

Striving for Simplicity: Simple Yet Effective Prior-Aware Pseudo-Labeling for Semi-Supervised Ultrasound Image Segmentation

Yaxiong Chen^{1,2}, Yujie Wang^{1*}, Zixuan Zheng³, Jingliang Hu³, Yilei Shi³, Shengwu Xiong^{1,2}, Xiao Xiang Zhu⁴, and Lichao Mou^{3(✉)}

¹ Wuhan University of Technology, Wuhan, China

² Shanghai Artificial Intelligence Laboratory, Shanghai, China

³ MedAI Technology (Wuxi) Co. Ltd., Wuxi, China
lichao.mou@medimagingai.com

⁴ Technical University of Munich, Munich, Germany

Abstract. Medical ultrasound imaging is ubiquitous, but manual analysis struggles to keep pace. Automated segmentation can help but requires large labeled datasets, which are scarce. Semi-supervised learning leveraging both unlabeled and limited labeled data is a promising approach. State-of-the-art methods use consistency regularization or pseudo-labeling but grow increasingly complex. Without sufficient labels, these models often latch onto artifacts or allow anatomically implausible segmentations. In this paper, we present a simple yet effective pseudo-labeling method with an adversarially learned shape prior to regularize segmentations. Specifically, we devise an encoder-twin-decoder network where the shape prior acts as an implicit shape model, penalizing anatomically implausible but not ground-truth-deviating predictions. Without bells and whistles, our simple approach achieves state-of-the-art performance on two benchmarks under different partition protocols. We provide a strong baseline for future semi-supervised medical image segmentation. Code is available at <https://github.com/WUTCM-Lab/Shape-Prior-Semi-Seg>.

Keywords: semi-supervised learning · segmentation · pseudo-labeling.

1 Introduction

Medical ultrasound imaging is one of the most ubiquitous diagnostic modalities used in clinical practice today owing to its non-invasive, radiation-free, and real-time nature [1,2]. This ubiquity has resulted in an explosion of ultrasound images. However, manual screening by clinical experts struggles to keep pace. Automated ultrasound image segmentation models can improve efficiency and consistency, but require large pixel-wise labeled datasets which are particularly scarce in the ultrasound domain compared to natural images or even other medical imaging modalities like CT and MRI [3].

* Work done during an internship at MedAI Technology (Wuxi) Co. Ltd.




Fig. 1. Examples of unrealistic and unnatural segmentations produced by a UNet.

A common *modus operandi* is to use semi-supervised learning, that is, leveraging both abundant unlabeled medical images and limited labeled ones simultaneously [4]. The majority of current state-of-the-art methods follow two schemes: consistency regularization and pseudo-labelling [5]. The consistency regularization-based approaches incorporate a regularization term in the loss to minimize the discrepancy between different predictions for the same image, which are derived by introducing perturbations to the input image or to the models involved. [6] proposes a network with two identical branches, each of which receives the same image with different perturbations, and designs a loss to encourage predictions of the two branches to be consistent. Subsequently, the authors of [7] introduce more branches, devise distinct perturbations for each of them, and enforce consistency regularization across branches. Further, [8] presents a sophisticated attention mechanism for the multi-branch architecture. On the other hand, there has been a significant amount of recent work based on pseudo-labelling. For instance, [9] proposes a method called mean teacher, which adopts a teacher-student architecture for semi-supervised learning. This framework comprises a teacher model and a student model. The student model is trained on unlabeled data using pseudo-labels generated by the teacher, and the teacher model itself is an exponential moving average of the student model’s weights. In [10], the authors present mutual learning based on the mean teacher framework and introduce an additional self-correction mechanism to refine pseudo-labels. [11] further devises an error localization module to help identify incorrect pixels in pseudo labels. Nevertheless, albeit successful, these state-of-the-art methods tend to increasingly complex designs as diverse as more network modules and additional losses.

Moreover, we observe that without sufficient labeled ultrasound images, models often latch onto artifacts rather than meaningful morphological cues. In addition, the flexibility of deep networks to fit arbitrary boundaries allows unrealistic and unnatural segmentations that violate basic anatomical principles (cf. Fig. 1). These pose challenges for semi-supervised image segmentation approaches. For example, consistency regularization-based methods encourage predictions to be invariant to perturbations, but are unable to correct errors due to anatomically implausible boundaries in the absence of ground truth constraints. The pseudo-labeling methodology relies on predictions from the teacher model to generate targets for the student, which may reinforce erroneous cues in areas of high uncertainty. To alleviate the issues, one possible solution is to utilize shape information. [12] and [13] are motivated to improve segmentation masks generated




Fig. 2. Pipeline of the proposed semi-supervised method for ultrasound image segmentation.

by fully supervised networks by better matching them to ground truths, but they do not take advantage of anatomical shape prior knowledge, which is crucial for semi-supervised ultrasound image segmentation where boundaries of targets are often ambiguous.

