

SHAN: Shape Guided Network for Thyroid Nodule Ultrasound Cross-Domain Segmentation

Ruixuan Zhang^{1,2,3,4}, Wenhuan Lu^{1,2}, Cuntai Guan⁴, Jie Gao^{1,2,3}, Xi Wei⁵,
and Xuewei Li^{1,2,3} (✉)

¹ College of Intelligence and Computing, Tianjin University, Tianjin, China

² Tianjin Key Laboratory of Cognitive Computing and Application, Tianjin, China

³ Tianjin Key Laboratory of Advanced Networking, Tianjin, China

{zrx_6566, wenhuan, gaojie, lixuewei}@tju.edu.cn

⁴ Nanyang Technological University, Singapore

ctguan@ntu.edu.sg

⁵ Tianjin Medical University Cancer Hospital, Tianjin, China

weixi@tmu.edu.cn

Abstract. Segmentation models for thyroid ultrasound images are challenged by domain gaps across multi-center data. Some methods have been proposed to address this issue by enforcing consistency across multi-domains or by simulating domain gaps using augmented single-domain. Among them, single-domain generalization methods offer a more universal solution, but their heavy reliance on the data augmentation causes two issues for ultrasound image segmentation. Firstly, the corruption in data augmentation may affect the distribution of grayscale values with diagnostic significant, leading to a decline in model's segmentation ability. The second is the real domain gap between ultrasound images is difficult to be simulated, resulting in features still correlate with domain, which in turn prevents the construction of the domain-independent latent space. To address these, given that the shape distribution of nodules is task-relevant but domain-independent, the SHape-prior Affine Network (SHAN) is proposed. SHAN serves shape prior as a stable latent mapping space, learning aspect ratio, size, and location of nodules through affine transformation of prior. Thus, our method enhances the segmentation capability and cross-domain generalization of model without any data augmentation methods. Additionally, SHAN is designed to be a plug-and-play method that can improve the performance of segmentation models with an encoder-decoder structure. Our experiments are performed on the public dataset TN3K and a private dataset TUI with 6 domains. By combining SHAN with several segmentation methods and comparing them with other single-domain generalization methods, it can be proved that SHAN performs optimally on both source and target domain data.

Keywords: Thyroid ultrasound · Shape prior · Segmentation · Generalization.

1 Introduction

The traditional mode of ultrasound diagnosis has been changed following the maturation of Deep Learning technology. More and more Computer Aided Diagnostic(CAD) systems [17] are involved in the ultrasonic diagnosis. Ultrasound image segmentation models suffer from domain shifts and have limited generalization across domains, as shown in Fig. 1(a). Domain Generalization(DG) is a method that improves the performance of the model on unseen domain (also called target domain). DG segmentation methods can be divided into Multi-DG (MDG, as shown in Fig. 1(b)-(d)) and Single-DG (SDG, as shown in Fig. 1(e)). MDG methods improve the generalization ability of the model by alignment, meta-learning, normalization, and disentanglement [20]. In contrast, SDG is a more general solution, since most of the SDG methods are able to cover the multi-domain situations. SDG usually treat images generated by data augmentation or style transfer methods as new domains to simulate the domain gap, and then models are trained based on MDG. Data augmentation methods for SDG include gamma correction and noise addition in BigAug [19], Bezier curve transformation in SLAug [16], photometric and geometric transforms in PDEN [10], frequency mixing in FreeSDG [9], global intensity non-linear augmentation in [14], etc.




Fig. 1. The domain gap and some domain generalization segmentation methods for thyroid ultrasound images. P_{XY}^i in blue denotes source data, P_{XY}^j in green denotes the unseen data. (a) shows the effect of domain gap on the model for nodule segmentation. (b)-(e) show some kinds of DG methods. (f) shows the idea of our method.

Most of the current segmentation methods for SDG rely on data augmentation methods, which have achieved impressive results in various scenarios. However, they still face two challenges in cross-domain segmentation tasks for ultrasound images. One is that existing methods overlook whether data augmentation affects the segmentation capability of the model, i.e., the corruption of grayscale values with diagnostic significance leads to a decline in the model's

performance in source domain. Second, the effectiveness of SDG is directly influenced by the data augmentation methods' ability to simulate the differences between real domains. Since the discrepancies between ultrasound images from different domains are mainly in resolution, probe scanning orientation, parameter settings and speckle noise, it is difficult to bridge the domain gap by simple data augmentation or style transfer methods. What we want to address is how to establish a stable domain-independent feature space without relying on data augmentation, so as to enhance the cross-domain generalization ability of the thyroid ultrasound image segmentation models while maintaining or even improving the segmentation performance on the source domain.

