

S³-Mamba: Small-Size-Sensitive Mamba for Lesion Segmentation

Anonymous submission

Abstract

Small lesions play a critical role in early disease diagnosis and intervention of severe infections. Popular models often face challenges in segmenting small lesions, as it occupies only a minor portion of an image, while down-sampling operations may inevitably lose focus on local features of small lesions. To tackle the challenges, we propose a **Small-Size-Sensitive Mamba (S³-Mamba)**, which promotes the sensitivity to small lesions across three dimensions: channel, spatial, and training strategy. Specifically, an Enhanced Visual State Space block is designed to focus on small lesions through multiple residual connections to preserve local features, and selectively amplify important details while suppressing irrelevant ones through channel-wise attention. A Tensor-based Cross-feature Multi-scale Attention is designed to integrate input image features and intermediate-layer features with edge features, and exploit the attentive support of features across multiple scales, thereby retaining spatial details of small lesions at multiple granularities. Finally, we introduce a novel regularized curriculum learning to automatically assess lesion size and sample difficulty, and gradually focus from easy samples to hard ones like small lesions. Extensive experiments on three medical image segmentation datasets show the superiority of our S³-Mamba, especially on segmenting small lesions.

Introduction

As the earliest indicators of severe diseases, small lesions play a pivotal role in clinical outcomes, where timely and precise identification can dramatically improve patient prognosis (Qureshi et al. 2023). These lesions, despite their small size, often signal the onset of significant health issues, making their segmentation essential for effective treatment planning (de Teresa-Trueba et al. 2023; Ma and Wang 2023). Many images from different modalities in medical imaging (Greenwald et al. 2022) contain numerous lesions that occupy less than 10% of the image area (Duering et al. 2023; Ali et al. 2023), *e.g.*, early tumors, micrometastases in lymph nodes, and small vascular abnormalities.

Existing models for medical image segmentation such as U-Net (Ronneberger, Fischer, and Brox 2015), ViTSeg (Du et al. 2023), and MedT (Valanarasu et al. 2021; Tang et al. 2022) have made significant strides in segmenting large objects in medical images. However, these models often struggle with the segmentation of small lesions. Fig. 1 shows the

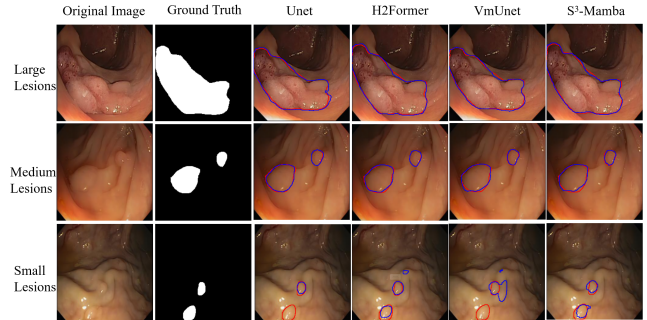


Figure 1: Segmentation results of CNN-based Unet (Ronneberger, Fischer, and Brox 2015), transformere-based H2Former (He et al. 2023), Mamba-based VmUnet (Ruan and Xiang 2024), and the proposed S³-Mamba for lesions of varying sizes. All models perform well in segmenting large and medium lesions, but S³-Mamba exhibits superior performance in segmenting small lesions, capturing fine details and contours of small lesions with higher accuracy.

results of several models for segmenting lesions. It is evident that small lesions are inaccurately segmented due to their subtle textures and morphology. A key challenge arises from the common down-sampling operations in network architectures, which unintentionally eliminates crucial local features associated with small lesions. As these lesions only occupy a small portion of the image, they may be overlooked or inadequately represented in the segmentation output, requiring more advanced techniques to handle small lesions.

To tackle the challenges, we propose a **Small-Size-Sensitive Mamba (S³-Mamba)** to effectively and efficiently segment small lesions. In view of the balance performance of Mamba (Gu and Dao 2023) between accuracy and speed, we choose Mamba as the baseline model. The proposed S³-Mamba enhances the segmentation of small lesions through three novel techniques: an **Enhanced Visual State Space (EnVSS)** block to selectively amplify key features, an **Tensor-based Cross-feature Multi-scale Attention (TCMA)** to integrate and preserve critical spatial details, and a novel curriculum learning strategy to adapt the training process based on lesion size and sample difficulty.

More specifically, the proposed EnVSS block aims to

tackle the problem of losing local features of small lesions during down-sampling through learnable channel weights and residual connections. Specifically, we insert the **Enhanced Channel Feature Block (EnCFBlock)** into the Visual State Space block (Liu et al. 2024), which integrates global context information through squeeze, excitation and scale operations to generate channel statistics and scale channel features, so that contribution of each channel is automatically adjusted and fine-tuned based on its relevance to small lesion features. In addition, two residual connections are applied to preserve important lesion fine details by amplifying crucial features while downplaying less relevant ones. The proposed EnVSS block enhances the model’s ability to capture and emphasize the subtle characteristics of small lesions that are often overlooked in existing models.

The TCMA is designed to retain the spatial characteristics of small lesions across multiple granularities, exploiting the attentive support from feature maps of multiple scales to enhance the sensitivity to small lesions. The novel tensor-based attention provides an effective mechanism to not only integrate multi-modality features at multiple scales but also dynamically adjust the attention based on the interplay between small lesion areas and their surrounding background context. Unlike existing methods that often treat each feature scale independently or focus solely on the lesion (Fiaz et al. 2024; Huang et al. 2024), the proposed TCMA uniquely emphasizes the relationship between the lesion and the larger surrounding regions at multiple granularities. This dual focus well preserves the spatial integrity of small lesions while simultaneously using the broader context to ensure that critical small lesions are not overlooked.

Traditional random sampling strategies (Singh et al. 2021) often prioritizes the accurate segmentation of larger lesions due to their dominance in the image, while neglecting small lesions. Although some workaround techniques such as data augmentation (Bosquet et al. 2023) have been explored to replicate small lesions to improve their representation, the bias towards larger objects remains a challenge. Inspired by (Wang, Chen, and Zhu 2021a), we propose a novel regularized curriculum learning strategy to dynamically adjust the training process based on lesion size and sample difficulty. More specifically, a Difficulty Measurer is designed to assess the the difficulty of samples, initialize sample weights based on lesion size and dynamically update them according to loss values during training. A Training Scheduler with regularization constraints is designed to gradually shift the focus from simpler samples to more challenging samples such as small lesions as training converges. This strategy enables the model to prioritize more challenging cases and adaptively focus on small lesion segmentation.

