This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Noise-Robust Tuning of SAM for Domain Generalized Ultrasound Image Segmentation

Zhikai Wei[1], Chao Wu[1], Hanyu Du[1], Rui Yu[2], Bo Du[1]✉, and Yongchao Xu[1]✉

[1] National Engineering Research Center for Multimedia Software, Institute of Artificial Intelligence, School of Computer Science, Medical Artificial Intelligence Research Institute of Renmin Hospital, Wuhan University, Wuhan, China
{dubo,yongchao.xu}@whu.edu.cn
[2] University of Louisville, KY, USA

Abstract. The Segment Anything Model (SAM) has achieved outstanding performance in both natural and medical image segmentation with extensive research validation. When applied to ultrasound images, which involve low contrast, indistinct boundaries and complex shapes, large models still suffer from significant performance degradation and limited generalization ability. We explore these challenges from a new perspective with the help of the segmentation foundation model SAM. In this paper, we propose Nora, a noise-robust fine-tuning framework for SAM to address domain generalized ultrasound image segmentation. Specifically, we introduce a feature-adaptive perturbation module, which applies well-designed noise to the fine-tuned features. We stimulate the model to segment the correct regions even under severe interference, thereby improving its robustness. Moreover, to further optimize SAM with prompts, we present an instance-aware prompt generation module. We introduce a set of tokens linked to distinct instances and then design a token-based augmentation strategy to prevent overcoupling and encourage tokens to capture more diverse information. Our Nora achieves state-of-the-art performance across extensive cross-domain experiments with three ultrasound image segmentation tasks, fully demonstrating its effectiveness and generalizability. The code is available at https://github.com/wkklavis/Nora.

Keywords: Ultrasound image segmentation · Domain generalization · Noise injection · SAM · Model robustness

1 Introduction

Ultrasound imaging, recognized for its non-invasive, safe, and widely accessible nature, plays a critical role in medical diagnosis and therapeutic interventions [26]. Numerous deep learning models have been proposed for ultrasound image segmentation, demonstrating significant potential. However, their performance often degrades when applied to other centers or hospitals, resulting in inconvenience for clinical use. Domain Generalized Ultrasound Image Segmentation (DGUIS) aims to train the model on a single source domain, enabling robust segmentation predictions on unseen target domains for ultrasound images.

Ultrasound images are intrinsically limited in quality, often exhibiting low resolution, poor contrast, and a low signal-to-noise ratio (SNR), which hinder the visibility of critical information [19]. Traditional generalization methods, primarily focused on style transformation, perform less effectively than expected on ultrasound images. The emerging vision foundation model [21,10], with strong generalization ability inherited from a large quantity pre-trained images, provides a new possible paradigm for domain generalized segmentation. Recently, many methods have adapted the Segment Anything Model (SAM) [21] to medical downstream tasks, achieving significant performances. However, SAM still struggles with low-quality ultrasound images, and there is a lack of generalized segmentation schemes specifically designed for them.

Inspired by adversarial attack [9], we start with noise robustness to approach the task of domain generalized ultrasound image segmentation. We draw an analogy between ultrasound generalization and adversarial defense. To be specific, we treat noisy ultrasound images as if they were already subjected to adversarial attacks. We tackle the problem by adopting noise injection [15] from adversarial defense strategies to enhance the model's robustness. In this paper, we present a novel **No**ise **R**obust fine-tuning S**A**M framework (namely "**Nora**") to address domain generalized ultrasound image segmentation. Its conceptual idea is to leverage the excellent segmentation capability of the large model to identify robust features under severe noise interference, thereby achieving generalization.

First, we propose a feature-adaptive perturbation module. When fine-tuning the SAM's encoder, we strategically inject well-designed noise into the intermediate features. We expect the model to not only segment low-quality ultrasound images successfully but also extract the correct regions of interest from features with higher levels of noisy corruption. With the strong segmentation potential of the foundation model, we fine-tune the disrupted features, facilitating the model to optimize and identify robust information. Thus, when faced with unseen ultrasound domains, our model still exhibits strong generalization performance.

