

Frequency-Dynamic Attention Modulation for Dense Prediction

Linwei Chen¹

Lin Gu^{2,3}

Ying Fu^{1*}

¹Beijing Institute of Technology

²RIKEN AIP

³The University of Tokyo

chenlinwei@bit.edu.cn; lin.gu@riken.jp; fuying@bit.edu.cn

Abstract

Vision Transformers (ViTs) have significantly advanced computer vision, demonstrating strong performance across various tasks. However, the attention mechanism in ViTs makes each layer function as a low-pass filter, and the stacked-layer architecture in existing transformers suffers from frequency vanishing. This leads to the loss of critical details and textures. We propose a novel, circuit-theory-inspired strategy called Frequency-Dynamic Attention Modulation (FDAM), which can be easily plugged into ViTs. FDAM directly modulates the overall frequency response of ViTs and consists of two techniques: Attention Inversion (AttInv) and Frequency Dynamic Scaling (FreqScale). Since circuit theory uses low-pass filters as fundamental elements, we introduce AttInv, a method that generates complementary high-pass filtering by inverting the low-pass filter in the attention matrix, and dynamically combining the two. We further design FreqScale to weight different frequency components for fine-grained adjustments to the target response function. Through feature similarity analysis and effective rank evaluation, we demonstrate that our approach avoids representation collapse, leading to consistent performance improvements across various models, including SegFormer, DeiT, and MaskDINO. These improvements are evident in tasks such as semantic segmentation, object detection, and instance segmentation. Additionally, we apply our method to remote sensing detection, achieving state-of-the-art results in single-scale settings. The code is available at <https://github.com/Linwei-Chen/FDAM>.

1. Introduction

In recent years, Vision Transformers (ViTs) have revolutionized computer vision, achieving state-of-the-art performance across various dense prediction tasks [21, 47, 55, 92].

However, when applying ViTs [21] to vision tasks, a critical issue emerges. The attention mechanism in ViTs exhibits a strong low-pass filtering characteristic [61, 83]. As shown in Figure 1(a), frequency analysis reveals that




Figure 1. Frequency analysis. We stack a model with pure 12 attention layer. (a) Attention frequency response analysis reveals that our modulated attention maintains a higher mean magnitude across all frequency bands compared to standard attention, while also exhibiting greater diversity in high-frequency regions. (b) Feature spectrum analysis shows that our modulated attention maintains a stable high-frequency ratio and consistently preserves high-frequency information across layers, unlike standard attention, which rapidly loses it and results in representation collapse.

this low-pass filtering limits spectral diversity, severely restricting the frequency representation power. This limitation is exacerbated by the stacked-layer architecture. As illustrated in Figure 1(b), feature spectrum analysis shows that standard attention rapidly loses high-frequency information across layers, with the high-frequency ratio dropping sharply from the initial layers to the deeper layers (nearly 0). This leads to frequency vanishing and blurred feature representations, negatively impacting performance on tasks requiring fine-grained visual understanding.

To address this issue, we propose Frequency-Dynamic Attention Modulation (FDAM), a novel, circuit-theory-inspired strategy to modulate the overall frequency response of ViTs. FDAM consists of two techniques: Attention Inversion (AttInv) and Frequency Dynamic Scaling (FreqScale). Drawing on circuit theory [2, 73, 85], which treats low-pass filters as fundamental building blocks, we introduce our first technique, AttInv. In AttInv, we view attention matrix as a set of low-pass filters and invert them (analogous to s-domain inversion in circuit design) to derive complementary high-pass filters. At each layer, both low-pass and high-pass filters are dynamically weighted. By cascading these layers over L levels, the architecture forms 2^L unique combinations of filter weights, enabling it to learn complex frequency responses, as shown in Figure 2.

*Corresponding Author




Figure 2. Analysis of frequency response fitting. From the center to the border are low- to high-frequency components. It is evident that the attention mechanism has strong low-pass characteristics, which makes it difficult to effectively fit high-pass, band pass/stop, and random filters. In contrast, our method, AttInv, demonstrates a superior capability in fitting these diverse frequency responses, indicating greater flexibility and effectiveness in handling a wide range of frequency characteristics.

The second technique, termed FreqScale, operates in a frequency-dynamic manner. It addresses the limitation of AttInv by providing fine-grained adjustments to the target response function. While AttInv effectively combines low-pass and high-pass filters, it lacks precise control over individual frequency components. FreqScale re-weights feature maps across separate frequency bands and dynamically amplifies high-frequency signals. This enables a more nuanced and adaptive adjustment of feature representations, enhancing the model’s ability to distinguish between different categories for dense prediction tasks such as segmentation.

These techniques are computationally efficient and can be *easily plugged into* existing ViT architectures. By incorporating these methods, we address the low-pass limitations in ViTs across various vision tasks, enabling full-spectrum feature representation. This results in significant performance improvements in dense prediction tasks, such as semantic segmentation (SegFormer +2.4 mIoU on ADE20K), object detection (Mask DINO +1.6 AP on COCO), and instance segmentation (Mask DINO +1.4 AP on COCO). Additionally, our method achieves 78.61 when applied to remote sensing object detection, surpassing previous state-of-the-art methods under single-scale training and testing settings. Furthermore, feature similarity analysis [61] and effective rank evaluation [37] confirm that our approach effectively prevents representation collapse.

Our contributions can be summarized as follows:

- We diagnose the low-pass filtering characteristic of ViTs’ attention mechanism through mathematical and spectral analysis. It reveals how this characteristic restricts frequency representation and leads to feature degradation, providing a clear understanding of the challenges faced by ViTs in fine-grained visual tasks.
- We introduce Frequency-Dynamic Attention Modulation (FDAM), comprising two techniques: AttInv and FreqScale. AttInv leverages the low-pass filtering characteristic to create a complementary high-pass fil-

ter, enabling full-spectrum feature representation. FreqScale provides fine-grained frequency adjustments by re-weighting and amplifying frequency signals. Together, these techniques effectively address the limitations of traditional attention mechanisms.

- Our methods are computationally efficient and can be easily integrated into existing ViT architectures. They achieve significant performance gains across diverse vision tasks. Extensive analysis, including effective rank evaluation [37], demonstrates that our approach avoids representation collapse.

2. Related Work

Vision Transformer. Following the success of attention-based architectures in neural machine translation [80], computer vision researchers explored Vision Transformers (ViT) [21]. These models tokenize image patches and process them using sequential attention mechanisms, achieving strong performance on classification benchmarks.

In recent years, various ViT variants have been proposed [24, 31, 42, 55, 70, 75, 77, 78, 95, 109, 110]. Hybrid architectures, such as [35, 86], combine convolutional inductive biases with attention mechanisms. Hierarchical designs like [84] use progressive resolution reduction for efficient multi-scale processing. To address quadratic complexity, windowed attention mechanisms [31, 55] and multi-granularity interactions [95] have been introduced. MetaFormer [101] demonstrated that the success of transformer stems from the residual MetaFormer framework rather than specific attention operators. TransNeXt [70] explores alternative spatial interaction paradigms inspired by biomimetic vision.

Beyond image classification, ViT variants have inspired the application of Transformers to other vision tasks, such as object detection [47, 114], semantic segmentation [72, 92, 106], and instance segmentation [15, 41].

Frequency-Domain Learning. Frequency-domain analysis has served as a foundational pillar in signal processing for decades [25, 65]. Recent advances have extended these principles to deep learning, where they enhance model optimization strategies [100] and improve the generalization capacity of deep neural networks (DNNs) [12, 81].