Our main contributions can be summarized as follows. 1) The proposed EnVSS block dynamically fine-tunes the contribution of each channel based on its relevance to small lesions through the channel-wise attention, and effectively preserves crucial lesion fine details through two residual connections. 2) The proposed TCMA employs a novel tensor-based multi-level attention strategy to integrate multi-modality features across scales and dynamically adjust the attention based on the interaction between small lesions and

their surrounding, preserving small lesion integrity while leveraging surrounding context to better segment small lesions. 3) The proposed regularized curriculum learning assesses the sample difficulty through a Difficulty Measurer and gradual shifts from easier samples to more challenging ones such as small lesions, through a Training Scheduler.

Related Work

Medical Image Segmentation. Existing models for medical image segmentation can be broadly categorized into four groups based on their architecture: 1) Models based on convolutional neural networks (CNNs), like UNet (Ronneberger, Fischer, and Brox 2015), VNet (Millettari, Navab, and Ahmadi 2016) and SegNet (Badrinarayanan, Kendall, and Cipolla 2017); 2) Transformer-based models such as UNETR++ (Shaker et al. 2024), MedT (Valanarasu et al. 2021) and H2Former (He et al. 2023); 3) Hybrid models combining CNN and Transformer architectures, *e.g.*, SwinUNet (Cao et al. 2022), SwinUNETR (Tang et al. 2022) and MedFormer (Gao et al. 2022); and 4) State space models, like VMUNet (Ruan and Xiang 2024), UMamba (Ma, Li, and Wang 2024), and SegMamba (Xing et al. 2024).

Although these models have made significant strides, small lesion segmentation remains challenging due to the down-sampling operations that commonly exist in these models, which may overlook small targets. In literature, there are some attempts for small object detection, *e.g.*, increase the input image size to generate high-resolution feature maps (Bosquet et al. 2023; Wang et al. 2024b), feature enhancement (Zhang et al. 2024), feature pyramids (Gao et al. 2024; Mei et al. 2023) and tuning loss functions (Miao et al. 2023). These methods offer valuable insights for segmenting small lesions. But the unique characteristics of medical images, *e.g.*, low contrast, complex anatomical backgrounds, and diverse imaging modalities, often complicate segmentation tasks and increase the difficulty of accurately segmenting small lesions, highlighting the urgent need for tailored models to address these challenges.

Very recently, Mamba models (Zhu et al. 2024; Liu et al. 2024) have gained traction, especially for medical image segmentation. U-Mamba (Ma, Li, and Wang 2024) extends the Mamba to the U-Net architecture, aiming to capture long-range dependencies using hybrid CNN-SSM (State Space Model) blocks. VM-UNET (Ruan and Xiang 2024) incorporates the Visual State Space (VSS) block to capture contextual information, and its asymmetric encoder-decoder structure demonstrates the capabilities of SSM for segmentation. These Mamba-based models predominantly emphasize long-range dependencies and contextual information, while sacrificing the accuracy of segmenting small lesions. To bridge this gap, we propose an enhanced VSS block that reduces the computational complexity while improving the segmentation accuracy of small lesions.

Curriculum Learning. Curriculum learning (Wang, Chen, and Zhu 2021a) was initially introduced as a training strategy to gradually increase sample complexity during training, mimicking the structured learning of humans and improving model generalization and convergence speed (Wang,