Moreover, we design a module to generate instance-aware prompts automatically, aiming to fully exploit the benefits of prompts for SAM. We introduce a set of learnable tokens that interact with image embeddings. These tokens, combined with cross-attention features, serve as sparse and dense prompts for subsequent network respectively. We aim for the tokens to interact with distinct instances, facilitating instance-level feature refinement [31]. Additionally, to better adapt to ultrasound image tasks, we design a token-based perturbation algorithm to prevent the learnable parameters from overcoupling to specific instances. Building on token-based feature refinement, we model inter-token dependencies and apply noise perturbation accordingly to mitigate overfitting, thereby further enhancing the model's generalization in ultrasound image segmentation.

In a nutshell, the main contributions of this paper are as follows:

1. We propose a novel domain generalization framework by fine-tuning the vision foundation model for ultrasound image segmentation. We apply a noise-based feature perturbation mechanism to stimulate SAM to learn more robust features.

2. We present a noise-robust prompt module, to efficiently and automatically harness SAM. We introduce a set of tokens to link distinct instances and generate fine-grained prompts, while incorporating a novel token perturbation algorithm to mitigate overfitting.
3. Extensive experiments across various DGUIS settings demonstrate the generalization of our fine-tuning scheme. Our Nora surpasses recent SAM-based and fine-tuning methods, achieving state-of-the-art results on cross-domain ultrasonic datasets covering three distinct organs.

2 Related works

Domain Generalized Medical Image Segmentation. Numerous approaches [23,4,16,8,25] have been proposed to address the model performance degradation due to distribution shift across varied application scenarios. The emergence of foundation vision models has led to numerous works to adapt them to medical tasks [28,24,14,5], achieving remarkable performance. Meanwhile, SAM, with its outstanding general segmentation capability, has introduced new perspectives for domain generalized segmentation. DAPSAM [30] proposes a domain-adaptive prompt framework to store and utilize source domain knowledge. Rein [31] introduces a robust fine-tuning approach to parameter-efficiently harness vision foundation model. In this paper, we consider the characteristics of ultrasound images and primarily focus on domain generalization, which is significant but lacks extensive research and exploration.

Noise Robust in Adversarial Attacks. Deep neural networks have achieved great success in a variety of applications. Despite the remarkable improvement, prior studies [7,3] have shown that networks are vulnerable to adversarial examples, which are intentionally perturbed inputs designed to cause erroneous prediction [13]. Such a fragility undermines the model's reliability and limits its deployment in critical fields like healthcare and autonomous driving. To address these vulnerabilities, researchers have proposed various defense methods. Among these, randomization techniques, particularly noise injection [15,18,32], have shown strong potential in enhancing adversarial robustness by introducing controlled uncertainty, helping the model counter adversarial inputs more effectively. In this paper, we treat ultrasonic images as attacked inputs and adopt adversarial defense strategies to tackle the problem. We address domain generalized ultrasound image segmentation from a novel perspective.

3 Method

For DGUIS, the objective is to learn a segmentation model for ultrasound images using only a source domain $\mathcal{S}$, and the trained model is expected to show good generalizable performance on any other unseen target domains $\mathcal{T} = \{\mathcal{T}_1, ..., \mathcal{T}_n\}$. We explore DGUIS based on fine-tuning SAM.

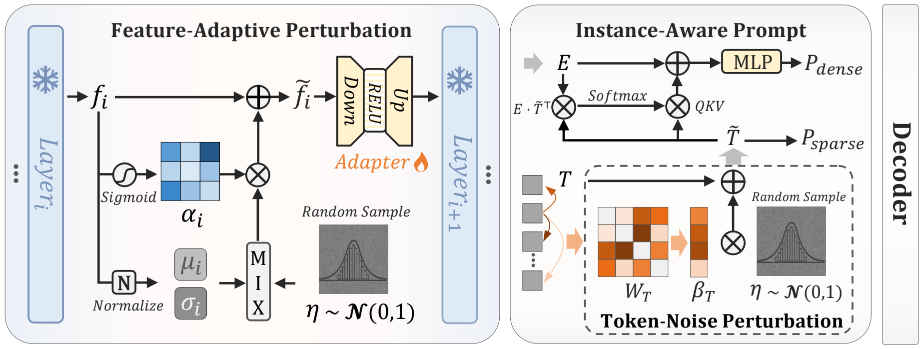

Fig. 1. The pipeline of the proposed Nora, a fine-tuning framework for DGUIS. We propose a feature-adaptive perturbation module (left) within the feature space of the adapter in each frozen block to encourage SAM to learn more robust features. To efficiently and automatically leverage SAM, we present a noise-robust prompt generation module (right), consisting of an instance-aware prompt and token-noise perturbation. The image embedding and the prompts are then fed into the decoder.