In this paper, we break the trend of recent state-of-the-art approaches that combine increasingly complex techniques, and introduce a simple and clean pseudo-labeling-based network. To address the problems caused by artifacts and ambiguous boundaries of lesions in ultrasound images, we learn a shape prior and inject this prior into the network. Specifically, our contributions are three-fold.

- Without complex designs, a simple yet effective encoder-twin-decoder network architecture is devised, which can readily achieve better performance.
- By integrating an adversarially learned shape prior, we push the limit of semi-supervised ultrasound image segmentation. The learned prior serves as a regularizer, penalizing the network only if its outputs are unrealistic, not if it deviates from ground truths.
- We provide a simple yet strong baseline for future works. Extensive experiments and ablation studies are conducted on two public ultrasound datasets to demonstrate its effectiveness.

2 Method

We first present the overall structure of our proposed network in Section 2.1. In Section 2.2, we then introduce shape prior modeling and the shape regularizer that is incorporated into the network. Finally, the details of the training process for the entire network are described in Section 2.3.

2.1 Semi-Supervised Image Segmentation Network

We propose a semi-supervised ultrasound image segmentation network that concurrently leverages labeled and unlabeled data for training. As illustrated in Fig. 2, our network is composed of three key components—an encoder $\mathcal{E}$, twin decoders $\mathcal{D}_l$ and $\mathcal{D}_p$, and a deep shape regularizer (DSR). $\mathcal{E}$ is a ResNet [17] that maps an input image into feature representations. The twin decoders, $\mathcal{D}_l$ and $\mathcal{D}_p$, share an identical architecture with five convolutional blocks, each containing two Conv-BN-ReLU layers, but serve distinct purposes. Specifically, $\mathcal{D}_l$ focuses exclusively on learning from limited labeled data to produce pseudo-labels for supervising $\mathcal{D}_p$. In contrast, $\mathcal{D}_p$ functions in a prior-guided manner to segment unlabeled images by integrating guidance from the pseudo-labels and regularization from DSR, without accessing any labeled data during training. Notably, only $\mathcal{E}$ and $\mathcal{D}_p$ are required at inference time. In essence, the proposed method seamlessly integrates semi-supervised learning and inductive bias modeling for segmentation in an end-to-end framework.

2.2 Shape Prior Modeling and Shape Regularization

Shape Prior Modeling We pre-train a generative adversarial network (GAN) to determine whether a shape is anatomically plausible. The GAN contains a generator that synthesizes shape masks from a randomly sampled latent vector, and a discriminator that distinguishes the generated masks from real ones. Through this adversarial training process, the discriminator learns to model the distribution of real shapes and has the ability to quantify the plausibility of any input mask. We therefore utilize the trained discriminator as a regularizer, namely DSR, for subsequent semi-supervised image segmentation. Note that the generator is discarded after pre-training.

Specifically, we use five transposed convolutional layers with batch normalization (BN) and ReLU for the generator, and five convolutional layers with leaky ReLU for the discriminator. Due to the limited number of ground truth segmentation masks, stably training the GAN is challenging. Additionally, the high resolution of the masks leads to training instability. To mitigate this, we uniformly resize all masks to 64×64 , which does not alter the distribution of shapes yet eases optimization. Further, following [14], we use the Wasserstein loss with a gradient penalty term to train our GAN. The loss is defined as:

$$\mathcal{L}_{\text{spm}} = \mathbb{E}_{\tilde{x} \sim P_g} [D(\tilde{x})] - \mathbb{E}_{x \sim P_r} [D(x)] + \lambda \mathbb{E}_{\hat{x} \sim P_x} [(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2], \quad (1)$$

where D is the discriminator, and P_g and P_r represent distributions of the generated and real masks, respectively. The last term refers to the gradient penalty introduced in [14].

DSR At training time, the discriminator aims to minimize $\mathcal{L}_{\text{spm}}$, while the generator attempts to maximize $\mathcal{L}_{\text{spm}}$ by fooling the discriminator via the first

term in Eq. (1). After training, the discriminator gains the ability to assess the plausibility of shape masks. Therefore, we can define DSR as

$$\ell_{\text{dsr}} = - \mathbb{E}_{\tilde{x} \sim P_m} [D(\tilde{x})], \quad (2)$$

where P_m denotes the distribution of segmentation masks predicted by the semi-supervised image segmentation network.

