

BA-SAM: Scalable Bias-Mode Attention Mask for Segment Anything Model

Yiran Song^{1*}, Qianyu Zhou^{1*}, Xiangtai Li², Deng-Ping Fan³, Xuequan Lu⁴, Lizhuang Ma^{1†}

¹Shanghai Jiao Tong University; ²S-Lab, Nanyang Technological University;

³Nankai University; ⁴La Trobe University

¹{songyiran, zhouqianyu, lzma}@sjtu.edu.cn,

²xiangtai.li@ntu.edu.sg, ³dengpfan@gmail.com ⁴xuequan.lu@deakin.edu.au

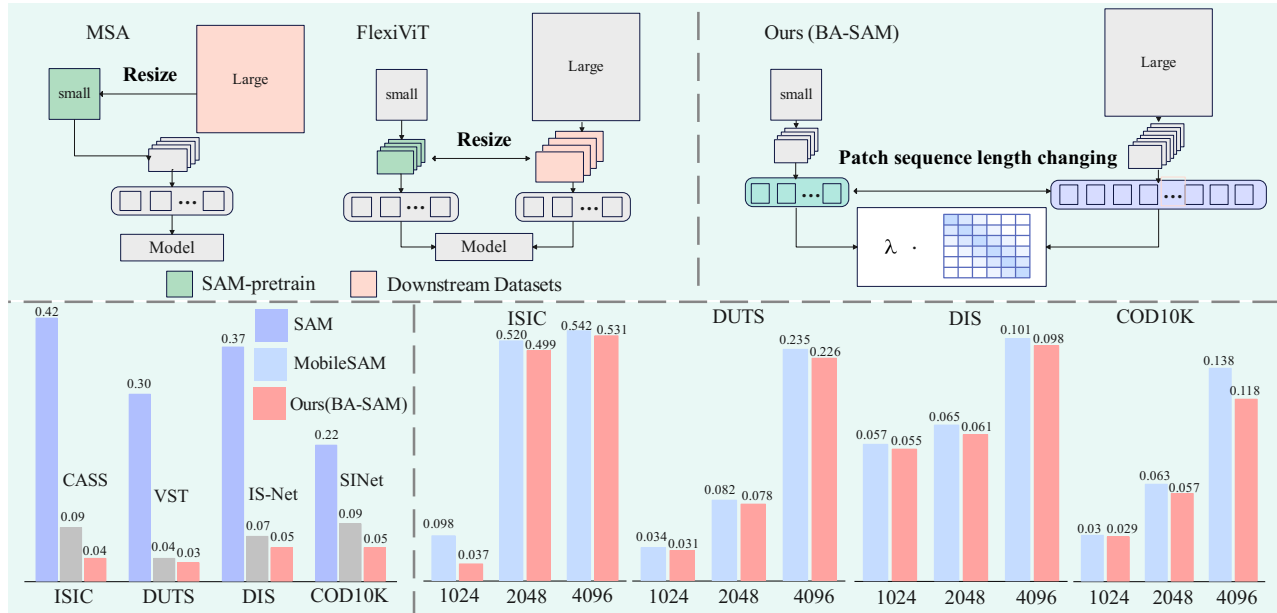


Figure 1. (a) Top: contrast between prior methods [4, 79] and BA-SAM. For large-scale datasets, previous approaches often resize images or change patch sizes to handle the issue of varying resolutions. In contrast, we propose a Scalable Bias-Mode Attention Mask (BA-SAM), which enhances SAM’s adaptability to varying image resolutions while eliminating structure modifications. Bottom (left): We introduce a generalized model that does not require any task-specific designs, yet outperforms state-of-the-art methods across four datasets. Bottom (right): With resolution variations, prior models’ performance degrades drastically. Instead, BA-SAM consistently alleviates this issue.

Abstract

In this paper, we address the challenge of image resolution variation for the Segment Anything Model (SAM). SAM, known for its zero-shot generalizability, exhibits a performance degradation when faced with datasets with varying image sizes. Previous approaches tend to resize the image to a fixed size or adopt structure modifications, hindering the preservation of SAM’s rich prior knowledge. Besides, such task-specific tuning necessitates a complete retraining of the model, which is cost-expensive and unacceptable for deployment in the downstream tasks. In this paper, we reformulate this issue as a length extrapolation problem, where token sequence length varies while maintaining a consistent patch size for images of different sizes. To

this end, we propose Scalable Bias-Mode Attention Mask (BA-SAM) to enhance SAM’s adaptability to varying image resolutions while eliminating the need for structure modifications. Firstly, we introduce a new scaling factor to ensure consistent magnitude in the attention layer’s dot product values when the token sequence length changes. Secondly, we present a bias-mode attention mask that allows each token to prioritize neighboring information, mitigating the impact of untrained distant information. Our BA-SAM demonstrates efficacy in two scenarios: zero-shot and fine-tuning. Extensive evaluation on diverse datasets, including DIS5K, DUTS, ISIC, COD10K, and COCO, reveals its ability to significantly mitigate performance degradation in the zero-shot setting and achieve state-of-the-art performance with minimal fine-tuning. Furthermore, we propose a generalized model and benchmark, showcasing BA-SAM’s generalizability across all four datasets simultaneously.

*Equal contribution.

†Corresponding author.

1. Introduction

Recently, the computer vision community [9, 22, 26, 29–31, 49, 50, 58, 66, 74, 83–90] has experienced a surge in the development of various foundation models [21, 35, 54]. Notably, Meta has introduced SAM (Segment Anything Model) [38], a prompt model that has made a significant impact. SAM can segment any object in an image or video by incorporating a single visual prompt, such as a box or a point, without requiring additional training. SAM is trained on an extensive SA-1B dataset [38], consisting of over 11 million images and one billion masks. Its emergence has undeniably showcased robust generalization capabilities across diverse images and objects, paving the way for new possibilities and avenues in intelligent image analysis and understanding [8, 34, 76, 79]. Based on SAM, some variants have been proposed, such as MobileSAM [76] and SAM-Adapter [8]. These efforts typically focus on improving SAM’s performance on specific datasets.

During the pre-training of SAM [38], the input image size is fixed at 1024. As a foundational model, SAM is expected to exhibit generalization capabilities across various downstream tasks, each associated with datasets featuring different image sizes. This is particularly crucial for high-resolution (HQ) datasets characterized by larger dimensions and more details. SAM performs well when the resolutions align with its training resolution of 1024. However, significant performance degradation is observed when inferring with resolutions larger than 1024. Hence, we aim to study a practical and realistic problem to enhance SAM’s adaptability to varying image resolutions of different datasets.