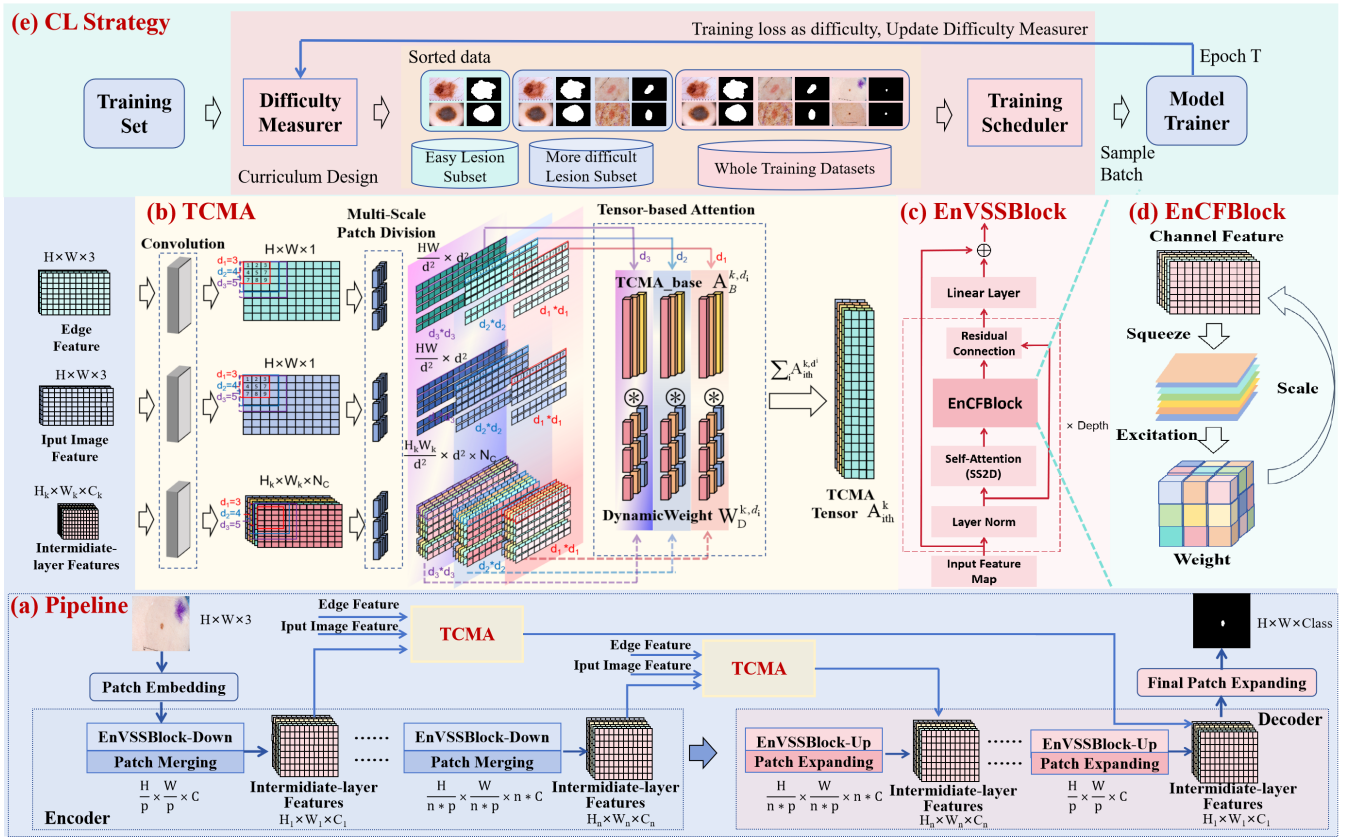


Figure 2: (a) Overview of the proposed S^3 -Mamba with *TCMA* and *EnVSSBlock*. (b) Detailed architecture of the *TCMA*, where input image features, intermediate-layer features, and edge features are divided into patches of three different sizes. A tensor-based attention derives the dynamic weights of these patches of three different scales, exploits their interaction and utilizes the *TCMA* features to modulate the features at decoder layers. (c) Detailed structure of the *EnVSSBlock*, which explicitly evaluates and adaptively adjusts the channel weights to enhance small lesion feature representation, and preserves fine details through residual connections. (d) Detailed structure of the Enhanced Channel Feature Block (*EnCFBlock*), which enhances the feature interaction through channel-wise interaction. (e) Architecture of the regularized curriculum learning strategy.

Chen, and Zhu 2021b). Recent work has extended curriculum learning to other tasks such as training Generative Adversarial Network (Doan et al. 2019) and image classification (Dogan et al. 2020). In this paper, a novel curriculum learning strategy with regularization is designed to help focus and more accurately segment challenging small lesions.

Proposed Method

Overview of Proposed S^3 -Mamba

As shown in Fig. 2(a), our S^3 -Mamba is built upon a U-shape encoder-decoder architecture, including an Enhanced Visual State Space Block (*EnVSSBlock*) and a Tensor-based Cross-feature Multi-scale Attention (*TCMA*). A patch embedding operation (Liu et al. 2024) is first applied to the original image to generate initial feature maps. The encoder comprises a series of *EnVSSBlocks*_{Down} and patch merging operations, progressively capturing global dependencies, extracting small lesion features, and compressing spatial dimensions. The *TCMA* makes use of edge features, input image features, and intermediate-layer features, and divides

them into patches of three distinct sizes, aiming to enrich the diversity of extracted features and focus on small lesions. Then, a novel tensor-based attention exploits the attentive supports from these multi-modal multi-scale patch features, and merges them to produce the final *TCMA* features, which are applied to tune the features from decoder layers. The decoder consists of a series of *EnVSSBlocks*_{Up} and patch expanding operations that progressively reconstruct the spatial dimensions while enhancing small lesion features. To further focus on small lesions, we introduce a regularized curriculum learning strategy, including a Difficulty Measurer to assess sample difficulty and a Training Scheduler to focus on segmenting challenges targets like small lesions.

Enhanced Visual State Space Block

VSSBlock is the core module of VMamba (Liu et al. 2024), capable of effectively capturing rich contextual information through deep convolution and *SS2D*, making it perform well in segmentation tasks. However, the *VSSBlock* has limitations in retaining fine details due to the local aggregation of

deep convolution and the indiscriminate handling of channel features, causing these fine features to be excessively smoothed or overshadowed by more prominent background information. To tackle these challenges, we propose the *EnVSSBlock* to better preserve image fine details.

The improvements over *VSSBlock* are three-fold. 1) We remove the *DWConv* layer in *VSSBlock*, thereby avoiding excessive smoothing and reducing propagation loss of information, enabling the model to retain more detailed features. 2) We introduce two residual connections as shown in Fig. 2(c), one from the input directly to the output that helps maintain the integrity of the initial input features, preventing the information loss in a deep network, and the other highlighting the important features to small lesions while preserving local details. 3) An *EnCFBlock* is added to enhance channel features through adaptive reweighting, suppressing irrelevant background context, and highlighting the key features to more accurately represent small lesions. These improvements collectively help the model focus on segmenting small lesions by preserving local details and adaptively highlighting the features relevant to small lesions.

The *EnCFBlock* exploits the channel-wise attention as in (Hu, Shen, and Sun 2018). Specifically, a global average pooling is first applied to capture the global context of each channel; The squeezed features are then passed through two fully connected layers to further reduce the dimensionality and then restore it, followed by a Sigmoid activation function; These adaptive weights are then applied to the previous channel features, enhancing important features while suppressing less relevant ones; Finally, a residual connection is employed to preserve locality. Hence, the *EnCFBlock* enhances the feature interaction through channel-wise attention to highlight local features relevant to small lesions.

Tensor-based Cross-feature Multi-scale Attention

Existing attention mechanisms such as Soft Attention (Omeroglu et al. 2023), Hard Attention (Jegham et al. 2023), Self-Attention (Cao et al. 2023), and Multi-Head Attention (Wang et al. 2024a) have demonstrated their power in various tasks, but they exhibit significant limitations when applied to small lesion segmentation, as they are not originally designed to effectively handle multi-scale information, nor can they fully integrate features of different types. What’s more, attention weights are often fixed after computation (Luvembe et al. 2023), nor adaptive to varying contexts.

To tackle the challenges, we propose a Tensor-based Cross-feature Multi-scale Attention (TCMA) to enhance the model’s ability to segment small lesions in three aspects. 1) The TCMA dynamically tunes the feature representations based on multi-feature multi-scale tensors through a novel tensor-based attention mechanism. The tensors encapsulate the input image features, intermediate prediction features, and edge features at multiple scales. The joint utilization of these features could help exploit the information relevant to segment small lesions in a wider spatial context and a broader range of modalities. 2) The novel tensor-based attention provides an effective mechanism to exploit the attentive supports from features of different modalities at different scales. This unique approach allows for simultaneous

interaction between spatial, categorical (class-related), and edge features, enabling the model to dynamically focus on the most relevant features for precisely segmenting small lesions. 3) Finally, the integration of TCMA tensors from low-level features in encoder layers with high-level abstract features in decoder layers provides an effective mechanism to exploit both sets of features, which captures the key features of small lesions in both global context and fine details in local neighborhoods. As a result, the proposed TCMA leads to more precise segmentation of small lesions.