Considering the fact that lesions in thyroid ultrasound images are approximately elliptic-shaped, we desire that the segmentation model could learn the shape distribution of the lesion which is a task-relevant but domain-independent feature. Liu et al. [11] have shown the cross-domain stability of lesion shape. But it was designed for images with high consistency of anatomy, which is not applicable to ultrasound images which have variable sections and flexible angles.

Therefore, in this paper, a plug-and-play method SHape-prior Affine Network (SHAN) is proposed for the cross-domain segmentation task of ultrasound images. The SHAN consists of 1) Nodule Shape Prior, a segmentation task-relevant but domain-independent prior knowledge to constrain a stable domain-invariant feature space in the segmentation model. 2) Global Affine Module, which builds a stable affine mapping relationship between nodule shape distribution and latent feature by learning the aspect ratio, size, and location of benign and malignant nodules in the global view, estimates the shape of nodules preliminarily. 3) Neighbor Affine Module refines the preliminary shape estimation based on the subtle grayscale variation of neighboring pixels in the local view, and completes the segmentation of the nodule. The segmentation capability and cross-domain generalization ability of SHAN has been proved by performing a number of experiments on the public thyroid ultrasound image dataset TN3K [5] and the private dataset which contains 6 domains.

2 Method

To describe the details of SHAN, as shown in Fig. 2, some symbolic definitions are given firstly. The input is denoted as $X \in \mathbb{R}^{H \times W \times 3}$. The $\hat{Y}$ and $Y \in \mathbb{R}^{H \times W \times C}$ represent the prediction and the Ground Truth(GT), where C is the number of classes. The encoder and decoder are denoted as $f_e(\cdot)$ and $f_d(\cdot)$. The $f_{seg}(\cdot)$ denotes the projection head of segmentation. The loss function of SHAN consists of the two cross-entropy losses between the elliptical prediction E from encoder, the prediction from decoder and the GT respectively, that is $\mathcal{L}_{all} = \mathcal{L}(E, Y) + \mathcal{L}(\hat{Y}, Y)$. Since SHAN is proposed as a plug-and-play method, we did not design complex loss function and the way it is integrated with the loss of baseline.




Fig. 2. An overview of the proposed SHAN during training in which the target domain in the dashed box on the left is kept unseen. The same shape prior is given for data from different domains. Details can be found in the main text.

2.1 Nodule Shape Prior

The SHAN attempts to find stable domain-invariant representation of thyroid nodules. Consider that a nodule is generally an approximate ellipse, which is domain-independent but task-relevant. More specifically, in Thyroid Imaging Reporting and Data System (TI-RADS) [18], the Axial/Transverse (A/T) of the nodule is an important diagnostic factor. Thus, SHAN initializes the shape prior $S \in \mathbb{R}^{H \times W}$ with the critical state $A/T = 1$. S is a fixed binary matrix in which pixels within a radius r from the center are filled with 1 and others are set to 0. The encoder builds the stable mapping between the latent features and the distributions of A/T from benign and malignant nodules based on this Nodule Shape Prior S .

2.2 Global Affine Module

The purpose of Global Affine Module (GAM) is to initialize the elliptic shape of nodules based on the Nodule Shape Prior by predicting the affine transformation matrix of each image under a global field, which constructs a mapping relationship between the latent features and the stable shape distribution of nodules. The features F_x and F_y of the lesion in the sagittal (vertical) and coronal (horizontal) planes in the ultrasound images are obtained by the GAM through the horizontal and vertical pooling of the latent features $U = f_e(X) \in \mathbb{R}^{H' \times W' \times C'}$. The global affine transformation matrix is denoted as $A_g^c \in \mathbb{R}^{2 \times 3}$, where c represents benign or malignant. Each A_g^c contains a total of 5 learnable parameters, including the zoom factor z_x^c controlling the x-axis direction, the zoom factor z_y^c controlling the y-axis direction, the offset factor v_x^c and v_y^c , and the rotation angle θ^c . These $5 \times (C - 1)$ parameters can be obtained from $fc(F_{x,y})$, and $fc(\cdot)$ is the fully connected layer. Inspired by [8], elliptic-shaped prediction E based on the Nodule Shape Prior is obtained by Grid Sample G and Affine Transformation ψ_1 , that is $E = G(S, \psi_1(A_g, P))$. $P = \{(x_i, y_i) | i = [1, \dots, H \times W]\}$ denotes the target coordinates of regular grid in E . The definition of ψ_1 is as follows,