Since SAM adopts the standard Vision Transformer [14] architecture, there are two common approaches to address the inconsistency between training and inference sizes for the ViT architecture. As depicted in Figure 1, the first approach, *e.g.*, MSA [79] and SAM-Adapter [8], involves directly resizing all datasets to match a predefined size. Conversely, the second approach, exemplified by FlexiViT [4], entails adjusting the patch size to accommodate larger image resolutions. Nevertheless, tuning the image or patch size necessitates a complete retraining of the model, which is cost-expensive and unacceptable for deployment in the downstream tasks. Besides, it prevents leveraging the rich prior knowledge reserved in the pre-trained model of SAM. As such, our objective is to explore a solution that enhances SAM’s adaptability to datasets of varying resolutions while avoiding structural modifications to SAM.

In this paper, we introduce a novel perspective that reframes the challenge of image resolution variation as a length extrapolation problem. Specifically, as depicted in Figure 1, for images of varying sizes, we employ different token sequence lengths while keeping a consistent patch size. It has been observed that the inconsistency in token length between training and prediction is a key factor in per-

formance degradation. This inconsistency manifests in two aspects: Firstly, changes in token length lead to variations in the magnitude of attention module values. When the dot product result becomes significantly large in magnitude, it can drive the subsequent Softmax layer into regions with minimal gradients. Consequently, the attention distribution after Softmax becomes highly concentrated, giving rise to the issue of vanishing gradients. Secondly, longer predictions rely on untrained information, such as additional position encodings. The introduction of untrained parameters brings a considerable amount of noise to the model, which, in turn, affects its performance.

To address these issues, we propose a Scalable Bias-Mode Attention Mask (BA-SAM) to enhance the length extrapolation capability of SAM. Our approach introduces two novel designs. Firstly, to ensure consistency in the attention layer’s dot product value, we present an improved scaling factor. This factor effectively regulates the magnitude of values within the attention layer, mitigating disruptive effects resulting from substantial changes in dot product operations and context length. Secondly, with a focus on maintaining consistency in attention focus areas, we introduce a novel bias-mode attention mask. This attention mask penalizes attention scores between distant query-key pairs, with the penalty increasing as the distance between the key and query grows. Consequently, when the context length varies, the influence of untrained distant information on each token diminishes. We achieve this mask by adding a bias after the query-key dot product, and this design is highly lightweight and could be seamlessly integrated into SAM-based models with minimal computational overhead.

Our approach demonstrates efficacy in two scenarios: zero-shot and fine-tuning. Extensive evaluation on datasets from five diverse tasks are conducted, including DIS5K [53], DUTS [64], ISIC [12], COD10K [16], and COCO [45]. These datasets vary in resolution, mostly exceeding SAM’s default resolution of 1024. In the zero-shot setting, our BA-SAM alleviates the model’s performance degradation caused by expanding the inference resolution without requiring any additional training. With a few fine-tuning epochs on downstream tasks, our BA-SAM consistently achieves state-of-the-art accuracy across all datasets. Additionally, to further demonstrate BA-SAM’s generalizability, we propose a generalized model and a new benchmark, which utilize one model to attain state-of-the-art performance across all four datasets simultaneously.

2. Related Work

Visual Foundation Models. Models that are trained on broad and can be adapted to numerous downstream tasks are called “Foundation Models” [5, 43, 67]. These models, Vision-Language Models (VLM) (CLIP [54] and DALL-E [55]) combine computer vision and natural language pro-

cessing to understand and generate descriptions or analyze visual content using textual and visual information. Masked Image Modeling [48, 73] (MIM) refers to a technique where parts of an image are masked during training to encourage a model to learn contextual information and complete missing regions. SAM [38] is a model designed for segmenting objects or areas in images, offering precise segmentation capabilities. In this paper, we use a variant of SAM called MobileSAM [76] as the baseline method. As our design optimizes the Transformer itself, it applies to all Transformer-based vision foundation models.

Resolution Variation Processing. To enable models to be more adaptable to variations in resolutions, traditional methods to deal with ViT have relied on adjustments to positional embeddings [40] and patch sizes [4, 6, 24, 30, 39, 44, 75]. Patch n’ Pack [13] employed sequence packing during training to handle inputs with arbitrary resolutions and aspect ratios. All of them necessitate training from scratch, incurring substantial computational and time costs. In contrast to previous approaches, we extend the concept of ‘length extrapolation’ from NLP into the context of addressing scale variations in CV. Length extrapolation refers to a model’s ability to generalize well to longer inputs than those it was trained on. In NLP, it has been successfully used, such as in ALIBI [52] and KERPLE [11], to enable models to adapt to longer sequences without significant performance degradation. Our approach seamlessly extends to two scenarios: zero-shot and finetuning, allowing us to leverage prior knowledge embedded in the SAM and significantly reduce training efforts.

Parameter Efficient Tuning. There have been some pioneering works for the Parameter Efficient Tuning (PETL) of visual models, such as AdaptFormer [7] and visual prompt tuning (VPT) [36]. He et al. [28] analyzed the unified view among PETL techniques such as prefixtuning [42], Prompt-tuning [36], and adapter [7]. Our method belongs to the category of Parameter Efficient Tuning.

Visual Attention Modeling. Various studies have incorporated attention mechanisms into neural network architectures designed for visual tasks [3, 33, 41, 63, 68, 93]. These mechanisms are employed in a channel-wise manner to capture cross-feature information [10, 66, 80]. They are also used for selecting paths in different branches of a network [61], or a combination of both strategies [78]. For instance, squeeze-and-excite network [32] features an attention-like module for modeling channel-wise relationships within layer features. Li *et al.* [41] utilize the attention mechanism to adapt the receptive field of neurons between network branches. The advent of transformers has led to hybrid architectures that integrate other modules. Bello’s work [2] introduces approximate content attention with a positional attention component. Child *et al.* [69] observe that many early layers in the network learn

locally connected patterns akin to convolutions, indicating that hybrid architectures inspired by both transformers and convolutional networks are a compelling design choice. Several recent studies explore this approach for various tasks [25, 59, 65, 70]. In contrast to prior work, we do not introduce a new attention structure. Instead, we offer theoretical proof for optimizing existing attention mechanisms. This resulting optimization approach is applicable across various attention designs and demonstrates strong performance across multiple datasets.