2.3 Training Strategy and Loss

We detail the training process of the proposed semi-supervised image segmentation network. In each iteration, we sample a batch of labeled images $\mathcal{B}_s = \{(x_i, y_i)\}_{i=1}^{|\mathcal{B}_s|}$ and a batch of unlabeled images $\mathcal{B}_u = \{(u_i)\}_{i=1}^{|\mathcal{B}_u|}$.

First, we pass each labeled image $x_i \in \mathcal{B}_s$ through the encoder $\mathcal{E}$ to attain feature maps, which are then fed into the decoder $\mathcal{D}_l$ to generate a segmentation mask. A cross entropy loss is used to measure the discrepancy between this mask and the corresponding ground truth:

$$\mathcal{L}_s = \frac{1}{|\mathcal{B}_s|} \sum_{i=1}^{|\mathcal{B}_s|} \ell_{\text{ce}}(\mathcal{D}_l \circ \mathcal{E}(x_i), y_i), \quad (3)$$

where $\circ$ is a composition function.

Next, for each unlabeled image $u_i \in \mathcal{B}_u$, the encoder $\mathcal{E}$ extracts feature representations, and $\mathcal{D}_l$ is used to produce a pseudo segmentation mask. Afterwards, we apply dropout on the feature representations and pass the perturbed features through $\mathcal{D}_p$ to predict a segmentation mask. A consistency loss is exploited to enforce agreement between the two masks. In addition, DSR evaluates the plausibility of the predicted segmentation map and outputs a score that imposes a constraint on learning $\mathcal{D}_p$:

$$\mathcal{L}_u = \frac{1}{|\mathcal{B}_u|} \sum_{i=1}^{|\mathcal{B}_u|} \ell_{\text{ce}}(\mathcal{D}_p \circ m_i \odot \mathcal{E}(u_i), \mathcal{D}_l \circ \mathcal{E}(u_i)) + \lambda \ell_{\text{dsr}}(\mathcal{D}_p \circ m_i \odot \mathcal{E}(u_i)), \quad (4)$$

where λ controls the strength of the regularization, and m_i is the dropout's mask.

Finally, the loss of our method is defined as

$$\mathcal{L} = \mathcal{L}_s + \gamma \mathcal{L}_u. \quad (5)$$

3 Experiments

3.1 Datasets and Evaluation Metrics

To validate the effectiveness and robustness of our network, we conduct experiments on two ultrasound image segmentation datasets. The TN3K dataset [15]

Table 1. Comparison with state-of-the-art methods on the TN3K dataset under different partition protocols. The results of the fully supervised baseline are marked in gray .

	1/8		1/4		1/2	
Methods	Dice	IoU	Dice	IoU	Dice	IoU
UNet [18]	73.98	62.95	80.06	70.61	81.75	72.89
FixMatch [19]	61.59	49.36	66.09	53.71	69.44	57.74
SemiMedSeg [20]	77.33	67.11	80.67	71.44	81.79	72.71
U2PL [21]	74.35	63.15	75.60	64.26	75.25	63.89
ST++ [22]	52.51	41.12	57.61	46.05	65.21	53.54
AugSeg [23]	72.82	61.38	74.10	63.22	76.09	65.28
UniMatch [24]	67.30	55.25	69.33	57.29	69.67	57.88
MTNet [8]	61.26	51.33	68.77	58.45	73.72	63.62
SS-Net [25]	70.99	59.90	72.93	62.80	73.49	63.89
Ours	82.21	73.16	83.24	74.47	84.17	75.70

Table 2. Comparison with state-of-the-art methods on the BUSI dataset.

	1/8		1/4		1/2	
Methods	Dice	IoU	Dice	IoU	Dice	IoU
UNet [18]	67.89	57.68	77.01	67.60	77.06	70.88
FixMatch [19]	42.76	34.96	50.35	41.10	56.44	46.87
SemiMedSeg [20]	72.04	63.27	77.73	69.05	77.03	68.71
U2PL [21]	68.50	59.03	68.97	60.55	72.76	64.62
ST++ [22]	52.81	40.78	61.67	50.10	63.09	52.28
AugSeg [23]	64.78	54.76	68.54	59.24	72.03	63.64
UniMatch [24]	53.96	43.69	55.87	45.53	57.03	46.52
MTNet [8]	51.30	41.10	60.57	50.47	68.89	58.60
SS-Net [25]	61.71	53.30	66.90	57.32	70.61	62.26
Ours	75.19	66.90	79.43	71.17	80.41	72.02

contains 3,493 thyroid nodule ultrasound images, of which we use 614 for testing, 301 for validation, and the remaining 2,578 for training. The BUSI dataset [16] collects 780 breast ultrasound images with an average size of 500×500 pixels. The images are categorized into normal, benign, and malignant classes. We use the 647 benign and malignant images, with 576 in the training set and 71 in the test set.