The integration of frequency-domain techniques into DNN architectures has demonstrated several advantages, such as capturing global contextual patterns through spectral operations [17, 27, 36, 49, 67], strengthening domain-generalizable representations through frequency-aware learning [10, 43, 50]. Researcher also utilize frequency-domain techniques in neural operations. For instance, FcaNet [66] enhances feature recalibration through frequency component analysis, while FreqFusion [8] optimizes multi-scale feature fusion using spectral properties. Subsequent work has further resolved aliasing arti-




Figure 3. Illustration of Frequency Dynamic Attention Modulation (FDAM), comprising AttInv for attention modulation and FreqScale for feature modulation. The original attention mechanism is predominantly influenced by low-frequency components due to its inherent low-pass filtering characteristics. **i) AttInv** inverts the low-pass filter, represented by the attention weights, to derive a high-pass filter. By dynamically combining these filters using a predicted weight, we achieve a balanced representation that retains both low- and high-frequency information. **ii) FreqScale** adaptively reweights different frequency bands, enhancing suppressed high-frequency components (e.g., edges, textures) while preserving structural low-frequency information. This integrated approach alleviates attention collapse and patch uniformity issues in Vision Transformers, facilitating full-spectrum feature representation for improved discriminability.

facts in downsampling operations through high-frequency suppression [9, 11, 26], and FADC [13] adapts convolutional dilation rates based on feature frequency profiles. [1, 62–64, 74, 104] also explore improving ViTs from frequency perspective.

Anti-Oversmoothing for ViT. Recent works [20, 61, 62, 83] have conducted Fourier domain analyses of the oversmoothing phenomenon in ViTs, demonstrating that the self-attention mechanism acts as a low-pass filter. This causes feature maps to lose high-frequency information and converge to a Direct-Current (DC) component, leading to patch uniformity and rank collapse in deep ViTs [20].

Existing solutions, such as AttnScale and FeatScale [83], mitigate this effect by adaptively scaling high-frequency components, while NeuTRENO [61] incorporates a regularizer to preserve token fidelity. However, these methods primarily focus on static enhancement of high-frequency signals and neglect the broader spectral context. They fail to dynamically adapt to the varying frequency requirements of different layers and tasks, resulting in incomplete feature representations. Our method addresses these limitations by combining adaptive inverted high-pass filters with fine-grained frequency control. This enables the model to dynamically adjust its attention to capture a richer spectrum of features across layers, thereby capturing subtle visual differences that benefit complex vision tasks.

3. Frequency-Dynamic Attention Modulation

In this section, we introduce our Frequency-Dynamic Attention Modulation (FDAM) mechanism, an approach designed to address the limitations of the low-pass filtering characteristic in Vision Transformers (ViTs). As shown

in Figure 3, our method consists of two key techniques, AttInv and FreqScale, which work together to enhance the frequency representation capabilities of ViTs.

3.1. Attention Inversion

Motivation. The self-attention mechanism is pivotal in ViTs [21], enabling the model to capture long-range dependencies and weigh the importance of different input elements. This mechanism is mathematically represented as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{C}}\right) \mathbf{V}, \quad (1)$$

where $\mathbf{Q}$, $\mathbf{K}$, and $\mathbf{V}$ represent the query, key, and value matrices, respectively, each with a shape of (C, HW) . These matrices are derived from the input feature $\mathbf{X}$ through a linear transformation. Here, c denotes the channel dimension of the vectors, and H and W represent the height and width.

The attention matrix $\mathbf{A} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{C}}\right)$ can be interpreted as a set of $H \times W$ linear filters [83], where each spatial location (p, q) has a corresponding filter $\mathbf{A}_{p,q} \in \mathbb{R}^{H \times W}$. It can be understood and mathematically proven to act as a low-pass filter due to its smoothing effect on feature maps [61, 83], which can be formulated as follows (see *supplementary material* for proof):

$$\begin{aligned} |\mathcal{F}(\mathbf{A}_{p,q})(u, v)| &= 1, \text{ if } (u, v) = (0, 0), \\ |\mathcal{F}(\mathbf{A}_{p,q})(u, v)| &< 1, \text{ if } (u, v) \neq (0, 0). \end{aligned} \quad (2)$$

where $\mathcal{F}(\cdot)$ denotes discrete Fourier transform (DFT), and (u, v) represents the frequency components. $|\mathcal{F}(\mathbf{A}_{p,q})(0, 0)| = 1$ ensures that the lowest direct current frequency is preserved, while $|\mathcal{F}(\mathbf{A}_{p,q})(u, v)| < 1$ for $(u, v) \neq (0, 0)$ indicates that higher frequency components are attenuated, thus confirming the low-pass filtering

effect. This limits its spectrum diversity and severely restricts its frequency representation power.

Now, consider a simple model with L layers of pure self-attention. Let $\mathcal{F}(\mathbf{X}^{(i)})$ denote the Fourier transformed spectrum of the feature map at layer i , and let $\mathcal{F}(\mathbf{A}^{(i)})$ denote the frequency response of the attention matrix at the same layer. The transformation across layers follows the recursive relation:

$$\mathcal{F}(\mathbf{X}^{(L)})(u, v) = \prod_{i=1}^L \mathcal{F}(\mathbf{A}^{(i)})(u, v) \cdot \mathcal{F}(\mathbf{X}^{(0)})(u, v). \quad (3)$$

Since $|\mathcal{F}(\mathbf{A}^{(i)})(u, v)| < 1$ for all nonzero frequencies $(u, v) \neq (0, 0)$, we observe that:

$$\lim_{L \rightarrow \infty} \prod_{i=1}^L |\mathcal{F}(\mathbf{A}^{(i)})(u, v)| = 0, \quad \forall (u, v) \neq (0, 0). \quad (4)$$

This means that, all high-frequency components are exponentially suppressed with layers, leaving only the lowest frequency component $(0, 0)$ dominant. Consequently, the model suffers from frequency vanishing, where fine-grained details and textures are lost, impairing the model to capture crucial information for dense prediction vision tasks.

To address the frequency representation limitations of the attention mechanism, we propose Attention Inversion (AttInv). The core idea is inspired by circuit theory [73, 85], which uses low-pass filters as fundamental elements to construct various filters, such as high-pass and band-pass filters. AttInv leverages the inherent low-pass filtering characteristic of the attention mechanism and inverts it to obtain a complementary high-pass filter. We then dynamically recombine these two types of filters in a spatially adaptive manner. This process enables us to capture high-frequency information that is typically lost in standard attention.

Overview of AttInv. AttInv involves two main steps: 1) *Attention Inversion*: Inverting filters in attention matrix $\mathbf{A}_{p,q}$ in the frequency domain to derive a complementary high-pass filter $\hat{\mathbf{A}}_{p,q}$. 2) *Dynamic Combination*: Predicting a spatial dynamic coefficient $\bar{\mathbf{S}}, \hat{\mathbf{S}} \in \mathbb{R}^{H \times W}$ for each attention head to combine complementary filters $\mathbf{A}_{p,q}, \hat{\mathbf{A}}_{p,q}$.

Attention Inversion. To obtain a complementary high-pass filter by inverting the low-pass filter, AttInv first computes the frequency response of $\mathbf{A}$ and subtracts it from an all-pass filter $\mathbf{I}_f$. The resulting high-pass filter $\hat{\mathbf{A}}$ is then transformed back into the spatial domain:

$$\hat{\mathbf{A}}_{p,q} = \mathcal{F}^{-1}(\mathbf{I}_f - \mathcal{F}(\mathbf{A}_{p,q})), \quad (5)$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier Transform and its inverse, respectively. This ensures that $\hat{\mathbf{A}}_{p,q}$ captures high-frequency components complementary to $\mathbf{A}_{p,q}$.