3. Preliminaries

SAM. Segment Anything Model (SAM) [38] consists of three core modules: image encoder, prompt encoder, and mask decoder. The SAM model has been trained on the extensive SA-1B dataset [38], which comprises more than 1 billion automatically generated masks. This dataset contains 400 times more masks than any existing segmentation dataset (COCO [45], LVIS [27] and ADE20K [82]) and is accompanied by 11 million images and 1 billion masks. Consequently, SAM exhibits valuable and robust zero-shot generalization to new data without necessitating further training. For further details regarding the SAM methodology, readers are encouraged to refer to [38]. Our Scalable Bias-Mode Attention Mask (BA-SAM) focuses its optimization on the image encoder while keeping the structures of the mask decoder and prompt encoder unchanged.

Scaling Factor in Attention. The two most commonly used attention functions are additive attention [1] and dot-product attention [62]. The vanilla Transformer chooses dot-product attention because it is much faster and more space-efficient in practice. However, for large values of the dimension of the tokens d_k , the dot products grow large in magnitude, pushing the softmax function into regions with minimal gradients. They use the *scaling factor* $\lambda_d = \frac{1}{\sqrt{d_k}}$ to scale the dot products. To better analyze the role of the scaling factor, we express the output element \mathbf{O}_i and the weight coefficient $a_{i,j}$ as follows:

$$\mathbf{O}_i = \sum_{j=1}^N a_{i,j} x_j v_j, \quad a_{i,j} = \frac{e^{\lambda_d q_i \cdot k_j}}{\sum_{j=1}^N e^{\lambda_d q_i \cdot k_j}}, \quad (1)$$

where λ_d represents the scaling factor.

Absolute & Relative Position Encoding in Attention.

The original Transformer [62] incorporates absolute non-parametric positional encoding $p = (p_1, \dots, p_n)$ with the input token embedding x as $x_i = x_i + p_i$, where $p_i, x_i \in \mathbb{R}^{d_x}$. Other works replace them with parametric encoding [23] or adopted Fourier-based kernelized versions [51]. Absolute position encoding enforces a fixed size for all inputs. Recent work [56] considers the pairwise relationships between elements. Relative relation is important for tasks




Figure 2. Illustration of the proposed BA-SAM method. (a) In the original SAM, when input token sequences length varies during testing, the magnitude of the Softmax outputs changes drastically. In contrast, we propose a new scaling factor to address this issue. (b) We introduce a bias-mode attention mask, which produces increasing penalties on attention scores as the distance between query and key grows.

where the relative ordering or distance of the elements matters. This method encodes the relative position between input elements x_i and x_j into vectors $p_{i,j}^v, p_{i,j}^q, p_{i,j}^k \in \mathbb{R}^{d_z}$. Then, we reformulate Eq. (1) as follows:

$$\mathbf{O}_i = \sum_{j=1}^n a_{i,j} (x_j v_j + p_{i,j}^v), \quad (2)$$

$$a_{i,j} = \frac{e^{\lambda(q_i + p_{i,j}^q) \cdot (k_j + p_{i,j}^k)}}{\sum_{j=1}^n e^{\lambda(q_i + p_{i,j}^q) \cdot (k_j + p_{i,j}^k)}}, \quad (3)$$

where $p_{i,j}^v, p_{i,j}^q, p_{i,j}^k$ is learned during training.

4. Methodology

From the above section, it is evident that the original SAM's design has a substantial limit to handling the length extrapolation problem. To address this, as shown in Figure 2, we present a Scalable Bias-mode Attention Mask (BA-SAM). In Sec. 4.1, we provide a theoretical explanation for the scaling factor used in the original Transformer and introduce a new scaling factor to regulate the magnitude inconsistency caused by length extrapolation. In Sec. 4.2 we design a bias-mode attention mask to place more focus on neighboring tokens, ensuring the sensitivity of the model when meeting positional encoding variations. Finally, we explain how we will embed our proposed BA-SAM into the SAM-based structure in Sec. 4.3.

4.1. New Scaling Factor

We discovered from the original attention module in SAM [62] that when the dot product becomes significantly large in magnitude, it can drive the Softmax layer into regions with minimal gradients. This is because the attention distribution after Softmax becomes highly concentrated, giving rise to the issue of vanishing gradients. Upon

closer examination of Eq. (1), it is obvious that the computation of the $q \cdot k$ term is intrinsically tied to both the token sequences length N and the dimension d_k . When the token sequence length N and/or the dimension d_k significantly increases, the overall efficacy of the attention is impacted, thus leading to a noticeable performance reduction.

To address this, we attempt to design a new scaling factor that allows the model to cope with significant variations in N and d_k values. When N or d_k grows significantly, we expect to regulate the magnitude of the values within the attention layer, maintaining a similar magnitude. While [62] introduced a scaling factor $\lambda = \frac{1}{\sqrt{d_k}}$ to counteract the effect of the large growth in magnitude due to the dot products, the authors did not explain how this scaling factor was designed in the original paper. Below we will provide a theoretical derivation of this scaling factor, and then elaborate on our proposed new scaling factor.

The dimension d_k . We assume that the components of q and k are independent Gaussian distribution variables with mean 0 and variance 1, then we want to prove that the $q \cdot k$ has mean 0 and variance d_k . The mean of $q \cdot k$ is:

$$\mathbb{E}[q \cdot k] = \mathbb{E} \left[\sum_{i=1}^{d_k} q_i k_i \right] = \sum_{i=1}^{d_k} \mathbb{E}[q_i] \mathbb{E}[k_i] = 0 \quad (4)$$

Similarly, we formulate the variance of $q \cdot k$ as follows:

$$\text{var}[q \cdot k] = \text{var} \left[\sum_{i=1}^{d_k} q_i k_i \right] = \sum_{i=1}^{d_k} \text{var}[q_i] \text{var}[k_i] = d_k \quad (5)$$

Given this, we can approximately consider the $q \cdot k$ values to be within the range of $-3\sqrt{d_k}$ to $3\sqrt{d_k}$, according to properties of Gaussian distribution. For larger models, d_k is generally a larger positive value, resulting in a significant

increase in the magnitude of numerical values of $q \cdot k$, compared to the additive attention option which has the range of $[-3, 3]$. Consequently, the attention distribution after Softmax becomes highly concentrated. This leads to severe gradient vanishing, which in turn hampers the effectiveness of training and can lead to less desired performance. As the $q \cdot k$ values lie in the range of $[-3\sqrt{d_k}, 3\sqrt{d_k}]$, the scaling factor can be simply defined as $\lambda_d = \frac{1}{\sqrt{d_k}}$, in order to maintain a similar magnitude.

Our new scaling factor. We provided the interpretation on how the original scaling factor was designed, which has been overlooked by the original authors. Now we explain the design of our new scaling factor.