$$\begin{pmatrix} x_i^s \\ y_i^s \end{pmatrix} = \psi_1(A_g, P_i) = \begin{bmatrix} \cos\theta \times z_x & -\sin\theta & v_x \\ \sin\theta & \cos\theta \times z_y & v_y \end{bmatrix} \begin{pmatrix} x_i \\ y_i \\ 1 \end{pmatrix} \quad (1)$$

where (x_i^s, y_i^s) are the coordinates of sample points in the Nodule Shape Prior S . In our method, the bilinear sampling is used as Grid Sample G , this is because the affine transformation changes the density of the grid which is defined as shown in Eq.2.

$$E_i = \sum_{h=1}^H \sum_{w=1}^W S_{hw} \max(0, 1 - |x_i^s - h|) \max(0, 1 - |y_i^s - w|) \quad (2)$$

The E , as the preliminary prediction of the nodule, provides the encoder with information about the position, A/T, growth direction of the nodules, and also provides hints for the segmentation task performed by the decoder.

2.3 Neighbor Affine Module

Unlike the GAM, which focuses on the global information of the nodule, the Neighbor Affine Module (NAM) establishes a mapping relationship between the initial elliptic prediction and the fine segmentation result in the local field. This allows the shape prior to further fit the edges of the nodule, improving the model's ability to extract adjacent local detail features. The neighbor transform matrix $A_n = f_n(D(E), U)$ defines the offset of the elliptic prediction pixel-by-pixel, where $f_n(\cdot)$ is the Neighbor Decoder, $D(\cdot)$ is down-sampling, and $A_n \in R^{(C-1) \times H \times W}$. The neighbor affine transformation is still performed by bilinear sampling G , i.e. $E' = G(E, \psi_2(A_n + P))$, where ψ_2 is a refined deformation function that gives each grid point a non-uniform affine transformation direction. The final segmentation prediction $\hat{Y} = f_{seg}(E', f_d(U))$.

3 Experiments

3.1 Datasets and Settings

Two thyroid ultrasound image datasets are used in the experiment. TN3K[5], a public dataset, includes 3493 images with pixel-labeled nodules. And a private Thyroid Ultrasound Image (TUI) dataset is collected from Tianjin Medical University Cancer Hospital, which includes 16780 images with pixel-labeled benign and malignant nodules. Each image in TUI contains exactly one nodule and has been approved by the Medical Ethics Committee of Tianjin Medical University Cancer Hospital. Written consent had been obtained from each patient after full explanation of the purpose and nature of all procedures used. Depending on the different ultrasound machine, the TUI can be divided into 6 domains, as shown in Fig. 3. The P1 of TUI is treated as source domain to train the models.

For TN3K and TUI, we all take the five-fold cross-validation. All the input images are resized to 224×224 . Our experiments are performed on one




Fig. 3. TUI dataset. (a)-(c) show the complexity, uniformity of gray scale, contrast and texture from different domains in TUI by different properties of Gray-Level Co-occurrence Matrix, respectively. (d)-(e) show the number and image of each domain.

NVIDIA A100 GPU. AdamW optimizer [12] is used with learning rate 10^{-3} for 200 epoches. Batch size is 24. The r of the shape prior is 20. In the following text, the "ours" refers to the SHAN utilizing Unet[15] as baseline.

3.2 Experiment on source domain

To show the effectiveness of our method on source domain for supervised segmentation task, we performed experiments on P1 of TUI and TN3K.

Table 1. The comparison of the segmentation results on source domain(P1).