3.1 Feature-Adaptive Perturbation

For simple implementation, we do not redesign the fine-tuning structure as described in the [24,31]. Instead, we apply AdaptFormer [6] to fine-tune SAM thanks to its efficiency and scalability, and further extend our work based on it.

We propose a feature-adaptive perturbation module within the feature space of the adapter. Different from previous style transformation methods based on AdaIN [17], we directly apply well-designed noise perturbations at the original features f_i produced by the i-th layer. Due to the uncontrollable perturbation impact of the vanilla Gaussian noise, we first align the distribution between the Gaussian noise η and feature f_i using the mean μ_i and the variance σ_i computed from f_i along the spatial dimension. Moreover, to better design the noise, we apply a sigmoid function to the features to generate the weight, which controls the magnitude of the perturbation. The entire feature-noise perturbation fine-tuning module can be formulated as:

$$\tilde{f}_i = Adapter\big(f_i + \alpha_i \cdot (\eta \cdot \sigma_i + \mu_i)\big), \quad \alpha_i = sigmoid(f_i), \quad \eta \sim \mathcal{N}(0,1), \quad (1)$$

where $Adapter(\cdot)$ denotes the lightweight tuning module from AdaptFormer [6], α_i is our noise perturbation weight, computed using the sigmoid function. We generate the weight based on the relative magnitude of the current features—high values allow stronger perturbations, while low-value features are less disturbed. This adaptive weight enables us to apply noise perturbations within a controlled range, which helps preserve feature integrity while encouraging the model to identify more robust features.

3.2 Noise-Robust Prompt

Building upon [31,2] but differing in design and structure, we introduce a set of learnable tokens $T \in \mathbb{R}^{N\times C}$ for prompting and a novel token-augmentation algorithm. These tokens interact with instances, refine the image embedding, and generate sufficient prompts to guide the subsequent decoder for segmentation. Specifically, for each image embedding $E \in \mathbb{R}^{H\times W\times C}$, our tokens act as queries through cross-attention to compute refined features, which serve as dense prompts $P_{dense} \in \mathbb{R}^{H\times W\times C}$. At the same time, the tokens also function as sparse prompts $P_{sparse} \in \mathbb{R}^{N\times C}$, providing supplementary information:

$$P_{dense} = Linear\Big(Softmax(\frac{E\cdot \tilde{T}^{\intercal}}{\sqrt{C}})\cdot \tilde{T} + E\Big), \quad P_{sparse} = \tilde{T} = TNP(T), \tag{2}$$

where $Linear(\cdot)$ is a simple linear layer used to project the feature space. TNP represents our proposed token noise perturbation module.

For each instance, we aim for tokens to learn more fine-grained features. FADA [2] employs instance normalization to mitigate tokens' preference for domain-specific styles. When handling low-quality ultrasound images, we adopt a more effective token-instance decoupling strategy to prevent overfitting. When a token exhibits excessively high similarity with other tokens, we consider its knowledge replaceable. Inspired by [27], but instead of pruning it, we design a token-enhancement algorithm tailored for ultrasonic features. We consider the correlations within tokens and introduce noise perturbation in a targeted manner to stimulate tokens T to capture richer instance-level information. Specifically, we calculate the similarity between each token and the other tokens to determine the degree of perturbation $\beta_T \in \mathbb{R}^N$:

$$TNP(T) = T + \beta_T \cdot \eta\,, \quad \beta_T = \frac{\sum_{i=1}^{N}(T\cdot T^{\intercal} - I)_{i,:}}{N\times C}, \quad \eta \sim \mathcal{N}(0,1), \tag{3}$$

where $\eta \in \mathbb{R}^{N\times C}$ is random noise from a standard Gaussian distribution and I denotes the identity matrix.