Our goal is to regulate the magnitude of the Softmax value a_{ij} (Eq. (1)) as the token sequences length N and the dimension d_k change. We further denote the token sequences length during the training as N_{train} and the token sequence length during the testing as N_{test} , where $N_{test} \gg N_{train}$. We now consider how to maintain a similar magnitude of a_{ij} during training and testing times. We introduce the scaling factor λ_N , which is related to N and controls the output scale of the Softmax value a_{ij} , namely, $S(a_{i,j}^{N_{test}}) = S(a_{i,j}^{N_{train}})$. Here, $S()$ represents the scale estimation function.

We next substitute Eq. (1) into $S(a_{i,j}^{N_{test}}) = S(a_{i,j}^{N_{train}})$, and simplify $q_i \cdot k_j$ into $x_{i,j}$ provided the magnitude of $q \cdot k$ has been already maintained. Now we have:

$$\begin{aligned} S\left(\frac{e^{\lambda_n x_{i,j}}}{\sum_{j=1}^n e^{\lambda_n x_{i,j}}}\right) &= S\left(\frac{e^{x_{i,j}}}{\sum_{j=1}^m e^{x_{i,j}}}\right) \\ \Rightarrow S(e^{\lambda_n}) &= S\left(\frac{N_{test}}{N_{train}}\right) \Rightarrow \lambda_n = \log_{N_{train}} N_{test} \end{aligned} \quad (6)$$

Considering both λ_d and λ_n , we can ultimately derive our new scaling factor as:

$$\lambda = \lambda_d \lambda_n = \frac{\log_{N_{train}} N_{test}}{\sqrt{d_k}} \quad (7)$$

Our new scaling factor in Eq. (7) ensures the magnitude of values during attention computation remains consistent, regardless of variations in d_k and N . It will enhance the extrapolative capacity of the model.

4.2. Bias-Mode Attention Mask

Another challenge is that changes in token sequence length will lead to variations in positional encoding. It is important to ensure the insensitivity of the model when such positional encoding variations occur during testing.

One possible way is absolute encoding without trainable parameters, such as Sinusoidal [62]. It requires the position encoding to have strong local-to-global inference capabilities. As mentioned by Su [60], this resembles Taylor series approximations: by knowing the values of several

derivatives at a particular point, effective estimates can be made within a neighborhood. Nevertheless, this assumes that the given function has high-order smoothness (higher-order derivatives exist and are bounded). Commonly used positional encodings are often combined with trigonometric functions, essentially high-frequency oscillatory functions with periodicity. However, these methods fail to satisfy the requirement of bounded high-order derivatives, making it less accurate to estimate the extrapolated results.

Another potential approach is using local attention, which constrains the model’s field of view and remains insensitive to variations in token sequence length. However, local attention is typically implemented using a local window, necessitating modifications to the SAM structure, which requires re-training from scratch and is unacceptable for deployment in the downstream tasks.

To this end, we propose to enable the attention layer to put more focus on neighboring tokens of the current token. As such, even with an increase in the length of a token sequence, each token is scarcely affected by the untrained tokens from distant positions. In particular, we design a simple yet effective bias-mode mask. We introduce a bias after the query-key dot product.

As shown in Figure 3, this mask exhibits a bias specified on the distance between the query-key pairs (i.e., $q \cdot k$). We expect that this proposed mask imposes penalties on attention scores between distant query-key pairs, and the penalty increases as the distance between a key q and a query k grows. To this end, we simply define the bias as the form of $b_{i,j} = \beta|i - j|$.

$$a_{i,j} = \frac{e^{\lambda(q_i \cdot k_j + b_{i,j})}}{\sum_{j=1}^n e^{\lambda(q_i \cdot k_j + b_{i,j})}}, \quad (8)$$

where β is a head-specific slope.

We further discuss the setting of β based on different cases. When conducting zero-shot transfer learning without fine-tuning, we set β to a static, non-learned fixed value. The specific value setting will be discussed in the experimental section (Sec. 5). When fine-tuning is required, we make β trainable. Since our Bias-Mode Attention Mask is lightweight relative to the model structure, it incurs negligible training time overhead.

4.3. BA-SAM Model

As shown in Figure 3, our BA-SAM is simple to implement, and can be seamlessly integrated into SAM [38] and its variants. Specifically, our design involves a new scaling factor (NSF) for the attention layer and a bias-mode attention mask (BM-AM). Our method does not involve any alterations to the model structure and is suited to both fine-tuning and non-fine-tuning cases. As for the fine-tuning case, it introduces negligible computational overhead as BM-AM incurs a very small amount of computation.




Figure 3. Embedding of our BA-SAM into a SAM backbone. NSF indicates our new scaling factor, and BM-AM denotes our designed bias-mode attention mask.

5. Experiments

We first describe the datasets and implementation details in Sec. 5.1. Then, we compare BA-SAM’s performance with state-of-the-art methods across multiple tasks in Sec. 5.2. Finally, extensive ablation studies and analysis are conducted in Sec. 5.3. to demonstrate the effectiveness.

5.1. Datasets and Implementations

Datasets. For a comprehensive evaluation of BA-SAM, we conduct extensive experiments on a wide range of segmentation tasks, *i.e.*, salient object segmentation [15, 18, 20], complex object segmentation [19], skin lesion segmentation [93], camouflaged object detection [17], which correspond to four datasets: DUTS [64], DIS-TE4 [53], ISIC [12] and COD10K [16]. Besides, we also verify it on the challenging COCO [45] instance segmentation benchmark. *More details about datasets can be referred to in the supplementary material.*

Implementation Details. In zero-shot settings that do not require fine-tuning, we use the original SAM [38] backbone. For fine-tuning scenarios, we employ MobileSAM [76] as the backbone. MobileSAM [76] is a SAM variant with a structure similar to ViT-Tiny [72], and further details can be found in [76]. MobileSAM [76] uses a ViT-H-based SAM as a teacher network for distillation, ultimately achieving competitive accuracy compared to the original SAM but with significantly fewer parameters. For various object segmentation tasks, a random point is extracted from the ground truth as the prompt input during the fine-tuning phase. For instance segmentation, we use the ViT-B [14] backbone and the state-of-the-art detector Deformable-DETR [91] trained on the COCO [45] dataset with Swin-L [47] backbone as box prompt generator. *More details are provided in the supplementary material.* The code will be released at this [link](#).

Evaluation metrics. In the experiments, we use the widely used Mean Absolute Error (MAE) and Average Precision (AP) for evaluation. A lower MAE score and a higher AP score indicate better model performance.