Generation of Multi-Feature Multi-Scale Tensors. We first construct a set of pyramid features, including the input image features embedded using a 7×7 convolutional layer and a 3×3 convolutional layer, the edge features extracted using a Sobel operator and embedded with a 3×3 convolutional layer, and intermediate prediction features. More specifically, Given the i -th level feature maps $F_i \in \mathbb{R}^{H_i \times W_i \times C_i}$, a 1×1 convolutional layer $\mathcal{F}_{1 \times 1}$ with N_c filters converts F_i into per-category prediction features as $R_i = \mathcal{F}_{1 \times 1}(F_i)$, where $R_i \in \mathbb{R}^{H_i \times W_i \times N_c}$ and N_c is the number of categories to predict. We hypothesize that $R_i^k \in \mathbb{R}^{H_i \times W_i \times 1}$ (for $1 \leq k \leq N_c$) contains only the information associated with the k -th category. We then divide the feature pyramid $R \in \mathbb{R}^{H \times W \times C}$ into patches of three different scales as,

$$P^{d_i} = \mathcal{F}_\pi(R, d_i), \quad (1)$$

where \mathcal{F}_π is the image partition operation, d_i represents the size of patches, and $P^{d_i} \in \mathbb{R}^{\frac{H_i W_i}{d_i^2} \times d_i^2 \times N_c}$. This multi-feature multi-scale tensor provides a detailed and comprehensive feature representation that enhances the model’s ability to accurately identify and segment small lesions.

Tensor-based Attention. While the tensor P^{d_i} encapsulates multi-feature multi-scale information, it is difficult to exploit the attentive support through the interaction among the tensor features. To achieve this, a novel tensor-based attention is designed to dynamically tune the weights of the tensor features. Specifically, an `einsum` operation is firstly adopted to exploit the relations between edge features, input image features and the k -th category-related features as,

$$A_B^{k, d_i} = \sum_{j=1}^{d_i^2} P_O^{d_i}(b, i, j, c) \cdot P_I^{k, d_i}(b, i, j, k) \cdot P_E^{d_i}(b, i, j, e), \quad (2)$$

where $P_O^{d_i}(b, i, j, c)$ represents the tensor value derived from the original input image features for the b -th sample in the batch, the i -th patch, the j -th pixel and the c -th channel; $P_I^{k, d_i}(b, i, j, k)$ represents the tensor value related to the k -th category from the intermediate prediction features; $P_E^{d_i}(b, i, j, e)$ represents the tensor value of edge features for the e -th channel. By summing over all pixel positions within the patch, we aggregate the information from all pixels of a patch into the TCMA base tensor A_B^{k, d_i} .

The element-wise multiplication between $P_O^{d_i}(b, i, j, c)$, $P_I^{k, d_i}(b, i, j, k)$, and $P_E^{d_i}(b, i, j, e)$ encapsulates the combined influence of the spatial, categorical (class-related), and

edge features, which are crucial for small lesion segmentation. In particular, the spatial features provide location-based context and the general appearance of the lesion, the categorical features ensure the model to focus on the relevant class (e.g., small lesions), and the edge features help to accurately delineate the boundaries of the lesion, which is often challenging due to blurring or low contrast. By combining all of them, the model generates a more robust and accurate representation A_B^{k,d_i} . $P_O^{d_i}$, P_I^{k,d_i} , and $P_E^{d_i}$ are then concatenated along the last dimension to form a comprehensive tensor $A_C^{k,d_i} \in \mathbb{R}^{B \times P \times d_i^2 \times (C+N_c+E)}$, and processed through an MLP to compute dynamic weights,

$$W_D^{k,d_i} = \mathcal{F}_{MLP}(\mathcal{F}_{Flatten}(A_C^{k,d_i})) \in \mathbb{R}^{B \times P \times d_i^2 \times O}, \quad (3)$$

where $\mathcal{F}_{Flatten}(\cdot)$ is the flatten operation, $\mathcal{F}_{MLP}(\cdot)$ represents fully connected layers with batch normalization, ReLU activation, and sigmoid activation functions, and O is the output dimension. The weighted TCMA tensor is then obtained by aggregating over the last dimension of channel as,

$$A_{ith}^{k,d_i} = \sum A_B^{k,d_i} \cdot W_D^{k,d_i}. \quad (4)$$

The TCMA tensor A_{ith}^k is obtained by integrating the tensors from different scales as $A_{ith}^k = \sum_i A_{ith}^{k,d_i}$, ensuring the model’s robustness for different sizes of lesions. The combination of different types of features helps more accurately segment small lesions, which are often subtle and difficult to be distinguished from the background. The utilization of multi-scale features allows the model to capture both fine details and broader contextual information, critical for handling lesions of varying sizes. The resulting tensor A_{ith}^k enables the model to dynamically adjust its focus on the most relevant features, leading to more precise segmentation.

Integration of TCMA Tensors with Decoder Features. The obtained TCMA tensor A_{ith}^k is then applied to modulate the output of each *EnVSSLayer_up*, enhancing the segmentation of small objects. Given a predicted mask $\hat{M} \in \mathbb{R}^{H \times W \times N_c}$ of an *EnVSSLayer_up* layer, it is first resized using fixed-size average pooling to match the dimensions of an intermediate feature map, and reshaped to $\tilde{M} \in \mathbb{R}^{H_i \times W_i \times N_c}$ to match the dimensions of TCMA tensor. Then, A_{ith}^k modulates the decoder features as,

$$\hat{M}_{ith}^k = A_{ith}^k \otimes \tilde{M}^k, \quad (5)$$

where \otimes denotes the einsum operation. This operation integrates the multi-modal multi-scale image fine details conveyed by the TCMA with the global context in the decoder features, refining the focus on small lesions in the network.

Regularized Curriculum Learning Strategy

We further refine the training strategy to fully unlock the model’s potential in segmenting small lesions through curriculum learning shown in Fig. 2(e). We establish a Difficulty Measurer based on the size of lesions. During training, we dynamically update the Difficulty Measurer based on the loss values, assigning higher weights to samples with lower loss and vice versa. As training progresses and loss stabilizes, the Training Scheduler shifts focus to difficult

samples with lower weights. Specifically, given a dataset $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, where \mathbf{x}_i represents features and \mathbf{y}_i denotes labels for each instance i , the model \mathcal{F}_ϕ , parameterized by ϕ , generates a prediction $\mathcal{F}_\phi(\mathbf{x}_i)$ for each \mathbf{x}_i and computes the corresponding loss $l_i = \mathcal{L}(\mathcal{F}_\phi(\mathbf{x}_i), \mathbf{y}_i)$, where $\mathcal{L}(\cdot)$ denotes the loss function. The primary objective is to minimize the empirical loss across the entire training set,

$$\min_{\phi} \sum_{i=1}^N v_i l_i + g(\mathbf{v}), \quad (6)$$

where $\mathbf{v} = [v_1, v_2, \dots, v_N] \in [0, 1]^N$ represents the vector of weights for training samples, determining by the Difficulty Measurer. Different from traditional curriculum learning (Wang, Chen, and Zhu 2021a), we introduce a novel regularization term $g(\mathbf{v})$ to influence the selection of v_i as,

$$g(\mathbf{v}) = \lambda \sum_{i=1}^N \left(\frac{1}{\mathcal{F}_{rank}(l_i)} \right) + (1 - \lambda) \sum_{i=1}^N v_i^2, \quad (7)$$

where λ is a weighting factor to balance the influence of weight assignment using loss ranking and the strength of the weight squared sum regularization, and $\mathcal{F}_{rank}(\cdot)$ is a rank function. Both terms serve as regularization constraints, with the former to regularize the weight distribution, ensuring the model to focus on easier samples first and gradually transit to more difficult ones such as small lesions, and the latter to regularize the sample weights v_i to prevent excessively large weights, ensuring a balanced distribution of weights.