Following common partition protocols in semi-supervised medical image segmentation, we perform experiments using 1/8, 1/4, and 1/2 of the total labeled data for each dataset. Moreover, we adopt two metrics, Dice similarity coefficient and intersection over union (IoU), to quantify segmentation performance.




Fig. 3. Qualitative comparison of segmentation results produced by our method and other competitors on the TN3K and BUSI test sets. The top two rows show examples from TN3K, while the bottom two rows show cases from BUSI.

3.2 Implementation Details

DSR Training We first train a GAN leveraging existing labels. Specifically, we resize all ground truth segmentation masks into a fixed size of 64×64 and set batch size to 16. We optimize the generator and discriminator using two RMSprop optimizers with a learning rate of 0.00005, and the total number of epochs is set to 5,000.

Segmentation Network Training After training DSR, we train the segmentation network by resizing all images to 320×320 pixels, followed by random cropping to 256×256 pixels and applying random data augmentation (rotation/scaling). We use a batch size of 16 and SGD with a momentum of 0.9 and a weight decay of 0.00001 to optimize the model. A polynomial learning rate schedule is adopted to adjust the learning rate $lr = init_lr \times \left(1 - \frac{iter}{max_iter}\right)^{power}$, where $init_lr = 0.001$ and $power = 0.9$. The total number of epochs is set to 200.

3.3 Comparison with State-of-the-Art Methods

In our experiments, we use a fully supervised UNet [18] (with ResNet34 backbone) as the baseline. We compare our approach with current state-of-the-art semi-supervised segmentation models including FixMatch [19], SemiMedSeg [20], U2PL [21], ST++ [22], AugSeg [23], UniMatch [24], MTNet [8], and SS-Net [25].

Table 3. Ablation study quantifying the contribution of DSR on the TN3K and BUSI datasets under three data partition protocols.

Datasets	Methods	1/8		1/4		1/2	
		Dice	IoU	Dice	IoU	Dice	IoU
TN3K	w/o DSR	81.19	72.17	82.26	73.34	83.08	74.63
	Full model	82.21	73.16	83.24	74.47	84.17	75.70
BUSI	w/o DSR	74.14	65.89	78.31	69.23	79.51	71.13
	Full model	75.19	66.90	79.43	71.17	80.41	72.02


**Fig. 4.** Qualitative assessment of our proposed DSR module on the TN3K and BUSI test sets. The top row shows an example from TN3K, while the bottom case is from BUSI.

Table 1 and Table 2 show experimental results on the TN3K and BUSI datasets. Our approach achieves the best performance, especially under sparse labeling (1/8). We improve 8.23% and 4.88% over the fully supervised baseline and second best semi-supervised image segmentation method on TN3K. Similar significant gains are observed on BUSI with limited labels.

Fig. 3 shows some visual examples. We can see that our approach produces high-quality segmentations on large objects, small objects, and with multiple instances, outperforming other methods which may occasionally produce segmentation errors.

3.4 Ablation Study

To validate the effectiveness of DSR, we conduct ablation studies on TN3K and BUSI. As summarized in Table 3, incorporating DSR consistently improves performance by approximately 1% on both datasets. This indicates that exploiting shape priors helps alleviate issues like porous. Qualitative results in Fig. 4 show-case segmentation masks with more natural, coherent shapes when using DSR.

4 Conclusion

In this work, we present a simple yet effective pseudo-labeling framework with an adversarially learned shape prior for semi-supervised medical ultrasound image segmentation. Unlike recent studies that tend to combine increasingly complex mechanisms, our approach utilizes a clean encoder-twin-decoder network architecture. We demonstrate the model on two public benchmarks and qualitatively and quantitatively establish the proficiency of the framework. The experimental results show that, without requiring complex architectures or augmentation strategies, our method achieves state-of-the-art performance under different data partition protocols.

Acknowledgments. This work is supported in part by the National Key Research and Development Program of China (2022ZD0160604), in part by the Natural Science Foundation of China (62101393/62176194), in part by the High-Performance Computing Platform of YZBSTCACC, and in part by MindSpore (<https://www.mindspore.cn>), a new deep learning framework.