5.2. Results

Results of Various Object Segmentation Tasks: Table 1 demonstrates the effectiveness of our approach across four diverse segmentation datasets. Δ_{diff} denotes the value of the performance degradation due to resolution changes during the inference. The upper and lower parts of the table indicate the results without and with fine-tuning, respectively due to input resolution variations. The best MAE performance is highlighted in bold, and the smallest degradation is underlined. We have three observations: Firstly, our proposed BA-SAM consistently outperforms both SAM [38] and MobileSAM [76] baselines on all four datasets. This is mainly because these baselines do not consider the issue of varying image resolutions. In contrast, our presented scaling factor and bias-mode attention mask explicitly handle this issue and further alleviate the performance degradation. Secondly, when testing on higher resolutions than the training size, SAM [38] and MobileSAM [76] baselines show less desirable results than the original image size. In contrast, our BA-SAM incurs significantly less performance drop in different datasets. Thirdly, during the experiment, we observe negligible computational overhead, whether fine-tuning is applied or not, which supports the claim in the method section. See Sec. 5.3 for details.

Results of Instance Segmentation: In Table 3, we evaluate the performance of our method on the COCO [45] instance segmentation benchmark. For a fair comparison, all experiments are conducted in a zero-shot manner, with the same initialization parameters for the comparative methods and without the use of any additional training data. Our BA-SAM consistently outperforms SAM [38] and MobileSAM [76] baselines, demonstrating better zero-shot generalization capability on instance segmentation.

Comparisons with State-of-the-Art Methods: To further demonstrate the superiority of the effectiveness and generalizability of our method, we compare the state-of-the-art approaches in Table 2. From the table, we have two following observations: Firstly, all the state-of-the-art approaches [16, 37, 46, 53, 57, 77, 81, 92] show less-desirable performances in each dataset. Instead, our BA-SAM (specialized models) consistently outperforms these methods when fine-tuned on each downstream dataset. Secondly, almost all of these state-of-the-art techniques are specifically designed for one task and cannot be generalized well to other tasks. Due to the strong zero-shot generalization capability of SAM [38], our proposed BA-SAM can also be employed as a generalized model, which fine-tunes with all of these downstream datasets in a unified and shared model. Importantly, unlike [12, 46, 57, 77], we eliminate the need for employing additional techniques to further enhance the performance. As shown in Table 2, our generalized model also consistently promotes the performance of SAM on all datasets, demonstrating its remarkable generalizability.

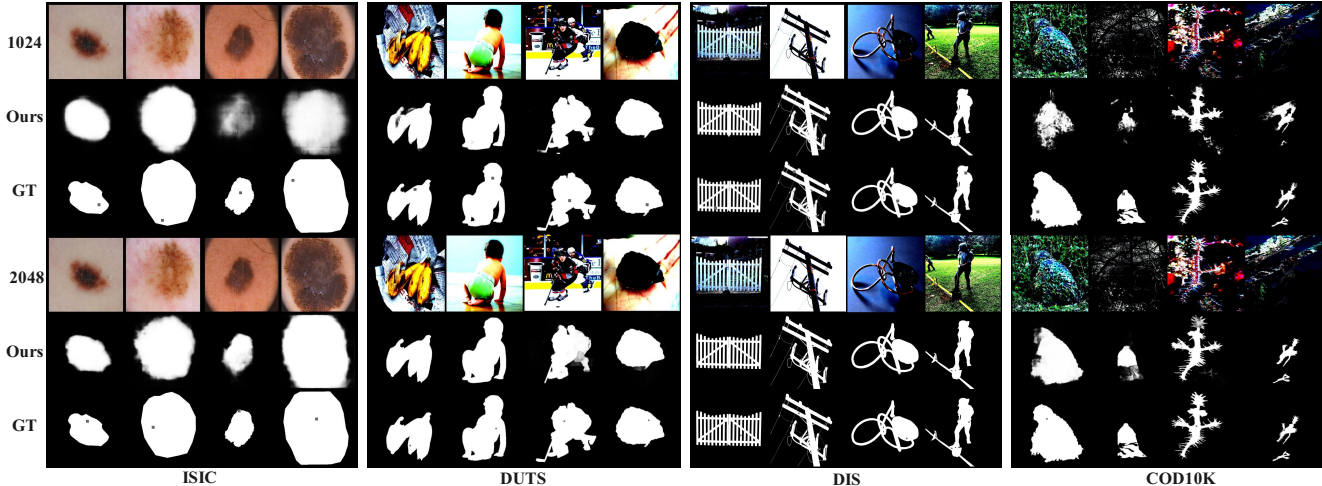


Figure 4. Visualization results of our BA-SAM on four object segmentation tasks, *i.e.*, skin lesion segmentation, salient object segmentation, complex object segmentation, camouflaged object detection, which corresponds to four datasets: ISIC [12], DUTS [64], DIS-TE4 [53], and COD10K [16]. Our BA-SAM can handle the issue of varying image resolutions and segments accurately in different tasks.

Method	Train Size	Test Size	ISIC [12]	$\Delta diff$	DUTS [64]	$\Delta diff$	DIS-TE4 [53]	$\Delta diff$	COD10K [16]	$\Delta diff$
Without fine-tuning										
SAM [38]	-	1024	0.421	-	0.298	-	0.362	-	0.217	-
	-	2048	0.601	18.0%	0.360	6.2%	0.411	4.9%	0.391	17.4%
Ours (w [38])	-	1024	0.417	-	0.294	-	0.356	-	0.208	-
	-	2048	0.589	17.2%	0.348	5.4%	0.406	5.0%	0.387	17.9%
MobileSAM [76]	-	1024	0.463	-	0.502	-	0.544	-	0.465	-
	-	2048	0.641	17.8%	0.437	6.5%	0.427	11.7%	0.346	11.9%
	-	4096	0.693	23.0%	0.328	17.4%	0.355	18.9%	0.300	16.5%
Ours (w [76])	-	1024	0.452	-	0.486	-	0.515	-	0.440	-
	-	2048	0.611	15.9%	0.413	7.3%	0.406	10.9%	0.321	11.9%
	-	4096	0.657	20.5%	0.283	20.3%	0.361	15.4%	0.246	19.4%
With fine-tuning										
MobileSAM [76]	1024	1024	0.098	-	0.034	-	0.057	-	0.030	-
		2048	0.520	42.2%	0.082	4.8%	0.065	0.8%	0.063	3.3%
		4096	0.542	44.4%	0.235	20.1%	0.101	4.4%	0.138	10.8%
	2048	2048	0.083	-	0.045	-	0.056	-	0.036	-
		4096	0.227	14.4%	0.091	4.6%	0.066	1.0%	0.059	2.3%
Ours (w [76])	1024	1024	0.037	-	0.031	-	0.055	-	0.029	-
		2048	0.499	40.4%	0.078	4.4%	0.061	0.6%	0.057	2.5%
		4096	0.531	43.6%	0.226	19.2%	0.098	4.3%	0.118	8.6%
	2048	2048	0.080	-	0.043	-	0.053	-	0.033	-
		4096	0.214	13.4%	0.088	4.4%	0.061	0.8%	0.056	2.3%

Table 1. Performance comparisons in varying image resolutions. We employed the widely-used MAE (Mean Absolute Error) score. Lower MAE scores indicate better model accuracy. $\Delta diff$ denotes the performance degradation due to resolution changes. Compared to the SAM [38] and MobileSAM [76] baselines, our BA-SAM achieves smaller degradation when encountering token sequences length changes. The best MAE performance is highlighted in bold, and the smallest performance degradation is underlined.