Experimental Results

Datasets. The ISIC2018 dataset (Azad et al. 2019) contains 2,694 dermoscopy images specifically designed for lesion segmentation. The CVC-ClinicDB dataset (Jha et al. 2019), a benchmark for colonoscopy image analysis, includes 612 high-resolution colonoscopy images with corresponding polyp annotations, focusing on polyp detection and segmentation. The in-house Lymph dataset consists of 5,344 multi-modal MRI images of prostate cancer lymph nodes, including 2,733 T2-weighted (T2WI) and 2,611 T1-weighted (T1WI) images, annotated by three skilled radiologists with cross-validation to ensure reliable ground truth.

We then analyze lesion sizes across the three datasets. We sort the lesion pixel distributions from smallest to largest in the ISIC2018 and CVC-ClinicDB datasets, dividing them into three groups: the smallest 30% as small lesions, 30%-60% as medium lesions, and 60% and above as large lesions. We randomly select 30% from each group to create three separate sets for model testing, each focusing on small, medium, and large lesions, respectively. The remaining 70% samples are used for model training. The Lymph dataset predominantly consists of very small and uniformly distributed objects. We utilize 70% for training and 30% for testing without further subdivision.

Compared Methods. Nine state-of-the-art segmentation models are compared, falling into four groups based on their architecture. 1) Traditional CNN-based models such as UNet (Ronneberger, Fischer, and Brox 2015); 2) Attention-based models, including UNETR++ (Shaker et al. 2024),

MODELS		MIOU			DSC			ACC			SPE			SEN			
		S	M	L	S	M	L	S	M	L	S	M	L	S	M	L	
ISIC2018	CNN	UNet _(MICCAI'15)	64.36	79.31	83.60	77.78	88.46	91.07	98.16	97.13	92.65	98.30	97.99	96.01	94.37	90.86	88.11
		UNETR++ _(TMI'24)	55.72	74.17	82.20	71.56	85.17	90.23	97.44	96.22	92.02	97.55	97.14	95.96	94.46	89.60	86.69
	Trans	MedT _(MICCAI'21)	56.19	73.19	81.57	71.93	84.52	89.85	97.52	96.04	91.62	97.67	96.95	94.91	33.00	89.38	87.18
		H2Former _(TMI'23)	63.51	77.61	83.26	77.68	87.39	90.87	98.18	96.85	92.50	98.36	97.76	96.03	93.06	90.24	87.73
		ViTSeg _(MICCAI'23)	55.12	73.83	80.50	71.07	84.94	89.19	97.51	96.18	91.20	97.79	97.16	95.47	89.76	89.05	85.42
	Hybrid	SwinUnet _(ECCV'22)	62.89	79.36	81.38	77.21	88.49	89.74	98.14	97.19	91.67	98.34	98.31	96.14	92.53	89.10	85.62
		SwinUNETR _(CVPR'22)	71.36	78.63	82.98	83.29	88.03	90.70	98.73	97.08	92.25	98.93	98.21	94.83	93.08	88.83	88.77
		MedFormer _(arXiv'23)	60.78	77.00	84.52	75.60	87.00	91.61	97.91	96.64	92.68	98.01	97.17	93.74	95.01	92.84	91.66
	Mam	VmUnet _(MICCAI'24)	64.04	78.90	84.67	78.07	88.21	91.70	98.21	97.03	93.07	98.37	97.76	95.42	93.67	91.75	89.91
		S³-Mamba (Ours)	77.13	81.36	83.28	87.09	89.72	90.87	99.07	97.55	92.60	99.34	98.81	97.01	91.66	88.39	86.65
CVC-ClinicDB	CNN	UNet _(MICCAI'15)	69.17	78.13	79.87	81.78	87.72	88.81	99.02	98.52	96.88	99.40	99.33	99.02	84.80	86.19	84.44
		UNETR++ _(TMI'24)	36.07	51.89	55.98	53.02	68.33	71.78	97.00	96.09	92.13	97.84	97.86	96.22	65.43	68.96	68.32
	Trans	MedT _(MICCAI'21)	23.56	39.05	47.15	38.13	56.16	64.08	94.67	93.59	90.33	95.50	95.31	95.72	63.42	67.13	58.91
		H2Former _(TMI'23)	65.82	81.02	86.61	79.39	89.52	92.82	98.85	98.73	97.93	99.21	99.41	99.11	85.40	88.41	91.12
		ViTSeg _(MICCAI'23)	13.48	22.42	30.50	23.76	36.62	49.06	91.57	90.03	87.50	92.66	92.83	90.54	50.72	47.09	41.10
	Hybrid	SwinUnet _(ECCV'22)	24.06	39.43	53.89	38.79	56.56	70.04	94.56	93.55	91.90	95.30	95.18	96.59	66.60	68.59	64.61
		SwinUNETR _(CVPR'22)	49.89	65.05	71.33	66.57	78.82	83.27	98.07	97.36	95.52	98.70	98.48	98.87	74.22	80.21	76.03
		MedFormer _(arXiv'23)	60.89	74.35	80.78	75.69	85.29	89.37	98.61	98.16	96.98	99.01	98.88	98.77	83.66	87.11	86.55
	Mam	VmUnet _(MICCAI'24)	58.26	82.24	87.94	73.63	90.25	93.58	98.37	98.79	98.14	98.66	99.36	99.14	87.73	91.83	92.34
		S³-Mamba (Ours)	75.40	83.81	85.86	85.97	91.19	92.39	99.25	99.46	99.08	99.53	99.81	99.58	88.78	89.61	91.37

Table 1: Performance comparison on ISIC2018 and CVC-ClinicDB datasets. **S**, **M** and **L** represent small, medium, and large lesions respectively. Our S³-Mamba performs the best for most metrics on most settings, especially for small lesions.

MedT (Valanarasu et al. 2021), H2Former (He et al. 2023), and ViTSeg (Du et al. 2023); 3) Hybrid models combining CNN and Transformer architectures, *e.g.*, SwinUNET (Cao et al. 2022), SwinUNETR (Tang et al. 2022), and MedFormer (Gao et al. 2022); and 4) State space models such as VMUNet (Ruan and Xiang 2024).