Method	IOU(%)			Dice(%)		
	Benign	Malignant	mean	Benign	Malignant	mean
TransDL[1]	62.29±0.99	70.70±0.39	66.50±0.47	76.76±0.75	82.83±0.27	79.80±0.36
TransDL*	76.73 ±0.20	79.10 ±0.15	77.91 ±0.03	86.83 ±0.06	88.33 ±0.03	87.58 ±0.02
SegNet[2]	70.23±3.74	74.45±1.64	72.34±2.43	82.45±2.62	85.35±1.07	83.90±1.67
SegNet*	77.38 ±1.67	79.74 ±0.26	78.56 ±0.39	87.24 ±0.66	88.73 ±0.16	87.98 ±0.25
TransUNet[3]	70.73±2.01	74.59±1.67	72.66±1.81	82.86±1.05	85.45±1.88	84.15±1.69
TransUNet*	75.61 ±1.94	78.05 ±0.86	77.48 ±0.66	86.02 ±0.49	88.37 ±0.64	87.29 ±0.44
UNet++[21]	79.79±0.31	80.44±0.26	80.12±0.28	88.76±0.19	89.16±0.16	88.96±0.17
UNet++*	80.46 ±0.42	81.10 ±0.51	80.78 ±0.46	89.77 ±0.26	89.56 ±0.30	89.37 ±0.28
AttUNet[13]	79.11±0.48	81.59 ±0.30	80.35±0.29	88.34±0.22	89.86±0.17	89.10±0.14
AttUNet*	80.88 ±0.29	81.50±0.23	81.19 ±0.18	89.43 ±0.09	89.81 ±0.13	89.62 ±0.10
baseline	78.35±1.06	80.38±0.48	79.36±0.46	87.86±0.67	89.12±0.30	88.49±0.29
SHAN(ours)	81.18 ±0.20	81.95 ±0.13	81.57 ±0.06	89.61 ±0.12	90.08 ±0.08	89.85 ±0.03

* stands for plug-and-play SHAN in that method.

TUI(P1). The 3-class segmentation performance of our method in TUI is shown in Tab. 1. It can be seen that the SHAN improves the segmentation of both benign and malignant nodules in the source domain. The reason is that the affine transformation matrix of the shape prior builds a relationship between the features and the A/T of different nodules, which is an effective representation. The differences between the true nodule shape and the predicted elliptic shape are emphatically learned by the neighbor affine transformation of SHAN.

TN3K. As the GT in TN3K only contains nodule and background, we give more metrics in Tab. 2 to show the performance of different methods for 2-class segmentation. The TRFE and TRFE+ are multi-task models trained on TN3K and combined thyroid-labeled images from TG3K[6]. We incorporated SHAN into the current reported best model TRFE+¹. The improvement our method brings to the original model is clearly observable.

Table 2. The segmentation results on TN3K.

Method	Accuracy(%)	IOU(%)	Dice(%)	HD95
Unet[15]	96.46 $\pm$ 0.11	65.99 $\pm$ 0.66	79.51 $\pm$ 1.31	18.44 $\pm$ 0.75
TRFE[5]	96.71 $\pm$ 0.07	68.33 $\pm$ 0.68	81.91 $\pm$ 1.35	17.96 $\pm$ 1.24
SegNet[2]	96.72 $\pm$ 0.12	66.54 $\pm$ 0.85	79.91 $\pm$ 1.69	17.13 $\pm$ 0.89
DeepLabv3+[4]	97.19 $\pm$ 0.05	70.60 $\pm$ 0.49	82.77 $\pm$ 0.98	13.92 $\pm$ 0.89
TRFE+[6]	97.04 $\pm$ 0.10	71.38 $\pm$ 0.43	83.30 $\pm$ 0.26	13.23 $\pm$ 0.63
TRFE+ w. SHAN	96.73 $\pm$ 0.06	73.59 $\pm$ 0.16	84.61 $\pm$ 0.04	4.05 $\pm$ 0.27

3.3 Experiment on target domain

To prove the generalization ability of our method, we treat P1 as the source domain and test model on the data from the other domains. The prediction results of the GAM and NAM on the target domain are shown in Fig.4. From the preliminary prediction E of the nodule, it can be seen that the SHAN could learn the stable nodule shape features in terms of location, size, and A/T.

The generalization abilities of different segmentation methods and their improvement with SHAN, the current SDG methods, and our method are shown in Tab. 3, respectively. The "Intra-domain" means training and testing the model on the specific target domain. We do not perform intra-domain training for P3 and T1 as the number of data in these domains is too small. First, the 3-12 rows in Tab. 3 show that our method improves the generalization ability of current segmentation models with encoding-decoding structure. The SHAN has a noticeable improvement on TransDL, up to 21.71% improvement in mIOU on P2. Second, comparison with other SDG methods shows that our method still reaches the optimum, as shown in 13-16 rows. This is because the data augmentation methods used by current SDG methods cannot fit the domain gap between

¹ <https://github.com/haifangong/TRFE-Net-for-thyroid-nodule-segmentation>




Fig. 4. Some examples of the input images, elliptic-shaped prediction, final prediction and ground truth. Left: Benign nodule images. Right: Malignant nodule images.

ultrasound images from multi-center. In contrast, our SHAN does not rely on data augmentation, but rather improves the generalization ability of the model by learning the nodule shape representation to obtain a domain-invariant latent features. More segmentation results can be seen in the supplementary material.