Dynamic Combination. Since different locations on the feature map exhibit varying local frequencies, it is essential to adaptively extract different frequency components across regions. For instance, edges require high-frequency information for accurate boundary preservation, whereas smooth

regions do not. To achieve this, we employ a spatially dynamic approach to combine the low-pass filter $\mathbf{A}$ and the high-pass filter $\hat{\mathbf{A}}$:

$$\tilde{\mathbf{A}}_{p,q} = \bar{\mathbf{S}}(p, q) \cdot \mathbf{A}_{p,q} + \hat{\mathbf{S}}(p, q) \cdot \hat{\mathbf{A}}_{p,q}, \quad (6)$$

where $\bar{\mathbf{S}}, \hat{\mathbf{S}}$ represent the combination weights obtained via a convolutional layer. This spatially adaptive combination dynamically balances low- and high-pass filtering, ensuring that each region retains the most relevant frequency components. When cascading L such layers, the frequency response across L layers can be expressed as:

$$\mathcal{F}(\mathbf{X}^{(L)}) = \prod_{i=1}^L [\bar{\mathbf{S}}^{(i)} \mathcal{F}(\mathbf{A}^{(i)}) + \hat{\mathbf{S}}^{(i)} \mathcal{F}(\hat{\mathbf{A}}^{(i)})] \cdot \mathcal{F}(\mathbf{X}^{(0)}). \quad (7)$$

This recursive composition of these hybrid filters expands into 2^L distinct weighted combinations of low- and high-pass operations, enabling flexible amplification or suppression of specific frequency bands. In contrast, stacking L standard attention layers monotonically attenuates high frequencies, leading to exponential vanishing as Equation (4) described. AttInv's quadratic complexity in frequency operations preserves both global structures (via low-pass) and fine details (via high-pass), overcoming the spectral limitations of conventional attention, as shown in Figure 1.

3.2. Frequency Dynamic Scaling

Motivation. AttInv effectively combines low-pass and high-pass filters but lacks precise control over individual frequency components. To address this, we propose Frequency Dynamic Scaling (FreqScale), which provides fine-grained adjustments to the target response function by reweighting feature maps across separate frequency bands and dynamically amplifying high-frequency signals.

Overview of FreqScale. FreqScale involves two main steps: 1) *Frequency Scaling Weight Generation*: Compute dynamic frequency band weights using an multi-layer perceptron (MLP) and combine them with learnable scaling weights to obtain final frequency scaling coefficients. 2) *Feature Frequency Modulation*: Transform features into the frequency domain, divide the spectrum into multiple bands, and modulate each band with frequency scaling weight.

Frequency Scaling Weight Generation. A straightforward approach to adjust frequency components within the feature map is to use static learnable parameters for each band. However, this static approach is suboptimal given the dynamic nature of attention mechanisms.

To address this, we propose a simple yet efficient method to generate dynamic frequency scaling weights based on the input. Specifically, as illustrated in Figure 4, we employ n learnable static scaling weights and use a multi-layer perceptron (MLP) with a tanh activation function to compute dynamic coefficients. These coefficients reassemble the static weights into dynamic frequency scaling weights of shape $C \times \mathbf{B}$, $\mathbf{B} \in \mathbb{R}^{b \times b}$ is number of frequency band.




Figure 4. Illustration of frequency scaling weight generation. The input feature map of dimensions $C \times H \times W$ is first and then fed into an MLP with a Tanh activation function to generate dynamic coefficients of dimensions $g \times n$. These dynamic coefficients are multiplied with n learnable static scaling weights $\in \mathbb{R}^{\frac{C}{g} \times b \times b}$ to produce the final scaling weights, which are upsampled to match the size of the feature map in the Fourier domain ($C \times H \times W$). This mechanism enables precise adjustment of various frequency components within the feature map dynamically.

Table 1. Quantitative comparisons using Vision Transformer [21] with *anti-oversmoothing* methods on the ADE20K *val* set [111].

Method <i>Segmentor</i> _{[NeurIPS2021] [72]}	Params	FLOPS	mIoU	
			SS	MS
DeiT-T _{[ICML2021] [78]}	6.7M	118G	35.7	36.7
+ AttScale _{[ICLR2022] [83]}	6.7M	118G	36.8 $\pm$ 1.1	37.8 $\pm$ 1.1
+ FeatScale _{[ICLR2022] [83]}	6.7M	118G	37.0 $\pm$ 1.3	37.9 $\pm$ 1.2
+ NeuTRENO _{[NeurIPS2023] [61]}	6.7M	118G	37.2 $\pm$ 1.5	38.1 $\pm$ 1.4
+ FDAM (Ours)	6.9M	120G	38.3$\pm$2.6	39.5$\pm$2.8

To reduce parameter cost and be parameter-efficient, we adopt a group-wise reassembly strategy. Each static scaling weight has a shape of $\frac{C}{g} \times b \times b$, and the dynamic coefficients have a shape of $g \times n$. The dynamic frequency scaling weights are generated via matrix multiplication:

$$\tilde{\mathbf{W}} = \mathbf{D} \cdot \text{stack}\{\mathbf{W}_1, \dots, \mathbf{W}_n\}, \quad (8)$$

where $\tilde{\mathbf{W}} \in \mathbb{R}^{C \times b \times b}$ represents the dynamic frequency scaling weights, $\mathbf{W}_i \in \mathbb{R}^{\frac{C}{g} \times b \times b}$ is the i -th static weight, and $\mathbf{D} \in \mathbb{R}^{g \times n}$ is the dynamic coefficient output by MLP.

Feature Frequency Modulation. With the frequency scaling weights obtained, we can modulate the feature as follows:

$$\mathbf{X} = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{X}) \odot \text{upsample}(\tilde{\mathbf{W}})), \quad (9)$$

where $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$ denotes the feature map. Note that we upsample the frequency scaling weights $\tilde{\mathbf{W}}$ to match the dimensions of $\mathbf{X}$, allowing us to effectively modulate the $\mathbf{B} \in \mathbb{R}^{b \times b}$ frequency bands within $\mathbf{X}$. This approach facilitates precise manipulation of frequency information within the feature map, thereby enhancing the model’s capacity to capture intricate details and textures.

4. Experiments

Datasets and Metrics. We evaluate our methods on challenging datasets, including ADE20K [111], COCO [51], and DOTA [88]. For segmentation tasks, we use mean Intersection over Union (mIoU) as the evaluation metric [7, 11, 23, 53, 57]. For object detection and instance segmentation, we use Average Precision (AP) [32, 68]. For panoptic segmentation, we use Panoptic Quality (PQ) [39].

Table 2. *Semantic segmentation* comparison with SegFormer [92] and UPerNet [91] on the ADE20K *val* set. SS and MS indicate single- and multi-scale test time settings.