Implementation Details. The input image size is 256×256 pixels. Common data augmentation techniques (Garcea et al. 2023) are applied to boost the performance. The backbone is initialized using VMamba-S pre-trained on ImageNet-1k. The AdamW optimizer is used with an initial learning rate of 0.0001 and a cosine scheduler. The batch size is 16. The maximum number of epochs is 600. Experiments are conducted on a Tesla V100 GPU with 32GB memory.

Evaluation Metrics. Five evaluation metrics are adopted: Mean Intersection over Union (mIoU), Dice Similarity Coefficient (DSC), Accuracy (ACC), Specificity (SPE), and Sensitivity (SEN), same as in (Hirling et al. 2024).

Comparison with State-of-the-Art Methods

ISIC2018 Dataset. The comparison results on the ISIC2018 dataset are summarized in Table 1. The following can be observed. 1) Our S³-Mamba achieves the highest mIoU of 77.13% and 81.36% for small and medium lesions respectively, and a competitive mIoU of 83.28% for large lesions. This robustness across lesion sizes is attributed to the EnVSS block’s dynamic fine-tuning of channel contributions and the TCMA’s multi-level attention strategy, both of which enhance the model’s ability to preserve small lesion details while leveraging the surrounding context. 2) The VmUnet

MODELS		MIOU	DSC	ACC	SPE	SEN
CNN	UNet _(MICCAI'15)	48.56	65.37	99.91	99.98	56.88
	UNETR++ _(TMI'24)	48.68	65.48	99.92	99.98	57.33
Trans	MedT _(MICCAI'21)	23.77	38.41	99.86	99.95	32.68
	H2Former _(TMI'23)	47.52	64.42	99.92	99.98	54.88
	ViTSeg _(MICCAI'23)	14.94	26.00	99.79	99.89	27.12
Hybrid	SwinUnet _(ECCV'22)	19.64	32.83	99.85	99.95	26.51
	SwinUNETR _(CVPR'22)	45.03	62.09	99.91	99.97	54.02
	MedFormer _(arXiv'23)	55.25	71.18	99.92	99.97	68.78
Mam	VmUnet _(MICCAI'24)	52.27	68.66	99.92	99.97	62.44
	S³-Mamba (Ours)	61.19	75.93	99.94	99.97	75.18

Table 2: Performance comparison on the Lymph dataset.

(Ruan and Xiang 2024), a predecessor of the Mamba series, excels in large lesion segmentation with a mIoU of 84.67% and a DSC of 91.70%, but it is less effective with small and medium lesions. This suggests that the VmUnet prioritizes broader spatial context, potentially at the cost of fine-grained details necessary for smaller lesions. The enhancements in our S³-Mamba, particularly in attention mechanisms, likely contribute to its improved performance across different lesion sizes. 3) UNet (Ronneberger, Fischer, and Brox 2015) achieves an mIoU of 64.36%, 79.31%, 83.60% for small, medium and large lesions respectively. While it is a foundational model in medical image segmentation, it lags behind newer models, particularly on small lesions, highlighting the need for ongoing innovation in segmentation techniques. 4) Transformer-based models like MedT achieves an mIoU




Figure 3: Segmentation results on the Lymph dataset.

of 56.19%, 73.19%, and 81.57% and H2Former achieves an mIoU of 63.51%, 77.61%, and 83.26% respective for small, medium and large lesions. They exhibit varied performance, with H2Former particularly strong in small lesion segmentation. This suggests that Transformers are highly dependent on their specific configuration and integration.

CVC-ClinicDB Dataset. The following can be observed on the CVC-ClinicDB dataset. 1) Some models such as UNETR++, MedT and ViTSeg experience a significant performance drop on this dataset compared to their performance on the ISIC2018 dataset, possibly due to the fact that the lesions in this dataset are generally small, with a lesion pixel ratio ranging from 0.44% to 50.17% only. In addition, the ISIC2018 dataset utilizes high-resolution dermoscopic images, while the CVC-ClinicDB dataset consists of colonoscopic images where lesion visual characteristics differ. 2) In contrast, our S^3 -Mamba continues to excel on the CVC-ClinicDB dataset, because of its optimization for small lesion segmentation, *e.g.*, the EnVSS block dynamically adjusts channel contributions, and the multi-level tensor-based attention integrates features across scales, preserving lesion integrity while leveraging the background information. 3) While the S^3 -Mamba exhibits strong overall performance, its sensitivity (SEN) is not the highest across both datasets. The slightly lower sensitivity could be due to the model’s balanced focus on both sensitivity and specificity.

Private Lymph Dataset. The comparison results for ultra-small lesions on the private prostate lymph dataset are summarized in Table 2. Our S^3 -Mamba outperforms all the models, achieving an mIoU of 61.19%, DSC of 75.93%, accuracy of 99.94%, and sensitivity of 75.18%, respectively. While the performance of all models decreases on the Lymph dataset compared to that on the ISIC2018 and CVC-ClinicDB datasets, our S^3 -Mamba still leads by about 5% in several metrics. Fig. 3 shows sample segmentation results, particularly focusing on small lesions. S^3 -Mamba demonstrates a closer alignment with the ground truth and produces fewer false positives compared to other models, as evident in the zoomed-in sections of the image. These observations demonstrate that S^3 -Mamba’s design innovations significantly contribute to its leading performance on

Methods	MIOU	DSC	ACC	SPE	SEN
Baseline	52.27	68.66	99.92	99.97	62.44
+EnCFBlock	55.70	71.55	99.93	99.97	67.06
+TCMA_S	53.08	69.35	99.93	99.98	62.67
+TCMA_M	55.66	71.51	99.93	99.97	68.67
+CL	55.15	71.09	99.92	99.97	67.82
S^3-Mamba	61.19	75.93	99.94	99.97	75.18

Table 3: Ablation of key components on the Lymph dataset.

the Lymph dataset, particularly in small lesion segmentation.