5.3. Ablation Study and Analysis

In this section, we first conduct ablation studies to study the contribution of each component. Then, we investigate the impact of the new scaling factor (NSF) and bias-mode attention mask (BM-AM) with a more detailed analysis.

Ablations Studies of Each Component. Table 4 sum-

marizes the effect of each designed component on the settings with and without fine-tuning, respectively. The baseline means using the MobileSAM [76] as the base network that uses the vanilla scaling factor (VSF) in the attention layer [62]. New Scaling Factor and Bias-Mode Attention Mask are abbreviated as NSF and BM-AM, respectively.

Methods	ISIC [12]	DUTS [64]	DIS-TE4 [53]	COD10K [16]
Specialized models				
CASS [57]	0.086	-	-	-
DINO [77]	0.081	-	-	-
MSA [71]	0.049	-	-	-
VST ₂₁ [46]	-	0.037	-	-
ICONet ₂₂ [92]	-	0.037	-	-
Gate ₂₀ [81]	-	-	0.109	-
IS-Net [53]	-	-	0.072	-
SINet [16]	-	-	-	0.092
SegMaR ₂₂ [37]	-	-	-	0.034
the same framework, 4 specialized models				
Ours (BA-SAM)	0.037	0.031	0.055	0.029
generalized model				
SAM [38]	0.419	0.298	0.373	0.217
Ours (BA-SAM)	0.054	0.030	0.054	0.054

Table 2. Comparison results (MAE) with state-of-the-art specialized models on various segmentation tasks.

Model	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
SAM [38]	42.5	69.6	44.7	29.7	47.0	56.7
Ours (w [38])	43.0	70.0	45.4	30.0	47.6	57.4
MobileSAM [76]	40.8	68.4	41.6	26.0	44.4	57.6
Ours (w [76])	41.6	69.0	42.3	26.2	44.8	58.4

Table 3. Results (AP) on COCO [45] instance segmentation.

From the table, we observe that NSF could achieve superior performances than the baseline with VSF. This is because the vanilla attention in SAM [38] and MobileSAM [76] does not consider maintaining the magnitude consistency of the Softmax outputs changes drastically when input resolutions vary during testing. In contrast, our NSF explicitly maintains the magnitude consistency and alleviates the performance degradation. Furthermore, by adding BM-AM, the performance could be further boosted when extrapolating to a larger test length. These improvements confirm that these individual components are complementary and together they significantly promote the performance.

Impact of Slope in Bias-Mode Attention Mask. In the Bias-Mode Attention Mask, the magnitude of the slope β determines penalty rates in different heads. In the zero-shot case, we use a fixed slope. Table 5 reveals the impact of different slopes of BM-AM on the ISIC [12] dataset. We found that the best performance is achieved when $\beta = 0.1$. Besides, our method is robust to different slope choices. In addition to the zero-shot setting, we also validate the BM-AM in fine-tuning scenarios on four datasets, shown in Table 1. The learnable parameters brought more performance growth compared with the fixed slope, while also bringing a slight increase in computation.

Computational Efficiency. In Table 6, we analyze the computational efficiency between the baselines and our BA-SAM. All the experiments are conducted on the same NVIDIA RTX 4090GPU to ensure fair comparisons. From the table, we observe that our BA-SAM is highly

Methods	ISIC [12]	DUTS [64]	DIS-TE4 [53]	COD10K [16]
Without fine-tuning				
Baseline [76]	17.8	6.5	11.7	11.9
+ NSF	16.4	7.5	11.3	12.0
+ BM-AM	16.8	7.2	11.7	11.9
+ Both	15.9	7.9	10.9	11.9
With fine-tuning (1024)				
Baseline [76]	42.2	4.8	0.8	3.3
+ NSF	40.9	4.6	0.7	3.0
+ BM-AM	41.2	4.5	0.8	2.7
+ Both	40.4	4.4	0.6	2.5
With fine-tuning (2048)				
Baseline [76]	14.4	4.6	1.0	2.3
+ NSF	13.1	4.6	0.8	2.4
+ BM-AM	13.7	4.5	0.9	2.3
+ Both	13.4	4.4	0.8	2.3

Table 4. Ablation study of each component on the settings with and without fine-tuning. Numbers indicate the performance degradation, $\Delta diff$. A lower $\Delta diff$ means a better performance.

Size	0.1	0.5	1	1.5	2	2.5
MAE	0.0375	0.0383	0.0381	0.0393	0.0388	0.0391
$\Delta diff$	40.4	40.8	40.4	40.9	41.2	40.7

Table 5. Impact of the slope β in BM-AM (zero-shot setting).

Model	Params (M)	Speed (ms)	Train Hours (h)
SAM [38]	81	113.9	-
Ours (w [38])	81	114.0	-
MobileSAM [76]	9.66	16.2	0.64
Ours (w [76])	9.67	16.5	0.65

Table 6. Comparisons of computational efficiency between the baselines and our BA-SAM. Params: number of parameters. Speed: inference speed. The top part is performed on the zero-shot setting, and the bottom part is the scenarios with fine-tuning.

lightweight, incurring negligible computational overhead for the models. The reasons lie in two aspects: firstly, the NSF exhibits nearly identical computational complexity with the vanilla one. Besides, the BM-AM is seamlessly incorporated by adding a mask matrix to the query-key dot product before applying the softmax operation. Although there is a slight increase in memory usage, it remains negligible compared to the memory occupied by large models.

6. Conclusion

In this paper, we address the important problem of varying image resolutions in SAM models by reformulating it as a problem of length extrapolation. To enhance the length extrapolation capability of SAM, we propose Scalable Bias-mode Attention Mask (BA-SAM). A new scaling factor (NSF) is introduced to maintain the consistent magnitude of attention. In addition, a bias-mode attention mask (BM-AM) is designed to prioritize neighboring information, mit-

igating the impact of untrained distant information. Extensive evaluation on diverse datasets reveals its ability to significantly alleviate performance degradation in the zero-shot setting and achieve state-of-the-art performance with minimal fine-tuning. Furthermore, we proposed a generalized model and benchmark, showcasing BA-SAM’s generalizability across all four datasets simultaneously.