Method (<i>SegFormer</i> <i>UPerNet</i>)	Params	FLOPS	mIoU	
			SS	MS
SegFormer-B0 _{[NeurIPS2021] [92]}	3.8M	8.6G	37.4	38.0
SegFormer-B0 + FDAM (Ours)	3.9M	8.9G	39.8$\pm$2.4	40.2$\pm$2.2
ResNet-50 _{[CVPR2016] [33]}	66M	947G	40.7	41.8
ResNet-101 _{[CVPR2016] [33]}	85M	1029G	42.9	44.0
DeiT-S _{[ICML2021] [78]}	52.1M	360G	42.9	43.8
DeiT-S + FDAM (Ours)	52.6M	363G	44.3$\pm$1.4	45.0$\pm$1.2
DeiT-B _{[ICML2021] [78]}	122M	787G	45.4	47.2
ViT-B-MLN _{[ECCV2020] [21, 71]}	144M	2007G	46.8	48.5
Swin-B _{[ICCV2021] [55]}	121M	1188G	48.1	49.7
NAT-B _{[CVPR2023] [31]}	123M	1137G	48.5	49.7
ConvNeXt-B _{[CVPR2022] [56]}	122M	1170G	49.1	49.9
ConvNeXt-B-dcls _{[ICLR2023] [38]}	122M	1170G	49.3	-
Swin-B-HAT _{[ECCV2022] [1]}	121M	1188G	48.9	50.3
DiNAT-B _{[arXiv2022] [30]}	123M	1137G	49.6	50.4
Focal-B _{[NeurIPS2022] [95]}	126M	1354G	49.0	50.5
DAT-B _{[CVPR2022] [89]}	121M	1212G	49.4	50.6
InceptionNeXt-B _{[CVPR2024] [103]}	115M	1159G	-	50.6
PeLK-B _{[CVPR2024] [5]}	126M	1237G	50.4	-
ViT-B-MLN + FDAM (Ours)	146M	2015G	48.0 $\pm$ 1.2	49.5 $\pm$ 1.0
MogaNet-L _{[ICLR2024] [44]}	113M	1176G	50.9	-
ConvFormer-M36 _{[TPAMI2024] [101]}	85M	1113G	51.3	-
OverLoCK-B _{[CVPR2025] [59]}	124M	1202G	51.7	52.3
DeiT3-B _{[ECCV2022] [79]}	144M	1283G	51.8	52.8
DeiT-B + FDAM (Ours)	124M	795G	46.5 $\pm$ 1.1	48.2 $\pm$ 1.0
ViT-B-MLN + FDAM (Ours)	146M	2015G	48.0 $\pm$ 1.2	49.5 $\pm$ 1.0
DeiT3-B + FDAM (Ours)	123M	1290G	52.6$\pm$0.8	53.4$\pm$0.6
MambaOut-B _{[arXiv2024] [102]}	112M	1178G	49.6	51.0
VMamba-B _{[arXiv2024] [112]}	122M	1170G	51.0	51.6
S.Mamba-B _{[ICLR2025] [90]}	127M	1176G	51.8	52.6
S.Mamba-B + FDAM (Ours)	129M	1180G	52.3$\pm$0.5	53.0$\pm$0.4
Swin-L _{[ICCV2021] [55]}	234M	3230G	52.1	53.5
ConvNeXt-XL _{[CVPR2022] [56]}	245M	2458G	53.2	53.7
MogaNet-XL _{[ICLR2024] [44]}	214M	2451G	54.0	-
DeiT3-L _{[ECCV2022] [79]}	354M	2231G	53.5	54.3
DeiT3-L + FDAM (Ours)	358M	2246G	54.1$\pm$0.6	54.8$\pm$0.5

Implementation Details. We follow the settings from the original papers for UPerNet [91], MaskDINO [41], and ViT [79]. On ADE20K [111], we train models for 160k iterations, following previous practice [55, 92]. On COCO [52] and DOTA [88], we adhere to standard practices [32, 46], training models for 12 epochs ($1 \times$ schedule). More details are provided in the *supplementary material*.

5. Main Results

In this section, we evaluate our method on a range of tasks, including object detection, instance segmentation, and semantic segmentation, using standard benchmarks such as COCO [51], ADE20K [111], and DOTA [88].

We first compare our method with recent anti-oversmoothing approaches [61, 83], followed by a comparison with state-of-the-art Vision Transformer (ViT [21], Swin [55], DAT [89], DeiT3 [79], DiNAT [30]), convolutional networks (ConvNeXt-B [56], Focal-B [95], ConvFormer [101], PeLK-B [5], InceptionNeXt [103],

Table 3. *Object detection and instance segmentation* comparison on the COCO validation set [51]. * indicates reproduced results.

Model (Backbone: $\frac{R50}{ViT}$)	Epochs	Params	FLOPs	AP ^{box}	AP ^{mask}
DETR _(ECCV2020) [4]	500	41M	86G	42.0	-
Deform. DETR _(ICLR2021) [113]	50	40M	173G	43.8	-
DAB-DETR _(ICLR2022) [54]	50	44M	94G	42.2	-
Mask-RCNN _(ICCV2017) [32]	36	40M	207G	40.9	37.1
HTC _(ICCV2019) [6]	36	80M	441G	44.9	39.7
QueryInst _(ICCV2021) [22]	36	-	-	45.6	40.6
Mask2Former _(CVPR2022) [15]	12	44M	226G	-	38.7
DDQ R-CNN _(CVPR2023) [108]	12	-	249G	44.6	41.2
Mask DINO* _(CVPR2023) [41]	12	52M	286G	45.5	41.2
Swin-T _(ICCV2021) [55]	12	48M	267G	42.7	39.3
ViTDet-B _(ECCV2022) [47]	12	90M	463G	43.8	39.9
ViTDet-L _(ECCV2022) [47]	12	308M	1542G	46.8	42.5
ViT-Adapter-B _(ICLR2023) [14]	12	120M	-	47.0	41.8
AdaptFormer-B _(ICLR2023) [114]	12	119M	733G	44.5	40.3
LoSA-B _(ICLR2023) [14]	12	117M	722G	45.1	41.8
META-B _(ICLR2023) [107]	12	115M	720G	45.4	42.3
ViT-CoMer-S _(CVPR2024) [87]	12	50M	-	45.8	40.5
PIIP-SBL _(NeurIPS2024) [114]	12	493M	727G	46.7	40.8
Mask DINO + FDAM (Ours)	12	53M	289G	47.1^{+1.6}	42.6^{+1.4}

Table 4. *Panoptic segmentation* comparison on the COCO validation set [51]. * indicates reproduced results.

Model (Backbone: R50)	Epochs	#Query	PQ	PQ Th	PQ St
DETR _(ECCV2020) [4]	500+25	100	43.4	48.2	36.0
PanopticFPN _(CVPR2019) [40]	36	-	42.5	50.3	30.7
PanopticFCN _(CVPR2021) [48]	36	-	44.3	53.0	36.5
MaskFormer _(NeurIPS2021) [16]	300	100	46.5	51.0	39.8
Mask2Former _(CVPR2022) [15]	12	100	46.9	52.5	38.4
Mask DINO* _(CVPR2023) [41]	12	300	48.7	54.6	40.0
Mask DINO + FDAM (Ours)	12	300	49.6^{+0.9}	55.5^{+0.9}	40.7^{+0.7}

MogaNet [44], OverLoCK [59]), and Mamba (MambaOut [102], Vision Mamba [112], Spatial Mamba [79]).

Our method achieves considerable improvements in dense prediction tasks, including object detection, semantic segmentation, instance segmentation, panoptic segmentation, and remote sensing object detection, by addressing the limitations of vision transformers. It does so with minimal additional parameters and FLOPS overhead. Our method is highly versatile, integrating seamlessly with state-of-the-art architectures such as ViT [21] and MaskDINO [41], as well as with Mamba models like Spatial Mamba [90]. Experiments demonstrate that our method consistently outperforms recent state-of-the-art baselines.

Comparing with Anti-Oversmoothing Methods. Table 1 compares various anti-oversmoothing methods for ViT [78]. Our proposed FDAM achieves the highest performance, with an mIoU of 38.3 for single-scale (SS) evaluation and 39.5 for multi-scale (MS) evaluation. This outperforms existing methods such as AttScale [83], FeatScale [83], and NeuTRENO [61]. Specifically, FDAM improves SS mIoU by 2.6 and MS mIoU by 2.8 compared to vanilla DeiT-T [78]. These results demonstrate the effectiveness of FDAM in addressing the oversmoothing problem and enhancing segmentation performance.

Semantic Segmentation. Table 2 shows that FDAM con-

sistently enhances performance across architectures and scales, with minimal computational overhead (0.5M/0.3G for SegFormer-B0, 0.5M/3G for DeiT-S). For SegFormer-B0, FDAM improves mIoU by +2.4 in single-scale (SS) (37.4→39.8) and +2.2 in multi-scale (MS) (38.0→40.2). Similar gains are observed for DeiT-S, with +1.4 mIoU (42.9→44.3) in SS and +1.2 (43.8→45.0) in MS.

Object Detection and Instance Segmentation. We evaluate FDAM’s effectiveness on object detection and instance segmentation tasks using the COCO validation set [51]. As shown in Table 3, FDAM is integrated into the state-of-the-art Mask DINO framework [41], achieving notable performance gains with minimal computational overhead.

Specifically, FDAM enhances the AP^{box} by +1.6 (45.5→47.1) and the AP^{mask} by +1.4 (41.2→42.6), while adding only 1M parameters and 3G FLOPs to the baseline. These highlight FDAM’s effectiveness and efficiency.