Ablation Study of Key Components

We ablate key components of our S^3 -Mamba on the Lymph dataset. For the baseline model, both *EnCFBlock* and TCMA are removed, and it is trained without curriculum learning. On top of the baseline model, the *EnCFBlock*, TCMA utilizing single-scale patches (denoted as +TCMA_S), TCMA utilizing multi-scale patches (denoted as +TCMA_M), and the baseline model trained using the curriculum learning (denoted as +CL) are added for ablation. As shown in Table 3, the baseline performs relatively poorly, with an mIoU of 52.27%. By adding the *EnCFBlock*, the mIoU increases to 55.70%, indicating its effectiveness in enhancing small lesion features. The single-scale TCMA provides only a slight improvement of 0.81% in terms of mIoU, suggesting that single-scale features are insufficient. The multi-scale TCMA significantly boosts the mIoU to 55.66%, underscoring the importance of multi-scale features and the power of the tensor-based attention. The curriculum learning strategy enhances the mIoU to 55.15%, showing its effectiveness in focusing on small lesions. The complete S^3 -Mamba model, which combines the *EnCFBlock*, TCMA, and curriculum learning, achieves the highest performance across all metrics, with an mIoU of 61.19%.

Conclusion

In this paper, we propose Small-Size-Sensitive Mamba (S^3 -Mamba), which significantly improves small lesion segmentation by introducing the *EnVSSBlock*, Tensor-based Cross-feature Multi-scale Attention, and a novel regularized curriculum learning strategy. In particular, the *EnVSSBlock* enhances the feature representation by adding residual connections to preserve local features and adaptively adjusting channel weights to amplify important features, particularly the subtle ones from small lesions through the channel-wise attention. The TCMA captures multi-scale information from patches of different modalities of different scales through a novel tensor-based attention mechanism, combining fine details with global context, thus improving the model’s capability in segmenting small lesions. Finally, the novel regularized curriculum learning strategy enables the model to gradually learn features from large to small lesions, enhancing the segmentation accuracy and robustness. The combination of these strategies achieves significant performance gains over state-of-the-art models on three benchmark datasets, significantly enhancing small lesion segmentation.

References

- Ali, S.; Jha, D.; Ghatwary, N.; Realdon, S.; Cannizzaro, R.; Salem, O. E.; Lamarque, D.; Daul, C.; Riegler, M. A.; Anonsen, K. V.; et al. 2023. A multi-centre polyp detection and segmentation dataset for generalisability assessment. *Scientific Data*, 10(1): 75.
- Azad, R.; Asadi-Aghbolaghi, M.; Fathy, M.; and Escalera, S. 2019. Bi-directional ConvLSTM U-Net with densely connected convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, 0–0.
- Badrinarayanan, V.; Kendall, A.; and Cipolla, R. 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12): 2481–2495.
- Bosquet, B.; Cores, D.; Seidenari, L.; Brea, V. M.; Mucientes, M.; and Del Bimbo, A. 2023. A full data augmentation pipeline for small object detection based on generative adversarial networks. *Pattern Recognition*, 133: 108998.
- Cao, H.; Wang, Y.; Chen, J.; Jiang, D.; Zhang, X.; Tian, Q.; and Wang, M. 2022. Swin-unet: Unet-like pure transformer for medical image segmentation. In *European conference on computer vision*, 205–218. Springer.
- Cao, M.; Wang, X.; Qi, Z.; Shan, Y.; Qie, X.; and Zheng, Y. 2023. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22560–22570.
- de Teresa-Trueba, I.; Goetz, S. K.; Mattausch, A.; Stojanovska, F.; Zimmerli, C. E.; Toro-Nahuelpan, M.; Cheng, D. W.; Tollrvey, F.; Pape, C.; Beck, M.; et al. 2023. Convolutional networks for supervised mining of molecular patterns within cellular context. *Nature Methods*, 20(2): 284–294.
- Doan, T.; Monteiro, J.; Albuquerque, I.; Mazouze, B.; Durand, A.; Pineau, J.; and Hjelm, R. D. 2019. On-line adaptive curriculum learning for gans. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3470–3477.
- Dogan, Ü.; Deshmukh, A. A.; Machura, M. B.; and Igel, C. 2020. Label-similarity curriculum learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, 174–190. Springer.
- Du, S.; Bayasi, N.; Hamarneh, G.; and Garbi, R. 2023. AViT: Adapting Vision Transformers for Small Skin Lesion Segmentation Datasets. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 25–36. Springer.
- Duering, M.; Biessels, G. J.; Brodtmann, A.; Chen, C.; Cordonnier, C.; de Leeuw, F.-E.; Debette, S.; Frayne, R.; Jouvent, E.; Rost, N. S.; et al. 2023. Neuroimaging standards for research into small vessel disease—advances since 2013. *The Lancet Neurology*, 22(7): 602–618.
- Fiaz, M.; Noman, M.; Cholakkal, H.; Anwer, R. M.; Hanna, J.; and Khan, F. S. 2024. Guided-attention and gated-aggregation network for medical image segmentation. *Pattern Recognition*, 110812.
- Gao, J.; Geng, X.; Zhang, Y.; Wang, R.; and Shao, K. 2024. Augmented weighted bidirectional feature pyramid network for marine object detection. *Expert Systems with Applications*, 237: 121688.
- Gao, Y.; Zhou, M.; Liu, D.; Yan, Z.; Zhang, S.; and Metaxas, D. N. 2022. A data-scalable transformer for medical image segmentation: architecture, model efficiency, and benchmark. *arXiv preprint arXiv:2203.00131*.
- Garcea, F.; Serra, A.; Lamberti, F.; and Morra, L. 2023. Data augmentation for medical imaging: A systematic literature review. *Computers in Biology and Medicine*, 152: 106391.
- Greenwald, N. F.; Miller, G.; Moen, E.; Kong, A.; Kagel, A.; Dougherty, T.; Fullaway, C. C.; McIntosh, B. J.; Leow, K. X.; Schwartz, M. S.; et al. 2022. Whole-cell segmentation of tissue images with human-level performance using large-scale data annotation and deep learning. *Nature biotechnology*, 40(4): 555–565.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- He, A.; Wang, K.; Li, T.; Du, C.; Xia, S.; and Fu, H. 2023. H2Former: An efficient hierarchical hybrid transformer for medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(9): 2763–2775.
- Hirling, D.; Tasnadi, E.; Caicedo, J.; Caroprese, M. V.; Sjögren, R.; Aubreville, M.; Koos, K.; and Horvath, P. 2024. Segmentation metric misinterpretations in bioimage analysis. *Nature methods*, 21(2): 213–216.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7132–7141.
- Huang, H.; Chen, Z.; Zou, Y.; Lu, M.; Chen, C.; Song, Y.; Zhang, H.; and Yan, F. 2024. Channel prior convolutional attention for medical image segmentation. *Computers in Biology and Medicine*, 178: 108784.
- Jegham, I.; Alouani, I.; Khalifa, A. B.; and Mahjoub, M. A. 2023. Deep learning-based hard spatial attention for driver in-vehicle action monitoring. *Expert Systems with Applications*, 219: 119629.
- Jha, D.; Smedsrud, P. H.; Riegler, M. A.; Johansen, D.; De Lange, T.; Halvorsen, P.; and Johansen, H. D. 2019. Resunet++: An advanced architecture for medical image segmentation. In *2019 IEEE international symposium on multimedia (ISM)*, 225–2255. IEEE.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024. Vmamba: Visual state space model 2024. *arXiv preprint arXiv:2401.10166*.
- Luvembe, A. M.; Li, W.; Li, S.; Liu, F.; and Xu, G. 2023. Dual emotion based fake news detection: A deep attention-weight update approach. *Information Processing & Management*, 60(4): 103354.
- Ma, J.; Li, F.; and Wang, B. 2024. U-mamba: Enhancing long-range dependency for biomedical image segmentation. *arXiv preprint arXiv:2401.04722*.
- Ma, J.; and Wang, B. 2023. Towards foundation models of biological image segmentation. *Nature Methods*, 20(7): 953–955.