The perturbation of tokens not only mitigates the overfitting of learnable tokens to source domain features, but also encourages tokens to learn instance-level robust features throughout the training process.

4 Experiments

4.1 Experimental Settings

Implementation Details. We employ AdaptFormer [6] to fine-tune the SAM as our baseline model. We adopt the same experimental setting as DAPSAM [30], including the warm-up strategy, the rank of adapter and the loss design. The initial learning rate is set to $5e^{-4}$, and the weight decay for the AdamW optimizer is set to 0.1. The hyperparameter N, representing the number of tokens, is set

Table 1. Quantitative comparison of our Nora and some state-of-the-art domain generalization methods on BUS datasets. The best and second-best are **bolded** and underlined, respectively.

Method	Type	*BUSI → DatasetB*				*BUSI → STU*			
		Dice↑	mIoU↑	ASD↓	HD↓	Dice↑	mIoU↑	ASD↓	HD↓
Baseline [6]	-	80.37	72.39	4.82	21.35	88.15	79.80	1.00	16.99
DSU [23]	Style Transfer	81.33	73.35	<u>4.23</u>	19.33	88.63	80.45	0.63	18.18
TriD [8]		81.65	<u>73.90</u>	4.30	<u>18.39</u>	89.21	81.41	0.91	17.35
Rein [31]	Fine Tuning	80.51	71.89	6.11	23.83	88.13	79.78	0.72	16.85
FADA [2]		81.36	73.55	4.67	21.48	88.67	80.39	0.59	16.27
DAPSAM [30]	SAM Based	81.67	73.42	5.51	22.82	<u>89.37</u>	<u>81.85</u>	<u>0.41</u>	**15.16**
SAMUS [24]		<u>82.23</u>	72.61	6.00	21.77	88.61	80.14	4.08	27.33
Nora (Ours)		**83.88**	**75.81**	**3.60**	**16.77**	**90.41**	**83.16**	**0.22**	<u>15.39</u>

to 50 for a suitable trade-off [31]. To quantitatively evaluate the segmentation performance, we adopt the following four commonly used metrics: dice coefficient (Dice), mean intersection over union (mIoU), average surface distance (ASD), and hausdorff distance (HD).

Datasets. We conduct three different domain generalized ultrasound image segmentation experiments, including breast cancer (BUS), thyroid nodule (Thyroid), and myocardium (MYO) ultrasound datasets. For BUS, we use BUSI [1], DatasetB [33] and STU [34]. BUSI is randomly split into 7:1:2 for training, validation, and testing, respectively. For Thyroid, we leverage TN3K [11] and DDTI [29]. The data in TN3K is partitioned into train, and test sets following TRFE [12]. For MYO, we adopt CAMUS [22] and HMC-QU [20]. We extract only the first frame from each video of CAMUS to ensure a dataset size similar to that of HMC-QU.

4.2 Comparison with SOTA Methods

Results on DGUIS are presented in Table 1 and Table 2. We compare our Nora with some recent state-of-the-art generalization approaches on DGUIS task. For BUS, our Nora surpasses the best-performing SAM-based methods, DAPSAM [30] and SAMUS [24] with an increase of 2.21% and 1.65% in Dice on the $S_{BUSI \to DatasetB}$, and also outperforms the latest generalization fine-tuning frameworks, Rein [31] and FADA [2]. Additionally, our method also surpasses TriD [8], an efficient and superior approach based on style augmentation, by approximately 2%. For Thyroid and MYO, Nora also achieves state-of-the-art results. Notably, our method maintains strong performance and exhibits stable segmentation capability across different ultrasound tasks for various diseases. In contrast, SAMUS, which fine-tunes a larger number of parameters (∼38M vs. ours ∼7M), suffers a substantial performance drop in the MYO task, where cross-domain style variations are more pronounced.

Table 2. Quantitative comparison of our Nora and some state-of-the-art domain generalization methods on Thyroid (left half) and MYO (right half) datasets. The best and second-best are **bolded** and underlined, respectively.