Compared to other state-of-the-art methods, FDAM consistently delivers better or comparable performance with significantly fewer parameters and FLOPs. For example, ViTDet-L [47] achieves an AP^{box} of 46.8 with 308M parameters and 1542G FLOPs, whereas FDAM-enhanced Mask DINO achieves an AP^{box} of 47.1 with just 53M parameters and 289G FLOPs. This demonstrates FDAM’s effectiveness as a lightweight yet powerful enhancement to unlock the potential of transformers.

Panoptic Segmentation. On more challenging panoptic segmentation, our proposed Mask DINO + FDAM achieves superior performance compared to existing state-of-the-art methods, as shown in Table 4. Specifically, it attains a PQ score of 49.6, outperforming Mask DINO (48.7) and Mask2Former (46.9). The improvements are also evident in the PQTh and PQSt metrics, with values of 55.5 and 40.7, respectively. These results demonstrate the effectiveness of FDAM in enhancing panoptic segmentation performance.

Remote Sensing Object Detection. Table 2 shows that applying FDAM to the LSKNet-S model improves the mAP by +1.12, from 77.49 to 78.61, with a minimal increase in model parameters (0.3M) and negligible additional computational cost, demonstrating FDAM’s efficiency.

Combination with Heavy Models. As shown in Table 2, FDAM also benefits larger models, achieving notable improvements on DeiT-B [78] (+1.1) and ViT-B [21] (+1.2). Specifically, DeiT3-Large with FDAM shows considerable mIoU improvements, with the single-scale (SS) mIoU increasing by +0.6 (53.5→54.1) and the multi-scale (MS) mIoU improving by +0.5 (54.3→54.8). It demonstrates that FDAM remains effective even when scaling up.

Combination with Mamba. Though not designed for Mamba, when combined with the recent Spatial Mamba-B [90], an improvement of +0.5 is observed (51.8→52.3), as shown in Table 2. These results confirm FDAM’s adaptability and effectiveness across various architectures.

Table 5. Remote sensing object detection results on the DOTA-v1.0 dataset [88] under a single-scale training and testing setting.

Method	#Params	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP
DETR-based																	
AO ² -DETR _[TCSVT2022] [18]	40.8M	87.99	79.46	45.74	66.64	78.90	73.90	73.30	90.40	80.55	85.89	55.19	63.62	51.83	70.15	60.04	70.91
O ² -DETR _[arXiv2021] [60]	-	86.01	75.92	46.02	66.65	79.70	79.93	89.17	90.44	81.19	76.00	56.91	62.45	64.22	65.80	58.96	72.15
ARS-DETR _[arXiv2022] [105]	41.6M	86.61	77.26	48.84	66.76	78.38	78.96	87.40	90.61	82.76	82.19	54.02	62.61	72.64	72.80	64.96	73.79
One-stage																	
SASM _[AAAI2022] [34]	36.6M	86.42	78.97	52.47	69.84	77.30	75.99	86.72	90.89	82.63	85.66	60.13	68.25	73.98	72.22	62.37	74.92
R3Det-GWD _[ICML2021] [97]	41.9M	88.82	82.94	55.63	72.75	78.52	83.10	87.46	90.21	86.36	85.44	64.70	61.41	73.46	76.94	57.38	76.34
R3Det-KLD _[NeurIPS] [99]	41.9M	88.90	84.17	55.80	69.35	78.72	84.08	87.00	89.75	84.32	85.73	64.74	61.80	76.62	78.49	70.89	77.36
O-RepPoints _[CVPR2022] [45]	36.6M	87.02	83.17	54.13	71.16	80.18	78.40	87.28	90.90	85.97	86.25	59.90	70.49	73.53	72.27	58.97	75.97
R-FCOS _[ICCV2019] [76]	31.9M	88.52	77.54	47.06	63.78	80.42	80.50	87.34	90.39	77.83	84.13	55.45	65.84	66.02	72.77	49.17	72.45
R3Det _[AAAI2021] [96]	41.9M	89.00	75.60	46.64	67.09	76.18	73.40	79.02	90.88	78.62	84.88	59.00	61.16	63.65	62.39	37.94	69.70
S2ANet _[TGRS2021] [28]	38.5M	89.11	82.84	48.37	71.11	78.11	78.39	87.25	90.83	84.90	85.64	60.36	62.60	65.26	69.13	57.94	74.12
Two-stage																	
SCRDet _[ICCV2019] [98]	41.9M	89.98	80.65	52.09	68.36	68.36	60.32	72.41	90.85	87.94	86.86	65.02	66.68	66.25	68.24	65.21	72.61
G-V _[TPAM2020] [94]	41.1M	89.64	85.00	52.26	77.34	73.01	73.14	86.82	90.74	79.02	86.81	59.55	70.91	72.94	70.86	57.32	75.02
CenterMap [58]	41.1M	89.02	80.56	49.41	61.98	77.99	74.19	83.74	89.44	78.01	83.52	47.64	65.93	63.68	67.07	61.59	71.59
ReDet _[CVPR2021] [29]	31.6M	88.79	82.64	53.97	74.00	78.13	84.06	88.04	90.89	87.78	85.75	61.76	60.39	75.96	68.07	63.59	76.25
Roi Trans _[CVPR2019] [19]	55.1M	89.01	77.48	51.64	72.07	74.43	77.55	87.76	90.81	79.71	85.27	58.36	64.11	76.50	71.99	54.06	74.05
R. F. R-CNN _[TPAM2016] [69]	41.1M	89.40	81.81	47.28	67.44	73.96	73.12	85.03	90.90	85.15	84.90	56.60	64.77	64.70	70.28	62.22	73.17
O-RCNN _[ICCV2024] [93]	74.4M	89.40	82.48	55.33	73.88	79.37	84.05	88.06	90.90	86.44	84.83	63.63	70.32	74.29	71.91	65.43	77.35
LSKNet-S _[ICCV2024] [93]	31.0M	89.66	85.52	57.72	75.70	74.95	78.69	88.24	90.88	86.79	86.38	66.92	63.77	77.77	74.47	64.82	77.49
GRA _[ECCV2024] [82]	41.7M	89.27	81.71	53.44	74.18	80.02	85.08	87.97	90.90	86.08	85.52	66.93	68.37	74.20	72.58	68.48	77.65
PKINet-S _[CVPR2024] [3]	30.8M	89.72	84.20	55.81	77.63	80.25	84.45	88.12	90.88	87.57	86.07	66.86	70.23	77.47	73.62	62.94	78.39
LSKNet-S + FDAM (Ours)	31.3M	89.85	83.86	55.12	78.84	79.71	85.10	88.35	90.88	88.77	86.12	68.31	66.57	76.77	72.49	68.36	78.61^{+1.12}




Figure 5. (a) Effective rank analysis for feature rank collapse. Higher *effective rank* [37] indicates a greater ability to capture complex patterns and nuanced details from the input data. FDAM maintains a consistently higher effective rank across all layers compared to the DeiT model using standard attention, demonstrating enhanced expressiveness of the attention mechanisms. (b) Feature similarity analysis. The cosine similarity increases with depth in the baseline DeiT model, indicating a loss of diversity in patch representations [61, 83]. The proposed FDAM method largely reduces this similarity, promoting more diverse features.

6. Analyses and Discussion

Here, we analyze effectiveness of the proposed method. More analyses are provided in *supplementary material*.

Rank Collapse Analysis. The inherent low-pass filtering characteristic of the attention mechanism lead to rank collapse [20], which degrades the model’s representational capacity. To analyze this, we employ the concept of *effective rank*, defined as the Shannon entropy of the normalized singular values of a matrix [37]. This measure provides a continuous and informative alternative to the traditional binary rank metric by capturing the distribution of singular values.