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

References

1. Xie, Y., Liao, H., Zhang, D., Chen, F.: Uncertainty-aware cascade network for ultrasound image segmentation with ambiguous boundary. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 268–278. (2022)
2. Ning, Z., Zhong, S., Feng, Q., Chen, W., Zhang, Y.: SMU-Net: Saliency-guided morphology-aware U-Net for breast lesion segmentation in ultrasound image. *IEEE Transactions on Medical Imaging* **41**(2), 476–490 (2022)
3. Chen, F., Chen, L., Kong, W., Zhang, W., Zheng, P., Sun, L., Zhang, D., Liao, H.: Deep semi-supervised ultrasound image segmentation by using a shadow aware network with boundary refinement. *IEEE Transactions on Medical Imaging* **42**(12), 3779–3793 (2023)
4. Xu, X., Sanford, T., Turkbey, B., Xu, S., Wood, B.J., Yan, P.: Shadow-consistent semi-supervised learning for prostate ultrasound segmentation. *IEEE Transactions on Medical Imaging* **41**(6), 1331–1345 (2022)
5. Xu, M.C., Zhou, Y., Jin, C., de Groot, M., Alexander, D.C., Oxtoby, N.P., Hu, Y., Jacob, J.: Bayesian pseudo labels: Expectation maximization for robust and efficient semi-supervised segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 580–590. (2022)
6. Bortsova, G., Dubost, F., Hogeweg, L., Katramados, I., De Bruijne, M.: Semi-supervised medical image segmentation via learning consistency under transformations. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 810–818. (2019)
7. Quali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12674–12684. (2020)

8. Zhong, L., Liao, X., Zhang, S., Wang, G.: Semi-supervised pathological image segmentation via cross distillation of multiple attentions. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 570–579. (2023)
9. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in Neural Information Processing Systems, pp. 1195–1204. (2017)
10. Zhang, P., Zhang, B., Zhang, T., Chen, D., Wen, F.: Robust mutual learning for semi-supervised semantic segmentation. arXiv preprint arXiv:2106.00609 (2021)
11. Kwon, D., Kwak, S.: Semi-supervised semantic segmentation with error localization network. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9957–9967. (2022)
12. Oktay, O., Ferrante, E., Kamnitsas, K., Heinrich, M., Bai, W., Caballero, J., Cook, S.A., De Marvao, A., Dawes, T., O'Regan, D.P., Kainz, B., Glocker, B., Rueckert, D.: Anatomically constrained neural networks (ACNNs): Application to cardiac image enhancement and segmentation. *IEEE Transactions on Medical Imaging* **37**(22), 384–395 (2017)
13. Larrazabal, A.J., Martinez, C., Ferrante, E.: Anatomical priors for image segmentation via post-processing with denoising autoencoders. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 585–593. (2019)
14. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: Advances in Neural Information Processing Systems, pp. 5767–5777. (2017)
15. Gong, H., Chen, J., Chen, G., Li, H., Li, G., Chen, F.: Thyroid region prior guided attention for ultrasound segmentation of thyroid nodules. *Computers in Biology and Medicine* **155**, 106389 (2023)
16. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. *Data Brief* **28**, 104863 (2020)
17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 770–778. (2016)
18. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 234–241. (2015)
19. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: FixMatch: Simplifying semi-supervised learning with consistency and confidence. In: Advances in Neural Information Processing Systems, pp. 596–608. (2020)
20. Zhang, R., Liu, S., Yu, Y., Li, G.: Self-supervised correction learning for semisupervised biomedical image segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 134–144. (2021)
21. Wang, Y., Wang, H., Shen, Y., Fei, J., Li, W., Jin, G., Wu, L., Zhao, R., Le, X.: Semi-supervised semantic segmentation using unreliable pseudo-labels. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4248–4257. (2022)
22. Yang, L., Zhuo, W., Qi, L., Shi, Y., Gao, Y.: ST++: Make self-training work better for semi-supervised semantic segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4268–4277. (2022)

23. Zhao, Z., Yang, L., Long, S., Pi, J., Zhou, L., Wang, J.: Augmentation matters: A simple-yet-effective approach to semi-supervised semantic segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11350–11359. (2023)
24. Yang, L., Qi, L., Feng, L., Zhang, W., Shi, Y.: Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7236–7246. (2023)
25. Wu, Y., Wu, Z., Wu, Q., Ge, Z., Cai, J.: Exploring smoothness and class separation for semi-supervised medical image segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention, pp. 34–43. (2022)