Method	*TN3K → DDTI*				*CAMUS3K → HMC-QU*			
	Dice↑	mIoU↑	ASD↓	HD↓	Dice↑	mIoU↑	ASD↓	HD↓
Baseline [6]	74.97	64.37	4.32	43.90	75.70	61.11	1.01	15.85
DSU [23]	76.05	65.26	4.33	45.26	76.58	62.30	1.08	16.26
TriD [8]	77.18	65.68	4.34	45.35	77.21	63.09	0.92	15.86
Rein [31]	75.87	64.69	5.00	45.90	77.37	63.27	**0.75**	15.28
FADA [2]	76.31	65.33	4.32	43.61	76.74	62.48	0.81	15.31
DAPSAM [30]	77.26	66.29	**3.71**	**43.12**	76.98	62.79	0.81	**13.71**
SAMUS [24]	76.32	64.53	19.33	62.98	47.41	31.41	9.61	44.07
Nora (Ours)	**77.73**	**66.54**	4.07	43.71	**77.74**	**63.77**	0.78	14.42

Qualitative Analysis. Fig. 2 showcases the segmentation results from various methods. Ultrasound image segmentation is particularly challenging due to low contrast, inconsistent textures, and indistinct object boundaries. Nevertheless, our method demonstrates robust performance when confronting these challenges, achieving superior performance.

Results on noise injection are presented in Fig. 3. To better demonstrate the effectiveness of our method, we also compare our feature-adaptive perturbation module with recent or classic noise injection methods derived from adversarial defense. Our method not only considers distribution alignment but also meticulously adjusts the noise to perturb intermediate features, fully leveraging SAM's potential segmentation capability. It surpasses the parameterized PNI [15], the learnable L2P [18], and also the latest NINE [32], which relaxes the Gaussian distribution to a non-informative prior arbitrary distribution.

4.3 Ablation Study

Table 3 is dedicated to examining the effectiveness of each component in Nora on the BUS datasets. Within the feature-adaptive perturbation (FAP), we inject elaborated noise into the ultrasound feature space to promote model's disturbance-resistant ability, improving generalizable performance in unseen domains and achieving a 1.24% gain in Dice on the $S_{BUSI \to DatasetB}$ over the baseline. Furthermore, we plumb the beneficial effect of prompts on SAM's performance, designing token-based prompts to capture fine-grained information and facilitate instance-level feature refinement. To curb overfitting between tokens and instances, token-noise perturbation (TNP) considers inter-token correlations to generate noise perturbations, stimulating richer feature learning and further improving the prompt module by 0.81% and 0.89% in Dice on target domains, respectively. Our Nora starts with noise to explore SAM's latent segmentation capacity systematically, enhancing the model's robustness and generalization.

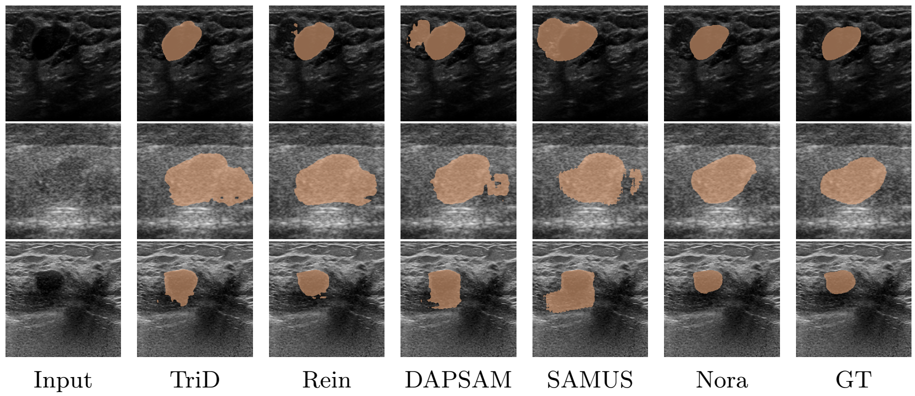

Fig. 2. Qualitative comparison between our Nora and some state-of-the-art domain generalization methods on various ultrasound images.