As shown in Figure 5(a), the effective rank of DeiT [78] model decreases rapidly with increasing depth, indicating a loss of feature anisotropy and hindering the model’s ability to capture complex patterns. In contrast, our FDAM main-




Figure 6. Visualization of attention modulation learned by AttInv. Warmer colors indicate higher values for high-pass filters. AttInv tends to assign higher values to foreground regions and semantic edges, emphasizing the focus on salient objects and boundaries.

tains a consistently higher effective rank across all layers, demonstrating that FDAM mitigates rank collapse and enhances expressiveness of attention mechanism.

Feature Similarity Analysis. We assess our model’s feature similarity using cosine similarity across layers in Figure 5(b). The DeiT shows a sharp rise in patch-wise cosine similarity with depth, hitting 0.70 by layer 11, signaling feature homogenization from repeated self-attention operations that erode discriminative spatial information. Our FDAM reduces late-layer similarity by up to 35%, enhancing robustness and task performance through more diverse representations. This analysis shows our methods curb over-smoothing, promote diversity, improve representational capacity, and boost performance on vision tasks.

Visualization of AttInv. Figure 6 shows that AttInv assigns higher high-pass filter values to foreground and semantic edges. This highlights the model’s focus on salient objects




Figure 7. Visualization of frequency modulation map learned by FreqScale. From the center to the border are low- to high-frequency components. Brighter colors highlight amplified frequency components. This demonstrates that FreqScale tends to enhance high-frequency components in the feature maps, effectively preventing over-smoothing caused by the attention mechanism.




Figure 8. Frequency response visualization for the *last three attention layers*. Warmer colors denote higher response. DeiT shows a strong response for the lowest direct current frequency. Our approach shows a balanced distribution across frequency bands, with a stronger emphasis on high-frequency components. This correlates with higher accuracy in downstream dense prediction tasks requiring fine-grained discriminative spatial details.

and boundaries, emphasizing the importance of these areas in capturing discriminative details and textures.

Visualization of FreqScale. Figure 7 demonstrates that FreqScale tends to amplify the high-frequency components in the feature maps, effectively preventing over-smoothing caused by the attention mechanism.

Feature Response Visualization. To further illustrate FDAM’s effectiveness, we visualize the frequency responses in Figure 8. The left spectrum of DeiT shows a strong concentration in low-frequency regions, indicating a bias toward coarse features. In contrast, FDAM (right) exhibits enhanced activation in mid-to-high frequencies, preserving fine-grained details and localized structures. This improved frequency response aligns with the higher accuracy observed in dense prediction tasks, where fine-grained spatial information is crucial.

Feature Visualization. As shown in Figure 9, DeiT features show a tendency to blur details and textures due to the model’s inherent low-pass filtering characteristic. This results in a loss of fine-grained details crucial for tasks requiring precise visual understanding. In contrast, our method generates feature maps with sharper, more discriminative details, effectively highlighting object structures. The feature spectrum of DeiT demonstrates a strong bias toward




Figure 9. Feature and spectrum visualization. (a) Input image. (b), (c) feature maps. (d), (e) feature spectrum. Our approach generates semantically focused activations (c). Compared to DeiT’s feature maps (b), our feature maps (c) capture sharper, more discriminative details, emphasizing object structures. The spectrum (d) shows DeiT’s dominance in low-frequency components, whereas our method exhibits stronger high-frequency components, indicating better edge and detail preservation.

low-frequency components. However, our feature spectrum shows a more balanced distribution across frequency bands, suggesting better preservation of fine-grained details and localized features critical for precise spatial discrimination.

7. Conclusion

We identified and addressed a key limitation of ViTs: their inherent low-pass filtering, which causes frequency vanishing and loss of fine-grained details, hindering dense prediction tasks. To overcome this, we proposed Frequency-Dynamic Attention Modulation (FDAM), which comprises Attention Inversion (AttInv) and Frequency Dynamic Scaling (FreqScale). AttInv restructures self-attention into a dynamic mix of high- and low-pass filters, enhancing frequency flexibility, while FreqScale further refines frequency response function by adaptively re-weighting frequency bands to preserve crucial frequency information. It can be *easily plugged into* existing ViT architectures.

Extensive experiments demonstrate that FDAM effectively avoids representation collapse and boosts performance in semantic segmentation, object detection, and instance segmentation with minimal cost. Additionally, FDAM achieves state-of-the-art results when applied to remote sensing object detection, showcasing its potential for real-world applications. Beyond its empirical success, FDAM offers new theoretical insights into self-attention’s spectral properties, paving the way for frequency-adaptive transformers and more robust vision architectures.

Acknowledgements

This work was supported by the National Key R&D Program of China (2022YFC3300705), the National Natural Science Foundation of China (62331006, 62171038, and 62088101), the Fundamental Research Funds for the Central Universities, and the JST Moonshot R&D Grant Number JPMJMS2011, Japan.

References

- [1] Jiawang Bai, Li Yuan, Shu-Tao Xia, Shuicheng Yan, Zhifeng Li, and Wei Liu. Improving vision transformers by revisiting high-frequency components. In *Proceedings of European Conference on Computer Vision*, pages 1–18. Springer, 2022. 3, 5
- [2] John Bird. *Electrical circuit theory and technology*. Routledge, 2017. 1
- [3] Xinhao Cai, Qiuxia Lai, Yuwei Wang, Wenguan Wang, Zeren Sun, and Yazhou Yao. Poly kernel inception network for remote sensing detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 27706–27716, 2024. 7
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of European Conference on Computer Vision*, pages 213–229. Springer, 2020. 6
- [5] Honghao Chen, Xiangxiang Chu, Yongjian Ren, Xin Zhao, and Kaiqi Huang. Pelk: Parameter-efficient large kernel convnets with peripheral convolution. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 5557–5567, 2024. 5
- [6] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiao-xiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 4974–4983, 2019. 6
- [7] Linwei Chen, Zheng Fang, and Ying Fu. Consistency-aware map generation at multiple zoom levels using aerial image. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:5953–5966, 2022. 5
- [8] Linwei Chen, Ying Fu, Lin Gu, Chenggang Yan, Tatsuya Harada, and Gao Huang. Frequency-aware feature fusion for dense image prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):10763–10780, 2024. 2
- [9] Linwei Chen, Ying Fu, Lin Gu, Dezhi Zheng, and Jifeng Dai. Spatial frequency modulation for semantic segmentation. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 1(1):1–18, 2025. 3
- [10] Linwei Chen, Ying Fu, Kaixuan Wei, Dezhi Zheng, and Felix Heide. Instance segmentation in the dark. *International Journal of Computer Vision*, 131(8):2198–2218, 2023. 2
- [11] Linwei Chen, Lin Gu, and Ying Fu. When semantic segmentation meets frequency aliasing. In *Proceedings of International Conference on Learning Representations*, pages 1–13, 2024. 3, 5
- [12] Linwei Chen, Lin Gu, Liang Li, Chenggang Yan, and Ying Fu. Frequency dynamic convolution for dense image prediction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 30178–30188, 2025. 2
- [13] Linwei Chen, Lin Gu, Dezhi Zheng, and Ying Fu. Frequency-adaptive dilated convolution for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3414–3425, 2024. 3
- [14] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. In *Proceedings of International Conference on Learning Representations*, pages 1–14, 2023. 6
- [15] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022. 2, 6
- [16] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *Proceedings of Advances in Neural Information Processing Systems*. 6
- [17] Lu Chi, Borui Jiang, and Yadong Mu. Fast fourier convolution. In *Proceedings of Advances in Neural Information Processing Systems*, volume 33, pages 4479–4488, 2020. 2
- [18] Linhui Dai, Hong Liu, Hao Tang, Zhiwei Wu, and Pinhao Song. Ao2-detr: Arbitrary-oriented object detection transformer. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(5):2342–2356, 2022. 7
- [19] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2849–2858, 2019. 7
- [20] Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *International conference on machine learning*, pages 2793–2803. PMLR, 2021. 3, 7
- [21] 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. In *Proceedings of International Conference on Learning Representations*, pages 1–12, 2020. 1, 2, 3, 5, 6
- [22] Yuxin Fang, Shusheng Yang, Xinggang Wang, Yu Li, Chen Fang, Ying Shan, Bin Feng, and Wenyu Liu. Instances as queries. In *Proceedings of IEEE International Conference on Computer Vision*, pages 6910–6919, 2021. 6
- [23] Ying Fu, Zheng Fang, Linwei Chen, Tao Song, and Defu Lin. Level-aware consistent multilevel map translation from satellite imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–14, 2022. 5
- [24] Yunyi Gao, Qiankun Liu, Lin Gu, and Ying Fu. Grayscale-assisted rgb image conversion from near-infrared images. *Tsinghua Science and Technology*, 30(5):2215–2226, 2025. 2