References

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014. 3
- [2] Irwan Bello. Lambdanetworks: Modeling long-range interactions without attention. In *CVPR*, 2021. 3
- [3] Irwan Bello, Barret Zoph, Quoc Le, Ashish Vaswani, and Jonathon Shlens. Attention augmented convolutional networks. In *ICCV*, 2019. 3
- [4] Lucas Beyer, Pavel Izmailov, Alexander Kolesnikov, Mathilde Caron, Simon Kornblith, Xiaohua Zhai, Matthias Minderer, Michael Tschannen, Ibrahim Alabdulmohsin, and Filip Pavetic. Flexivit: One model for all patch sizes. In *CVPR*, 2023. 1, 2, 3
- [5] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 2
- [6] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once for all: Train one network and specialize it for efficient deployment. *ICLR*, 2020. 3
- [7] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *NeurIPS*, 2022. 3
- [8] Tianrun Chen, Lanyun Zhu, Chaotao Ding, Runlong Cao, Shangzhan Zhang, Yan Wang, Zejian Li, Lingyun Sun, Papa Mao, and Ying Zang. Sam fails to segment anything?-sam-adapter: Adapting sam in underperformed scenes: Camouflage, shadow, and more. In *arXiv preprint arXiv:2304.09148*, 2023. 2
- [9] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, 2021. 2
- [10] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *CVPR*, 2020. 3
- [11] Ta-Chung Chi, Ting-Han Fan, Peter J Ramadge, and Alexander Rudnicky. Kerple: Kernelized relative positional embedding for length extrapolation. *NeurIPS*, 35, 2022. 3
- [12] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *ISBI*, 2018. 2, 6, 7, 8
- [13] Mostafa Dehghani, Basil Mustafa, Josip Djolonga, Jonathan Heek, Matthias Minderer, Mathilde Caron, Andreas Steiner, Joan Puigcerver, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Patch n’pack: Navit, a vision transformer for any aspect ratio and resolution. *arXiv preprint arXiv:2307.06304*, 2023. 3
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 6
- [15] Deng-Ping Fan, Ming-Ming Cheng, Jiang-Jiang Liu, Shang-Hua Gao, Qibin Hou, and Ali Borji. Salient objects in clutter: Bringing salient object detection to the foreground. In *ECCV*, 2018. 6
- [16] Deng-Ping Fan, Ge-Peng Ji, Guolei Sun, Ming-Ming Cheng, Jianbing Shen, and Ling Shao. Camouflaged object detection. In *CVPR*, 2020. 2, 6, 7, 8
- [17] Deng-Ping Fan, Ge-Peng Ji, Ming-Ming Cheng, and Ling Shao. Concealed object detection. *TPAMI*, 2021. 6
- [18] Deng-Ping Fan, Zheng Lin, Zhao Zhang, Menglong Zhu, and Ming-Ming Cheng. Rethinking rgb-d salient object detection: Models, data sets, and large-scale benchmarks. *TNNLS*, 2021. 6
- [19] Deng-Ping Fan, Tengpeng Li, Zheng Lin, Ge-Peng Ji, Dingwen Zhang, Ming-Ming Cheng, Huazhu Fu, and Jianbing Shen. Re-thinking co-salient object detection. *TPAMI*, 2022. 6
- [20] Deng-Ping Fan, Jing Zhang, Gang Xu, Ming-Ming Cheng, and Ling Shao. Salient objects in clutter. *TPAMI*, 2022. 6
- [21] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *CVPR*, 2023. 2
- [22] Zhengyang Feng, Qianyu Zhou, Qiqi Gu, Xin Tan, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Dmt: Dynamic mutual training for semi-supervised learning. *Pattern Recognition*, 2022. 2
- [23] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *ICML*, 2017. 3
- [24] Chengyue Gong, Dilin Wang, Meng Li, Xinlei Chen, Zhicheng Yan, Yuandong Tian, Vikas Chandra, et al. Nasvit: Neural architecture search for efficient vision transformers with gradient conflict aware supernet training. In *ICLR*, 2021. 3
- [25] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *ICCV*, 2021. 3
- [26] Qiqi Gu, Qianyu Zhou, Minghao Xu, Zhengyang Feng, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Pit: Position-invariant transform for cross-fov domain adaptation. In *ICCV*, 2021. 2
- [27] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 3

- [28] Junxian He, Chunting Zhou, Xueze Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *ICLR*, 2021. 3
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2
- [30] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollar, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 3
- [31] Lu He, Qianyu Zhou, Xiangtai Li, Li Niu, Guangliang Cheng, Xiao Li, Wenxuan Liu, Yunhai Tong, Lizhuang Ma, and Liqing Zhang. End-to-end video object detection with spatial-temporal transformers. In *ACM MM*, 2021. 2
- [32] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 3
- [33] Ge-Peng Ji, Deng-Ping Fan, Peng Xu, Ming-Ming Cheng, Bowen Zhou, and Luc Van Gool. Sam struggles in concealed scenes—empirical study on “segment anything”. *arXiv preprint arXiv:2304.06022*, 2023. 3
- [34] Wei Ji, Jingjing Li, Qi Bi, Wenbo Li, and Li Cheng. Segment anything is not always perfect: An investigation of sam on different real-world applications. *arXiv preprint arXiv:2304.05750*, 2023. 2
- [35] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 2
- [36] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022. 3
- [37] Qi Jia, Shuilian Yao, Yu Liu, Xin Fan, Risheng Liu, and Zhongxuan Luo. Segment, magnify and reiterate: Detecting camouflaged objects the hard way. In *CVPR*, 2022. 6, 8
- [38] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 2, 3, 5, 6, 7, 8
- [39] Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, and Ali Farhadi. Matryoshka representations for adaptive deployment. *arXiv preprint arXiv:2205.13147*, 2022. 3
- [40] Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. Pix2struct: Screenshot parsing as pretraining for visual language understanding. In *ICML*, 2023. 3
- [41] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *CVPR*, 2019. 3
- [42] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL*, 2021. 3
- [43] Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations and recent trends in multimodal machine learning: Principles, challenges, and open questions. *arXiv preprint arXiv:2209.03430*, 2022. 2
- [44] Mingbao Lin, Mengzhao Chen, Yuxin Zhang, Ke Li, Yunhang Shen, Chunhua Shen, and Rongrong Ji. Super vision transformer. *IJCV*, 2022. 3
- [45] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 3, 6, 8
- [46] Nian Liu, Ni Zhang, Kaiyuan Wan, Ling Shao, and Junwei Han. Visual saliency transformer. In *ICCV*, 2021. 6, 8
- [47] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 6
- [48] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *CVPR*, 2022. 3
- [49] Shaocong Long, Qianyu Zhou, Chenhao Ying, Lizhuang Ma, and Yuan Luo. Diverse target and contribution scheduling for domain generalization. *arXiv preprint arXiv:2309.16460*, 2023. 2
- [50] Shaocong Long, Qianyu Zhou, Chenhao Ying, Lizhuang Ma, and Yuan Luo. Rethinking domain generalization: Discriminability and generalizability. *arXiv preprint arXiv:2309.16483*, 2023. 2
- [51] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, 2018. 3
- [52] Ofir Press, Noah A Smith, Mike Lewis, et al. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2021. 3
- [53] Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, and Luc Van Gool. Highly accurate dichotomous image segmentation. In *ECCV*, 2022. 2, 6, 7, 8
- [54] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2
- [55] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [56] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *Proceedings of NAACL-HLT*, 2018. 3
- [57] et al Singh, Pranav. Cass: cross architectural self-supervision for medical image analysis. *arXiv preprint arXiv:2206.04170*, 2022. 6, 8
- [58] Yiran Song, Qianyu Zhou, and Lizhuang Ma. Rethinking implicit neural representations for vision learners. In *ICASSP*, 2023. 2
- [59] Aravind Srinivas, Tsung-Yi Lin, Niki Parmar, Jonathon Shlens, Pieter Abbeel, and Ashish Vaswani. Bottleneck transformers for visual recognition. In *CVPR*, 2021. 3
- [60] JianLin Su. Transformer upgrade road: 7, length extrapolation and local attention, 2023. 5

