *Medical Image Analysis*, 73:102141, 2021. 1
- [4] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *Proceedings of the European conference on computer vision (ECCV)*, pages 185–200, 2018. 1
- [5] Cheng Chen, Quande Liu, Yueming Jin, Qi Dou, and Pheng-Ann Heng. Source-free domain adaptive fundus image segmentation with denoised pseudo-labeling. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 225–235. Springer, 2021. 1, 2, 3, 5, 6, 7
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 801–818, 2018. 6
- [7] Francisco Fumero, Silvia Alayón, José L Sanchez, Jose Sigut, and M Gonzalez-Hernandez. Rim-one: An open retinal image database for optic nerve evaluation. In *2011 24th International Symposium on Computer-based Medical Systems (CBMS)*, pages 1–6. IEEE, 2011. 6
- [8] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, pages 1180–1189. PMLR, 2015. 1, 2
- [9] Tejas Gokhale, Swaroop Mishra, Man Luo, Bhavdeep Singh Sachdeva, and Chitta Baral. Generalized but not robust? comparing the effects of data modification methods on out-of-domain generalization and adversarial robustness. *arXiv preprint arXiv:2203.07653*, 2022. 3
- [10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 3
- [11] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 1, 2
- [12] Konstantinos Kamnitsas, Christian Baumgartner, Christian Ledig, Virginia Newcombe, Joanna Simpson, Andrew Kane, David Menon, Aditya Nori, Antonio Criminisi, Daniel Rueckert, et al. Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In *International Conference on Information Processing in Medical Imaging*, pages 597–609. Springer, 2017. 1, 2
- [13] Jogendra Nath Kundu, Naveen Venkat, R Venkatesh Babu, et al. Universal source-free domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4544–4553, 2020. 1, 2
- [14] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial Intelligence Safety and Security*, pages 99–112. Chapman and Hall/CRC, 2018. 1, 3
- [15] Vinod K Kurmi, Venkatesh K Subramanian, and Vinay P Namboodiri. Domain impression: A source data free domain adaptation method. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 615–625, 2021. 2
- [16] Yingwei Li, Zhuotun Zhu, Yuyin Zhou, Yingda Xia, Wei Shen, Elliot K Fishman, and Alan L Yuille. Volumetric medical image segmentation: a 3d deep coarse-to-fine framework and its adversarial examples. In *Deep Learning and Convolutional Neural Networks for Medical Imaging and Clinical Informatics*, pages 69–91. Springer, 2019. 1
- [17] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 6028–6039. PMLR, 2020. 2
- [18] Kai Lin, Thomas H Li, Shan Liu, and Ge Li. Real photographs denoising with noise domain adaptation and attentive generative adversarial network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1
- [19] Peng Liu, Charlie T Tran, Bin Kong, and Ruogu Fang. Cada: Multi-scale collaborative adversarial domain adaptation for unsupervised optic disc and cup segmentation. *Neurocomputing*, 469:209–220, 2022. 2
- [20] Xiaofeng Liu, Fangxu Xing, Georges El Fakhri, and Jonghye Woo. Self-semantic contour adaptation for cross modality brain tumor segmentation. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022. 2
- [21] Yuang Liu, Wei Zhang, and Jun Wang. Source-free domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1215–1224, 2021. 1, 2, 4
- [22] Ziwei Liu, Zhongqi Miao, Xingang Pan, Xiaohang Zhan, Dahua Lin, Stella X Yu, and Boqing Gong. Open compound domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12406–12415, 2020. 1, 5, 6, 7, 8
- [23] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning*, pages 97–105. PMLR, 2015. 2
- [24] Xingjun Ma, Yuhao Niu, Lin Gu, Yisen Wang, Yitian Zhao, James Bailey, and Feng Lu. Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recognition*, 110:107332, 2021. 1

- [25] Xingjun Ma, Yuhao Niu, Lin Gu, Yisen Wang, Yitian Zhao, James Bailey, and Feng Lu. Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recognition*, 110:107332, 2021. 3
- [26] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 3, 4, 8
- [27] José Ignacio Orlando, Huazhu Fu, João Barbosa Breda, Karel van Keer, Deepti R Bathula, Andrés Diaz-Pinto, Ruogu Fang, Pheng-Ann Heng, Jeyoung Kim, JoonHo Lee, et al. Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical Image Analysis*, 59:101570, 2020. 6
- [28] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3764–3773, 2020. 1
- [29] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*, pages 372–387. IEEE, 2016. 1, 3
- [30] Magdalini Paschali, Sailesh Conjeti, Fernando Navarro, and Nassir Navab. Generalizability vs. robustness: investigating medical imaging networks using adversarial examples. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 493–501. Springer, 2018. 3
- [31] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? *Advances in Neural Information Processing Systems*, 33:3533–3545, 2020. 3
- [32] Ali Shafahi, Parsa Saadatpanah, Chen Zhu, Amin Ghiasi, Christoph Studer, David Jacobs, and Tom Goldstein. Adversarially robust transfer learning. *arXiv preprint arXiv:1905.08232*, 2019. 3, 8
- [33] Jayanthi Sivaswamy, S Krishnadas, Arunava Chakravarty, G Joshi, A Syed Tabish, et al. A comprehensive retinal image dataset for the assessment of glaucoma from the optic nerve head analysis. *JSM Biomedical Imaging Data Papers*, 2(1):1004, 2015. 6
- [34] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018. 1, 2
- [35] Vibashan VS, Jeya Maria Jose Valanarasu, and Vishal M Patel. Target and task specific source-free domain adaptive image segmentation. *arXiv preprint arXiv:2203.15792*, 2022. 2, 5, 6, 7
- [36] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019. 1, 2, 5, 6, 7
- [37] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2020. 5, 6, 7
- [38] Shujun Wang, Lequan Yu, Kang Li, Xin Yang, Chi-Wing Fu, and Pheng-Ann Heng. Boundary and entropy-driven adversarial learning for fundus image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 102–110. Springer, 2019. 1, 4, 5, 6
- [39] Chen Yang, Xiaoqing Guo, Zhen Chen, and Yixuan Yuan. Source free domain adaptation for medical image segmentation with fourier style mining. *Medical Image Analysis*, 79:102457, 2022. 1, 3, 5, 6, 7
- [40] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 1, 2
- [41] Mingyang Yi, Lu Hou, Jiacheng Sun, Lifeng Shang, Xin Jiang, Qun Liu, and Zhiming Ma. Improved ood generalization via adversarial training and pretraing. In *International Conference on Machine Learning*, pages 11987–11997. PMLR, 2021. 3, 8
- [42] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pages 7472–7482. PMLR, 2019. 8
- [43] Yuyang Zhao, Zhun Zhong, Zhiming Luo, Gim Hee Lee, and Nicu Sebe. Source-free open compound domain adaptation in semantic segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(10):7019–7032, 2022. 1, 3