- [25] Rafael C Gonzalez. *Digital image processing*. Pearson education india, 2009. 2
- [26] Julia Grabinski, Steffen Jung, Janis Keuper, and Margret Keuper. Frequencylowcut pooling-plug and play against catastrophic overfitting. In *Proceedings of European Conference on Computer Vision*, pages 36–57, 2022. 3
- [27] John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro. Adaptive fourier neural operators: Efficient token mixers for transformers. In *Proceedings of International Conference on Learning Representations*, pages 1–12, 2022. 2
- [28] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2021. 7
- [29] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2786–2795, 2021. 7
- [30] Ali Hassani and Humphrey Shi. Dilated neighborhood attention transformer. 2022. 5
- [31] Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 6185–6194, 2023. 2, 5
- [32] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of IEEE International Conference on Computer Vision*, pages 2961–2969, 2017. 5, 6
- [33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 5
- [34] Liping Hou, Ke Lu, Jian Xue, and Yuqiu Li. Shape-adaptive selection and measurement for oriented object detection. In *Association for the Advancement of Artificial Intelligence*, volume 36, pages 923–932, 2022. 7
- [35] Qibin Hou, Cheng-Ze Lu, Ming-Ming Cheng, and Jiashi Feng. Conv2former: A simple transformer-style convnet for visual recognition. *IEEE Transactions Pattern Analysis and Machine Intelligence*, pages 8274–8283, 2024. 2
- [36] Zhipeng Huang, Zhizheng Zhang, Cuiling Lan, Zheng-Jun Zha, Yan Lu, and Baining Guo. Adaptive frequency filters as efficient global token mixers. In *Proceedings of IEEE International Conference on Computer Vision*, pages 1–11, 2023. 2
- [37] Minyoung Huh, Hossein Mobahi, Richard Zhang, Brian Cheung, Pulkit Agrawal, and Phillip Isola. The low-rank simplicity bias in deep networks. *Transactions on Machine Learning Research*, pages 1–15, 2023. 2, 7
- [38] Ismail Khalfaoui-Hassani, Thomas Pellegrini, and Timothée Masquelier. Dilated convolution with learnable spacings. In *Proceedings of International Conference on Learning Representations*, pages 1–13, 2023. 5
- [39] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019. 5
- [40] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. 6
- [41] Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M Ni, and Heung-Yeung Shum. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3041–3050, 2023. 2, 5, 6
- [42] Hesong Li and Ying Fu. Fcdfusion: A fast, low color deviation method for fusing visible and infrared image pairs. *Computational Visual Media*, 11(1):195–211, 2025. 2
- [43] Hesong Li, Ziqi Wu, Ruiwen Shao, Tao Zhang, and Ying Fu. Noise calibration and spatial-frequency interactive network for stem image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Conference*, pages 21287–21296, June 2025. 2
- [44] Siyuan Li, Zedong Wang, Zicheng Liu, Cheng Tan, Haitao Lin, Di Wu, Zhiyuan Chen, Jiangbin Zheng, and Stan Z Li. Moganet: Multi-order gated aggregation network. In *ICLR*, pages 1–19, 2024. 5, 6
- [45] Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1829–1838, 2022. 7
- [46] Yuxuan Li, Xiang Li, Yimain Dai, Qibin Hou, Li Liu, Yongxiang Liu, Ming-Ming Cheng, and Jian Yang. Lsknet: A foundation lightweight backbone for remote sensing. *International Journal of Computer Vision*, pages 1–22, 2024. 5
- [47] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *Proceedings of European Conference on Computer Vision*, pages 280–296. Springer, 2022. 1, 2, 6
- [48] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 214–223, 2021. 6
- [49] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *Proceedings of International Conference on Learning Representations*, pages 1–12, 2021. 2
- [50] Shiqi Lin, Zhizheng Zhang, Zhipeng Huang, Yan Lu, Cuiling Lan, Peng Chu, Quanzeng You, Jiang Wang, Zicheng Liu, Amey Parulkar, et al. Deep frequency filtering for domain generalization. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 11797–11807, 2023. 2
- [51] 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 *Proceedings of European Conference on Computer Vision*, pages 740–755, 2014. 5, 6

- [52] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. [5](#)
- [53] Songlin Liu, Linwei Chen, Li Zhang, Jun Hu, and Ying Fu. A large-scale climate-aware satellite image dataset for domain adaptive land-cover semantic segmentation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 205:98–114, 2023. [5](#)
- [54] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. 2022. [6](#)
- [55] 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 *Proceedings of IEEE International Conference on Computer Vision*, pages 10012–10022, 2021. [1](#), [2](#), [5](#), [6](#)
- [56] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. [5](#)
- [57] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 3431–3440, 2015. [5](#)
- [58] Yang Long, Gui-Song Xia, Shengyang Li, Wen Yang, Michael Ying Yang, Xiao Xiang Zhu, Liangpei Zhang, and Deren Li. On creating benchmark dataset for aerial image interpretation: Reviews, guidances, and million-aid. *IEEE Journal of selected topics in applied earth observations and remote sensing*, 14:4205–4230, 2021. [7](#)
- [59] Meng Lou and Yizhou Yu. Overlock: An overview-first-look-closely-next convnet with context-mixing dynamic kernels. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2025. [5](#), [6](#)
- [60] Teli Ma, Mingyuan Mao, Honghui Zheng, Peng Gao, Xiaodi Wang, Shumin Han, Errui Ding, Baochang Zhang, and David Doermann. Oriented object detection with transformer. *arXiv preprint arXiv:2106.03146*, 2021. [7](#)
- [61] Tam Nguyen, Tan Nguyen, and Richard Baraniuk. Mitigating over-smoothing in transformers via regularized nonlocal functionals. *Proceedings of Advances in Neural Information Processing Systems*, 36:80233–80256, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [62] Namuk Park and Songkuk Kim. How do vision transformers work? In *Proceedings of International Conference on Learning Representations*, pages 1–14, 2021. [3](#)
- [63] Badri Patro and Vijay Agneeswaran. Scattering vision transformer: Spectral mixing matters. In *Proceedings of Advances in Neural Information Processing Systems*. [3](#)
- [64] Badri N Patro, Vinay P Nambodiri, and Vijay Srinivas Agneeswaran. Spectformer: Frequency and attention is what you need in a vision transformer. *arXiv preprint arXiv:2304.06446*, 2023. [3](#)
- [65] Ioannis Pitas. *Digital image processing algorithms and applications*. John Wiley & Sons, 2000. [2](#)
- [66] Zequn Qin, Pengyi Zhang, Fei Wu, and Xi Li. Fcanet: Frequency channel attention networks. In *Proceedings of IEEE International Conference on Computer Vision*, pages 783–792, 2021. [2](#)
- [67] Yongming Rao, Wenliang Zhao, Zheng Zhu, Jiwen Lu, and Jie Zhou. Global filter networks for image classification. In *Proceedings of Advances in Neural Information Processing Systems*, volume 34, pages 980–993, 2021. [2](#)
- [68] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proceedings of Advances in Neural Information Processing Systems*, pages 91–99, 2015. [5](#)
- [69] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016. [7](#)
- [70] Dai Shi. Transnext: Robust foveal visual perception for vision transformers. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 17773–17783, 2024. [2](#)
- [71] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *Transactions on Machine Learning Research*, pages 1–13, 2022. [5](#)
- [72] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 7262–7272, 2021. [2](#), [5](#)
- [73] G Szentirmai. Electronic filter design handbook. *Proceedings of the IEEE*, 70(3):317–317, 1982. [1](#), [4](#)
- [74] Yuki Tatsunami and Masato Taki. Fft-based dynamic token mixer for vision. In *Association for the Advancement of Artificial Intelligence*, volume 38, pages 15328–15336, 2024. [3](#)
- [75] Ye Tian, Ying Fu, and Jun Zhang. Transformer-based under-sampled single-pixel imaging. *Chinese Journal of Electronics*, 32(5):1151–1159, 2023. [2](#)
- [76] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of IEEE International Conference on Computer Vision*, pages 9627–9636, 2019. [7](#)
- [77] Hesong Li Tianyu Zhu and Ying Fu. Trim-sod: A multi-modal, multi-task, and multi-scale spacecraft optical dataset. *Space: Science & Technology*, 1(1):1–27, 2025. [2](#)
- [78] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proceedings of International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. [2](#), [5](#), [6](#), [7](#)
- [79] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. In *European conference on computer vision*, pages 516–533. Springer, 2022. [5](#), [6](#)
- [80] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser,