Table 3. The mIOU(%) comparison of the SHAN, segmentation methods and SDG methods on target domains.

Method	P2	P3	M3	T1	T2
Intra-domain	70.67	—	64.72	—	72.93
TransDL	45.66 $\pm$ 1.41	56.50 $\pm$ 2.19	63.38 $\pm$ 0.74	59.04 $\pm$ 0.78	45.75 $\pm$ 1.73
TransDL*	67.37 $\pm$ 0.39	62.04 $\pm$ 0.98	66.89 $\pm$ 0.08	69.25 $\pm$ 0.29	53.97 $\pm$ 1.23
TransUNet	56.20 $\pm$ 1.75	57.43 $\pm$ 3.18	63.92 $\pm$ 2.66	63.41 $\pm$ 1.14	51.47 $\pm$ 1.92
TransUNet*	66.93 $\pm$ 1.20	63.18 $\pm$ 1.95	66.47 $\pm$ 0.93	68.99 $\pm$ 0.32	54.02 $\pm$ 0.87
SegNet	60.87 $\pm$ 2.85	62.05 $\pm$ 1.86	65.54 $\pm$ 1.68	64.51 $\pm$ 0.69	53.97 $\pm$ 1.21
SegNet*	71.44 $\pm$ 0.36	70.40 $\pm$ 1.92	69.26 $\pm$ 0.74	70.10 $\pm$ 0.16	59.25 $\pm$ 0.18
UNet++	72.99 $\pm$ 0.65	68.56 $\pm$ 0.21	68.59 $\pm$ 2.04	73.01 $\pm$ 2.52	61.33 $\pm$ 0.07
UNet++*	74.16 $\pm$ 0.72	70.25 $\pm$ 0.40	69.74 $\pm$ 1.63	74.11 $\pm$ 1.67	62.08 $\pm$ 0.12
AttUNet	73.01 $\pm$ 0.72	66.66 $\pm$ 0.35	65.22 $\pm$ 1.56	73.29 $\pm$ 0.96	60.71 $\pm$ 1.00
AttUNet*	75.80 $\pm$ 0.33	70.10 $\pm$ 0.91	66.23 $\pm$ 1.24	75.98 $\pm$ 0.53	62.15 $\pm$ 0.07
FreeSDG‡ [9]	54.86 $\pm$ 0.56	55.12 $\pm$ 1.80	61.55 $\pm$ 1.37	57.24 $\pm$ 0.71	48.26 $\pm$ 1.06
SLAug‡ [16]	53.59 $\pm$ 0.29	55.53 $\pm$ 1.94	64.17 $\pm$ 1.15	56.36 $\pm$ 0.55	53.80 $\pm$ 0.89
C ² SDG‡ [7]	55.06 $\pm$ 0.41	56.71 $\pm$ 2.23	67.08 $\pm$ 0.57	62.11 $\pm$ 0.59	54.44 $\pm$ 0.28
GIN‡ [14]	71.40 $\pm$ 0.16	68.15 $\pm$ 1.70	66.16 $\pm$ 0.40	72.75 $\pm$ 0.93	60.38 $\pm$ 0.93
SHAN(ours)	75.19 $\pm$ 0.31	72.05 $\pm$ 0.29	70.91 $\pm$ 0.35	75.13 $\pm$ 0.42	64.98 $\pm$ 0.97

* stands for plug-and-play SHAN in that method. ‡stands for it's a SDG method.

4 Discussion

In this paper, the SHAN is proposed for the cross-domain segmentation of thyroid ultrasound images, which learns the shape prior distributions of thyroid nodules by learnable affine matrix, to improve the performance of the model

in source and target domains. The workflow of the SHAN is similar to that of human segmenting an object, which refines the segmentation result based on the coarse-grained preset. We prove the effectiveness of the SHAN by extensive experiments on TN3K and a private dataset with 6 domains. Exploring the scalability of this method on other diseases is our future work.