- Mei, Y.; Fan, Y.; Zhang, Y.; Yu, J.; Zhou, Y.; Liu, D.; Fu, Y.; Huang, T. S.; and Shi, H. 2023. Pyramid attention network for image restoration. *International Journal of Computer Vision*, 131(12): 3207–3225.
- Miao, J.; Zhou, S.-P.; Zhou, G.-Q.; Wang, K.-N.; Yang, M.; Zhou, S.; and Chen, Y. 2023. SC-SSL: Self-correcting Collaborative and Contrastive Co-training Model for Semi-Supervised Medical Image Segmentation. *IEEE Transactions on Medical Imaging*.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, 565–571. Ieee.
- Omeroglu, A. N.; Mohammed, H. M.; Oral, E. A.; and Aydin, S. 2023. A novel soft attention-based multi-modal deep learning framework for multi-label skin lesion classification. *Engineering Applications of Artificial Intelligence*, 120: 105897.
- Qureshi, I.; Yan, J.; Abbas, Q.; Shaheed, K.; Riaz, A. B.; Wahid, A.; Khan, M. W. J.; and Szczuko, P. 2023. Medical image segmentation using deep semantic-based methods: A review of techniques, applications and emerging trends. *Information Fusion*, 90: 316–352.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Ruan, J.; and Xiang, S. 2024. Vm-unet: Vision mamba unet for medical image segmentation. *arXiv preprint arXiv:2402.02491*.
- Shaker, A. M.; Maaz, M.; Rasheed, H.; Khan, S.; Yang, M.-H.; and Khan, F. S. 2024. UNETR++: delving into efficient and accurate 3D medical image segmentation. *IEEE Transactions on Medical Imaging*.
- Singh, D. K.; Rai, S. N.; Joseph, K.; Saluja, R.; Balasubramanian, V. N.; Arora, C.; Subramanian, A.; and Jawahar, C. 2021. Order: Open world object detection on road scenes. In *Proc. NeurIPS Workshops*, volume 1, 3.
- Tang, Y.; Yang, D.; Li, W.; Roth, H. R.; Landman, B.; Xu, D.; Nath, V.; and Hatamizadeh, A. 2022. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20730–20740.
- Valanarasu, J. M. J.; Oza, P.; Hacihaliloglu, I.; and Patel, V. M. 2021. Medical transformer: Gated axial-attention for medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, 36–46. Springer.
- Wang, J.; Lai, C.; Wang, Y.; and Zhang, W. 2024a. EMAT: Efficient feature fusion network for visual tracking via optimized multi-head attention. *Neural Networks*, 172: 106110.
- Wang, P.; Ma, Z.; Dong, B.; Liu, X.; Ding, J.; Sun, K.; and Chen, Y. 2024b. Generative data augmentation by conditional inpainting for multi-class object detection in infrared images. *Pattern Recognition*, 153: 110501.
- Wang, X.; Chen, Y.; and Zhu, W. 2021a. A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(9): 4555–4576.
- Wang, X.; Chen, Y.; and Zhu, W. 2021b. A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(9): 4555–4576.
- Xing, Z.; Ye, T.; Yang, Y.; Liu, G.; and Zhu, L. 2024. Segmamba: Long-range sequential modeling mamba for 3d medical image segmentation. *arXiv preprint arXiv:2401.13560*.
- Zhang, Y.; Ye, M.; Zhu, G.; Liu, Y.; Guo, P.; and Yan, J. 2024. FFCA-YOLO for small object detection in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.

Reproducibility Checklist

This paper:

- Includes a conceptual outline and/or pseudocode description of AI methods introduced? (yes)
- Clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results. (yes)
- Provides well marked pedagogical references for less-familare readers to gain background necessary to replicate the paper. (yes)

Does this paper make theoretical contributions? (yes)

If yes, please complete the list below.

- All assumptions and restrictions are stated clearly and formally. (yes)
- All novel claims are stated formally (e.g., in theorem statements). (yes)
- Proofs of all novel claims are included. (yes)
- Proof sketches or intuitions are given for complex and/or novel results. (yes)
- Appropriate citations to theoretical tools used are given. (yes)
- All theoretical claims are demonstrated empirically to hold. (yes)
- All experimental code used to eliminate or disprove claims is included. (yes)

Does this paper rely on one or more datasets? (yes)

If yes, please complete the list below.

- A motivation is given for why the experiments are conducted on the selected datasets? (yes)
- All novel datasets introduced in this paper are included in a data appendix. (yes)

- All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes. (NA)
- All datasets drawn from the existing literature (potentially including authors' own previously published work) are accompanied by appropriate citations. (yes)
- All datasets drawn from the existing literature (potentially including authors' own previously published work) are publicly available. (yes)
- All datasets that are not publicly available are described in detail, with explanation why publicly available alternatives are not scientifically satisfying. (yes)

Does this paper include computational experiments?

(yes)

If yes, please complete the list below.

- Any code required for pre-processing data is included in the appendix. (yes).
- All source code required for conducting and analyzing the experiments is included in a code appendix. (yes)
- All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes. (yes)
- All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from (yes)
- If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results. (yes)
- This paper specifies the computing infrastructure used for running experiments (hardware and software), including GPU/CPU models; amount of memory; operating system; names and versions of relevant software libraries and frameworks. (yes)
- This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics. (yes)
- This paper states the number of algorithm runs used to compute each reported result. (yes)
- Analysis of experiments goes beyond single-dimensional summaries of performance (e.g., average; median) to include measures of variation, confidence, or other distributional information. (yes)
- The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank). (yes)
- This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments. (yes)
- This paper states the number and range of values tried per (hyper-) parameter during development of the paper, along with the criterion used for selecting the final parameter setting. (yes)