- and Illia Polosukhin. Attention is all you need. In *Proceedings of Advances in Neural Information Processing Systems*, volume 30, pages 1–11, 2017. [2](#)
- [81] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 8684–8694, 2020. [2](#)
- [82] Jiangshan Wang, Yifan Pu, Yizeng Han, Jiayi Guo, Yiru Wang, Xiu Li, and Gao Huang. Gra: Detecting oriented objects through group-wise rotating and attention. In *Proceedings of European Conference on Computer Vision*, pages 298–315. Springer, 2024. [7](#)
- [83] Peihao Wang, Wenqing Zheng, Tianlong Chen, and Zhangyang Wang. Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. In *International Conference on Learning Representations*. [1](#), [3](#), [5](#), [6](#), [7](#)
- [84] 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 *Proceedings of IEEE International Conference on Computer Vision*, pages 568–578, 2021. [2](#)
- [85] Steve Winder. *Analog and digital filter design*. Elsevier, 2002. [1](#), [4](#)
- [86] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *Proceedings of IEEE International Conference on Computer Vision*, pages 22–31, 2021. [2](#)
- [87] Chunlong Xia, Xinliang Wang, Feng Lv, Xin Hao, and Yifeng Shi. Vit-comer: Vision transformer with convolutional multi-scale feature interaction for dense predictions. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 5493–5502, 2024. [6](#)
- [88] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3974–3983, 2018. [5](#), [7](#)
- [89] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. [5](#)
- [90] Chaodong Xiao, Minghan Li, Zhengqiang Zhang, Deyu Meng, and Lei Zhang. Spatial-mamba: Effective visual state space models via structure-aware state fusion. In *ICLR*, pages 1–13, 2025. [5](#), [6](#)
- [91] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of European Conference on Computer Vision*, pages 418–434, 2018. [5](#)
- [92] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Proceedings of Advances in Neural Information Processing Systems*, volume 34, pages 12077–12090, 2021. [1](#), [2](#), [5](#)
- [93] Xingxing Xie, Gong Cheng, Jiabao Wang, Ke Li, Xiwen Yao, and Junwei Han. Oriented r-cnn and beyond. *International Journal of Computer Vision*, pages 1–23, 2024. [7](#)
- [94] Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 43(4):1452–1459, 2020. [7](#)
- [95] Jianwei Yang, Chunyuan Li, Xiyang Dai, and Jianfeng Gao. Focal modulation networks. *Proceedings of Advances in Neural Information Processing Systems*, 35:4203–4217, 2022. [2](#), [5](#)
- [96] Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In *Association for the Advancement of Artificial Intelligence*, volume 35, pages 3163–3171, 2021. [7](#)
- [97] Xue Yang, Junchi Yan, Qi Ming, Wentao Wang, Xiaopeng Zhang, and Qi Tian. Rethinking rotated object detection with gaussian wasserstein distance loss. In *Proceedings of International Conference on Machine Learning*, pages 11830–11841. PMLR, 2021. [7](#)
- [98] Xue Yang, Jirui Yang, Junchi Yan, Yue Zhang, Tengfei Zhang, Zhi Guo, Xian Sun, and Kun Fu. Scrddet: Towards more robust detection for small, cluttered and rotated objects. In *Proceedings of IEEE International Conference on Computer Vision*, pages 8232–8241, 2019. [7](#)
- [99] Xue Yang, Xiaojiang Yang, Jirui Yang, Qi Ming, Wentao Wang, Qi Tian, and Junchi Yan. Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. *Proceedings of Advances in Neural Information Processing Systems*. [7](#)
- [100] Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier perspective on model robustness in computer vision. In *Proceedings of Advances in Neural Information Processing Systems*, volume 32, 2019. [2](#)
- [101] Weihao Yu, Chenyang Si, Pan Zhou, Mi Luo, Yichen Zhou, Jiashi Feng, Shuicheng Yan, and Xinchao Wang. Metaformer baselines for vision. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 46(2):896–912, 2024. [2](#), [5](#)
- [102] Weihao Yu and Xinchao Wang. Mambaout: Do we really need mamba for vision? *arXiv preprint arXiv:2405.07992*, pages 1–16, 2024. [5](#), [6](#)
- [103] Weihao Yu, Pan Zhou, Shuicheng Yan, and Xinchao Wang. Inceptionnext: When inception meets convnext. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 5672–5683, 2024. [5](#)
- [104] Guhnoo Yun, Juhan Yoo, Kijung Kim, Jeongho Lee, and Dong Hwan Kim. Spanet: Frequency-balancing token mixer using spectral pooling aggregation modulation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 1–16, 2023. [3](#)

- [105] Y Zeng, X Yang, Q Li, Y Chen, and J Yan. Ars-detr: Aspect ratio sensitive oriented object detection with transformer. arxiv 2023. *arXiv preprint arXiv:2303.04989*. 7
- [106] Bowen Zhang, Zhi Tian, Quan Tang, Xiangxiang Chu, Xiaolin Wei, Chunhua Shen, et al. Segvit: Semantic segmentation with plain vision transformers. *Proceedings of Advances in Neural Information Processing Systems*, 35:4971–4982, 2022. 2
- [107] Dong Zhang, Rui Yan, Pingcheng Dong, and Kwang-Ting Cheng. Memory efficient transformer adapter for dense predictions. In *Proceedings of International Conference on Learning Representations*, pages 1–15, 2025. 6
- [108] Shilong Zhang, Xinjiang Wang, Jiaqi Wang, Jiangmiao Pang, Chengqi Lyu, Wenwei Zhang, Ping Luo, and Kai Chen. Dense distinct query for end-to-end object detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 7329–7338, 2023. 6
- [109] Tao Zhang, Ying Fu, Jun Zhang, and Chenggang Yan. Deep guided attention network for joint denoising and demosaicing in real image. *Chinese Journal of Electronics*, 33(1):303–312, 2024. 2
- [110] Yingkai Zhang, Zeqiang Lai, Tao Zhang, Ying Fu, and Chenghu Zhou. Unaligned rgb guided hyperspectral image super-resolution with spatial-spectral concordance. *International Journal of Computer Vision*, pages 1–21, 2025. 2
- [111] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 633–641, 2017. 5
- [112] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model, 2024. 5, 6
- [113] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *Proceedings of International Conference on Learning Representations*, pages 1–12, 2021. 6
- [114] Xizhou Zhu, Xue Yang, Zhaokai Wang, Hao Li, Wenhan Dou, Junqi Ge, Lewei Lu, Yu Qiao, and Jifeng Dai. Parameter-inverted image pyramid networks. *Proceedings of Advances in Neural Information Processing Systems*, 37:132267–132288, 2024. 2, 6