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

References

1. Azad, R., Heidari, M., Shariatnia, M., Aghdam, E.K., Karimijafarbigloo, S., Adeli, E., Merhof, D.: Transdeeplab: Convolution-free transformer-based deeplab v3+ for medical image segmentation. In: Predictive Intelligence in Medicine - 5th International Workshop Held in Conjunction with MICCAI. Lecture Notes in Computer Science, vol. 13564, pp. 91–102. Springer (2022)
2. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(12), 2481–2495 (2017)
3. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306 (2021), <https://arxiv.org/abs/2102.04306>
4. Chen, L., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: European Conference on Computer Vision, Proceedings, Part VII. Lecture Notes in Computer Science, vol. 11211, pp. 833–851. Springer (2018)
5. Gong, H., Chen, G., Wang, R., Xie, X., Mao, M., Yu, Y., Chen, F., Li, G.: Multi-task learning for thyroid nodule segmentation with thyroid region prior. In: 18th IEEE International Symposium on Biomedical Imaging, ISBI. pp. 257–261. IEEE (2021)
6. Gong, H., Chen, J., Chen, G., Li, H., Li, G., Chen, F.: Thyroid region prior guided attention for ultrasound segmentation of thyroid nodules. *Comput. Biol. Medicine* **155**, 106389 (2023)
7. Hu, S., Liao, Z., Xia, Y.: Devil is in channels: Contrastive single domain generalization for medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention - MICCAI Proceedings, Part IV. Lecture Notes in Computer Science, vol. 14223, pp. 14–23. Springer (2023)
8. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: Conference on Neural Information Processing Systems, NeurIPS. pp. 2017–2025 (2015)
9. Li, H., Li, H., Zhao, W., Fu, H., Su, X., Hu, Y., Liu, J.: Frequency-mixed single-source domain generalization for medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention - MICCAI Proceedings, Part VI. Lecture Notes in Computer Science, vol. 14225, pp. 127–136. Springer (2023)
10. Li, L., Gao, K., Cao, J., Huang, Z., Weng, Y., Mi, X., Yu, Z., Li, X., Xia, B.: Progressive domain expansion network for single domain generalization. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. pp. 224–233. IEEE (2021)

11. Liu, Q., Chen, C., Dou, Q., Heng, P.: Single-domain generalization in medical image segmentation via test-time adaptation from shape dictionary. In: Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI. pp. 1756–1764. AAAI Press (2022)
12. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: 7th International Conference on Learning Representations, ICLR. OpenReview.net (2019)
13. Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M.C.H., Heinrich, M.P., Misawa, K., Mori, K., McDonagh, S.G., Hammerla, N.Y., Kainz, B., Glocker, B., Rueckert, D.: Attention u-net: Learning where to look for the pancreas. arXiv preprint arXiv:1804.03999 (2018), <http://arxiv.org/abs/1804.03999>
14. Ouyang, C., Chen, C., Li, S., Li, Z., Qin, C., Bai, W., Rueckert, D.: Causality-inspired single-source domain generalization for medical image segmentation. *IEEE Trans. Medical Imaging* **42**(4), 1095–1106 (2023)
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI Proceedings, Part III. Lecture Notes in Computer Science, vol. 9351, pp. 234–241. Springer (2015)
16. Su, Z., Yao, K., Yang, X., Huang, K., Wang, Q., Sun, J.: Rethinking data augmentation for single-source domain generalization in medical image segmentation. In: Thirty-Seventh Conference on Artificial Intelligence, AAAI. pp. 2366–2374. AAAI Press (2023)
17. Tang, L., Tian, C., Yang, H., Cui, Z., Hui, Y., Xu, K., Shen, D.: TS-DSANN: texture and shape focused dual-stream attention neural network for benign-malignant diagnosis of thyroid nodules in ultrasound images. *Medical Image Anal.* **89**, 102905 (2023)
18. Tessler, F.N., Middleton, W.D., Grant, E.G., Hoang, J.K., Berland, L.L., Teefey, S.A., Cronan, J.J., Beland, M.D., Desser, T.S., Frates, M.C., et al.: Acr thyroid imaging, reporting and data system (ti-rads): white paper of the acr ti-rads committee. *Journal of the American college of radiology* **14**(5), 587–595 (2017)
19. Zhang, L., Wang, X., Yang, D., Sanford, T., Harmon, S.A., Turkbey, B., Wood, B.J., Roth, H., Myronenko, A., Xu, D., Xu, Z.: Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation. *IEEE Trans. Medical Imaging* **39**(7), 2531–2540 (2020)
20. Zhou, K., Liu, Z., Qiao, Y., Xiang, T., Loy, C.C.: Domain generalization: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(4), 4396–4415 (2023)
21. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support Workshop Held in Conjunction with MICCAI. pp. 3–11. Springer (2018)