- [61] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 3
- [62] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 3, 4, 5, 7
- [63] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, 2017. 3
- [64] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *CVPR*, 2017. 2, 6, 7, 8
- [65] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021. 3
- [66] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 2, 3
- [67] Xiao Wang, Guangyao Chen, Guangwu Qian, Pengcheng Gao, Xiao-Yong Wei, Yaowei Wang, Yonghong Tian, and Wen Gao. Large-scale multi-modal pre-trained models: A comprehensive survey. *MIR*, 2023. 2
- [68] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *ECCV*, 2018. 3
- [69] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph Gonzalez, Kurt Keutzer, and Peter Vajda. Generating long sequences with sparse transformers. *arXiv preprint arXiv:2006.03677*, 2020. 3
- [70] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph Gonzalez, Kurt Keutzer, and Peter Vajda. Visual transformers: Token-based image representation and processing for computer vision. *arXiv preprint arXiv:2006.03677*, 2020. 3
- [71] Junde Wu, Rao Fu, Huihui Fang, Yuanpei Liu, Zhaowei Wang, Yanwu Xu, Yueming Jin, and Tal Arbel. Medical sam adapter: Adapting segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.12620*, 2023. 8
- [72] Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Tinyvit: Fast pretraining distillation for small vision transformers. In *ECCV*, 2022. 6
- [73] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, 2022. 3
- [74] Hongyi Xu, Fengqi Liu, Qianyu Zhou, Jinkun Hao, Zhijie Cao, Zhengyang Feng, and Lizhuang Ma. Semi-supervised 3d object detection via adaptive pseudo-labeling. In *ICIP*, 2021. 2
- [75] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xi-aodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In *ECCV*, 2020. 3
- [76] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. 2, 3, 6, 7, 8
- [77] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 6, 8
- [78] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Haibin Lin, Zhi Zhang, Yue Sun, Tong He, Jonas Mueller, R Manmatha, et al. Resnest: Split-attention networks. In *CVPR*, 2022. 3
- [79] Kaidong Zhang and Dong Liu. Customized segment anything model for medical image segmentation. In *arXiv preprint arXiv:2304.13785*, 2023. 1, 2
- [80] Hengshuang Zhao, Jiaya Jia, and Vladlen Koltun. Exploring self-attention for image recognition. In *CVPR*, 2020. 3
- [81] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Suppress and balance: A simple gated network for salient object detection. In *ECCV*, 2020. 6, 8
- [82] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 2019. 3
- [83] Qianyu Zhou, Zhengyang Feng, Qiqi Gu, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Uncertainty-aware consistency regularization for cross-domain semantic segmentation. *CVIU*, 2022. 2
- [84] Qianyu Zhou, Zhengyang Feng, Qiqi Gu, Jiangmiao Pang, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Context-aware mixup for domain adaptive semantic segmentation. *TCSVT*, 2022.
- [85] Qianyu Zhou, Ke-Yue Zhang, Taiping Yao, Ran Yi, Shouhong Ding, and Lizhuang Ma. Adaptive mixture of experts learning for generalizable face anti-spoofing. In *ACM MM*, 2022.
- [86] Qianyu Zhou, Ke-Yue Zhang, Taiping Yao, Ran Yi, Kekai Sheng, Shouhong Ding, and Lizhuang Ma. Generative domain adaptation for face anti-spoofing. In *ECCV*, 2022.
- [87] Qianyu Zhou, Chuyun Zhuang, Ran Yi, Xuequan Lu, and Lizhuang Ma. Domain adaptive semantic segmentation via regional contrastive consistency regularization. In *ICME*, 2022.
- [88] Qianyu Zhou, Qiqi Gu, Jiangmiao Pang, Xuequan Lu, and Lizhuang Ma. Self-adversarial disentangling for specific domain adaptation. *TPAMI*, 2023.
- [89] Qianyu Zhou, Xiangtai Li, Lu He, Yibo Yang, Guangliang Cheng, Yunhai Tong, Lizhuang Ma, and Dacheng Tao. Transvod: end-to-end video object detection with spatial-temporal transformers. *TPAMI*, 2023.
- [90] Qianyu Zhou, Ke-Yue Zhang, Taiping Yao, Xuequan Lu, Ran Yi, Shouhong Ding, and Lizhuang Ma. Instance-aware domain generalization for face anti-spoofing. In *CVPR*, 2023. 2

- [91] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. 6
- [92] Mingchen Zhuge, Deng-Ping Fan, Nian Liu, Dingwen Zhang, Dong Xu, and Ling Shao. Salient object detection via integrity learning. *PAMI*, 2022. 6, 8
- [93] Mingchen Zhuge, Deng-Ping Fan, Nian Liu, Dingwen Zhang, Dong Xu, and Ling Shao. Salient object detection via integrity learning. *TPAMI*, 2023. 3, 6