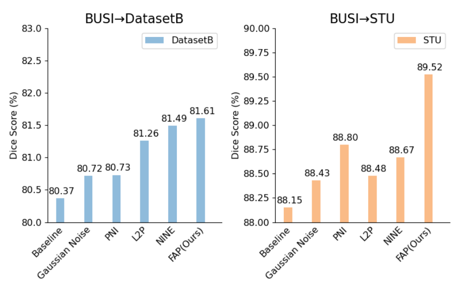

Fig. 3. Comparison results of our FAP module and related noise injection methods on BUS datasets. All noise injection methods are based on baseline.

Table 3. Ablation study on different components of our Nora. We use the BUS datasets, with BUSI as the source domain for training. The first row corresponds to the baseline AdaptFormer [6].

Components			DatasetB		STU	
FAP	Prompt	TNP	Dice↑	HD↓	Dice↑	HD↓
×	×	×	80.37	21.35	88.15	16.99
✓	×	×	81.61	21.42	89.52	**14.09**
×	✓	×	81.35	21.08	88.50	16.77
×	✓	✓	82.16	19.35	89.39	17.81
✓	✓	✓	**83.88**	**16.77**	**90.41**	15.39

5 Conclusion

In this paper, we propose a fine-tuning framework, Nora, which leverages the powerful segmentation model SAM to address domain generalization ultrasound image segmentation. Ultrasound plays an increasingly important role in medical diagnosis. Conventional generalization approaches mainly focus on style variation and regularization. However, their performance often falls short when applied to low-quality ultrasound images in different domains. Driven by this, we propose an efficient fine-tuning method specifically designed for DGUIS. Nora primarily focuses on noise-driven optimization schemes, strategically applying perturbations to the fine-tuned features and instance-aware prompts. We enhance the model's disturbance-resistant capability and then improve its cross-domain generalization, achieving superior performance compared to SOTA methods. Extensive experiments across various settings validate the effectiveness of proposed Nora in efficiently fine-tuning SAM for DGUIS.

Acknowledgments. This work was supported partially by the National Key Research and Development Program of China (2023YFC2705700), NSFC 62222112, 62176186, and 62225113, the NSF of Hubei Province of China (2024AFB245), the Special Fund for Central Guidance on Local Science and Technology Development from the Sichuan Provincial Department of Science and Technology (2024ZYD0285), the Key Research and Development Program of Dazhou Science and Technology Bureau (24ZDYF0005). Rui Yu received no funding in support of this work.

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

References

1. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. Data in Brief **28**, 104863 (2020)
2. Bi, Q., Yi, J., Zheng, H., Zhan, H., Huang, Y., Ji, W., Li, Y., Zheng, Y.: Learning frequency-adapted vision foundation model for domain generalized semantic segmentation. NeurIPS **37**, 94047–94072 (2025)
3. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: Security and Privacy (SP). pp. 39–57 (2017)
4. Chen, C., Li, Z., Ouyang, C., Sinclair, M., Bai, W., Rueckert, D.: Maxstyle: Adversarial style composition for robust medical image segmentation. In: MICCAI. pp. 151–161 (2022)
5. Chen, C., Miao, J., Wu, D., Zhong, A., Yan, Z., Kim, S., , et al: Ma-sam: Modality-agnostic sam adaptation for 3d medical image segmentation. Medical Image Analysis **98**, 103310 (2024)
6. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: Adaptformer: Adapting vision transformers for scalable visual recognition. NeurIPS **35**, 16664–16678 (2022)
7. Chen, Y., Dai, Z., Yu, H., Low, B.K.H., Ho, T.H.: Recursive reasoning-based training-time adversarial machine learning. Artificial Intelligence **315**, 103837 (2023)
8. Chen, Z., Pan, Y., Ye, Y., Cui, H., Xia, Y.: Treasure in distribution: A domain randomization based multi-source domain generalization for 2D medical image segmentation. In: MICCAI. vol. 14223, pp. 89–99 (2023)
9. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: ICML. pp. 2206–2216 (2020)
10. Fang, Y., Wang, W., Xie, B., Sun, Q., Wu, L., Wang, X., Huang, T., Wang, X., Cao, Y.: Eva: Exploring the limits of masked visual representation learning at scale. In: CVPR. pp. 19358–19369 (2023)
11. 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)
12. 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)

13. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: ICLR (2015)
14. Gowda, S.N., Clifton, D.A.: Cc-sam: Sam with cross-feature attention and context for ultrasound image segmentation. In: ECCV. pp. 108–124 (2024)
15. He, Z., Rakin, A.S., Fan, D.: Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack. In: CVPR. pp. 588–597 (2019)
16. Hu, S., Liao, Z., Xia, Y.: Devil is in channels: Contrastive single domain generalization for medical image segmentation. In: MICCAI. vol. 14223, pp. 14–23 (2023)
17. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV. pp. 1501–1510 (2017)
18. Jeddi, A., Shafiee, M.J., Karg, M., Scharfenberger, C., Wong, A.: Learn2perturb: an end-to-end feature perturbation learning to improve adversarial robustness. In: CVPR. pp. 1241–1250 (2020)
19. Jiao, J., Zhou, J., Li, X., Xia, M., Huang, Y., Huang, L., Wang, N., Zhang, X., Zhou, S., Wang, Y., et al.: Usfm: A universal ultrasound foundation model generalized to tasks and organs towards label efficient image analysis. Medical Image Analysis **96**, 103202 (2024)
20. Kiranyaz, S., Degerli, A., Hamid, T., Mazhar, R., et al.: Left ventricular wall motion estimation by active polynomials for acute myocardial infarction detection. IEEE Access **8**, 210301–210317 (2020)
21. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: ICCV. pp. 3992–4003 (2023)
22. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., et al: Deep learning for segmentation using an open large-scale dataset in 2D echocardiography. IEEE Transactions on Medical Imaging **38**(9), 2198–2210 (2019)
23. Li, X., Dai, Y., Ge, Y., Liu, J., Shan, Y., Duan, L.: Uncertainty modeling for out-of-distribution generalization. In: ICLR (2022)
24. Lin, X., Xiang, Y., Yu, L., Yan, Z.: Beyond adapting sam: Towards end-to-end ultrasound image segmentation via auto prompting. In: MICCAI. pp. 24–34 (2024)
25. Liu, F., Ye, M., Du, B.: Learning a generalizable re-identification model from unlabelled data with domain-agnostic expert. Visual Intelligence **2**(1), 28 (2024)
26. Liu, S., Wang, Y., Yang, X., Lei, B., Liu, L., Li, S.X., Ni, D., Wang, T.: Deep learning in medical ultrasound analysis: a review. Engineering **5**(2), 261–275 (2019)
27. Luo, Y., Liu, P., Yang, Y.: Kill two birds with one stone: Domain generalization for semantic segmentation via network pruning. IJCV pp. 1–18 (2024)
28. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. Nature Communications **15**, 1–9 (2024)
29. Pedraza, L., Vargas, C., Narváez, F., Durán, O., Muñoz, E., Romero, E.: An open access thyroid ultrasound image database. In: 10th International Symposium on Medical Information Processing and Analysis. vol. 9287, pp. 188–193 (2015)
30. Wei, Z., Dong, W., Zhou, P., Gu, Y., Zhao, Z., Xu, Y.: Prompting segment anything model with domain-adaptive prototype for generalizable medical image segmentation. In: MICCAI. pp. 533–543 (2024)
31. Wei, Z., Chen, L., Jin, Y., Ma, X., Liu, T., Ling, P., Wang, B., Chen, H., Zheng, J.: Stronger fewer & superior: Harnessing vision foundation models for domain generalized semantic segmentation. In: CVPR. pp. 28619–28630 (2024)
32. Yang, H., Wang, M., Wang, Q., Yu, Z., Jin, G., Zhou, C., Zhou, Y.: Non-informative noise-enhanced stochastic neural networks for improving adversarial robustness. Information Fusion **108**, 102397 (2024)

33. Yap, M.H., Pons, G., Marti, J., Ganau, S., Sentis, M., et al: Automated breast ultrasound lesions detection using convolutional neural networks. IEEE Journal of Biomedical and Health Informatics **22**(4), 1218–1226 (2017)
34. Zhuang, Z., Li, N., Joseph Raj, A.N., Mahesh, V.G., Qiu, S.: An rdau-net model for lesion segmentation in breast ultrasound images. PloS one **14**(8), e0221535 (2019)