

SparX: A Sparse Cross-Layer Connection Mechanism for Hierarchical Vision Mamba and Transformer Networks

Meng Lou, Yunxiang Fu, Yizhou Yu*

School of Computing and Data Science, The University of Hong Kong
loumeng@connect.hku.hk, yunxiang@connect.hku.hk, yizhouy@acm.org

Abstract

Due to the capability of dynamic state space models (SSMs) in capturing long-range dependencies with linear-time computational complexity, Mamba has shown notable performance in NLP tasks. This has inspired the rapid development of Mamba-based vision models, resulting in promising results in visual recognition tasks. However, such models are not capable of distilling features across layers through feature aggregation, interaction, and selection. Moreover, existing cross-layer feature aggregation methods designed for CNNs or ViTs are not practical in Mamba-based models due to high computational costs. Therefore, this paper aims to introduce an efficient cross-layer feature aggregation mechanism for vision backbone networks. Inspired by the Retinal Ganglion Cells (RGCs) in the human visual system, we propose a new sparse cross-layer connection mechanism termed SparX to effectively improve cross-layer feature interaction and reuse. Specifically, we build two different types of network layers: ganglion layers and normal layers. The former has higher connectivity and complexity, enabling multi-layer feature aggregation and interaction in an input-dependent manner. In contrast, the latter has lower connectivity and complexity. By interleaving these two types of layers, we design a new family of vision backbone networks with sparsely cross-connected layers, achieving an excellent trade-off among model size, computational cost, memory cost, and accuracy in comparison to its counterparts. For instance, with fewer parameters, SparX-Mamba-T improves the top-1 accuracy of VMamba-T from 82.5% to 83.5%, while SparX-Swin-T achieves a 1.3% increase in top-1 accuracy compared to Swin-T. Extensive experimental results demonstrate that our new connection mechanism possesses both superior performance and generalization capabilities on various vision tasks.

Code — <https://github.com/LMMMEng/SparX>

Extended version — <https://arxiv.org/abs/2409.09649>

1 Introduction

Contextual modeling plays a vital role in computer vision tasks, including image classification, object detection, and semantic segmentation. Recent advancements have explored large-kernel convolutions to enlarge receptive fields (Liu

*Corresponding Author
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

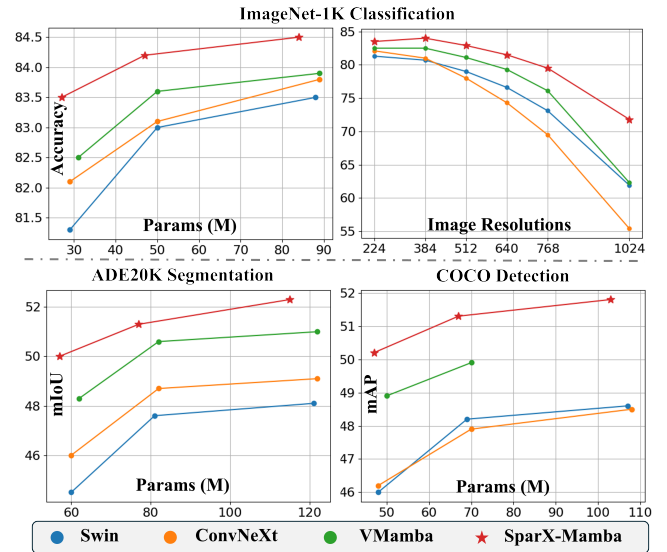


Figure 1: Performance comparison between SparX-Mamba and other methods on various vision tasks.

et al. 2022; Ding et al. 2022b; Liu et al. 2023; Ding et al. 2024) and Vision Transformers (ViTs) (Dosovitskiy et al. 2021; Touvron et al. 2021; Liu et al. 2021; Wang et al. 2021) to possess powerful long-range modeling capability. However, both CNNs with large convolution kernels and Vision Transformers exhibit limitations. The performance of CNNs still lags behind advanced ViTs while the latter suffers from substantial computational costs when processing high-resolution inputs, due to its quadratic complexity.

Empowered by dynamic State Space Models (SSMs) capable of modeling long-range dependencies with near-linear complexity, Mamba (Gu and Dao 2024) has demonstrated promising performance in natural language processing (NLP) tasks. Building upon this success, researchers have proposed Mamba-based models for vision tasks (Zhu et al. 2024; Yang et al. 2024; Liu et al. 2024; Huang et al. 2024; Behrouz, Santacatterina, and Zabih 2024; Pei, Huang, and Xu 2024). Such work has made successful attempts to introduce linear-time long-range modeling capability into vision models by leveraging SSMs, which are primarily used

to capture spatial contextual information within the same layer. Nonetheless, there might exist both complementarity and redundancy among token features in different layers as they capture image characteristics and semantics at different levels and granularities. Thus features across different layers need to be distilled to extract useful information and remove redundancy through aggregation, interaction, and selection. However, existing Mamba-based vision models do not possess such an ability to conduct feature distillation across layers, hindering the realization of the full power of SSMs.

Although cross-layer feature interaction and reuse have been shown to effectively improve performance in both CNNs and Vision Transformers, such as DenseNet (Huang et al. 2017) and FcaFormer (Zhang, Hu, and Wang 2023), these methods are not directly applicable to Mamba-based vision models. As shown in Table 5, when VMamba is restructured as a DenseNet-style Network (DSN), its throughput drops by almost 50%, accompanied by a 1GB increase in GPU memory usage. Meanwhile, as shown in Table 6, using Spatial Cross-Attention (SCA) in FcaFormer as a multi-layer feature interaction module results in a more than 80% increase in GPU memory usage compared to simple feature concatenation, due to the quadratic complexity of attention calculations. Therefore, designing an efficient connectivity pattern to facilitate cross-layer feature interaction in Mamba-based models poses a significant challenge.

Unlike existing works that primarily focus on improving Mamba-based token mixers, in this work, we propose a novel architectural design that significantly boosts the performance of SSM-based vision models. Our design draws inspiration from Retinal Ganglion Cells (RGCs) in the human visual system. RGCs serve as the information transmission hub in the human visual system, bridging the retinal input and the visual processing centers in the central nervous system (Kim et al. 2021). Due to the complex neural architecture and diverse cell composition within the RGC layer (Curcio et al. 1990; Watson 2014), it has a larger number of connections with other cell layers in the visual system, thereby fostering intricate neural interactions with other layers. Inspired by this biological mechanism, we propose a novel sparse cross-layer connection mechanism named SparX. We define two types of SSM-based basic layers, namely, normal layers and ganglion layers. The former has less information flow, taking a single input from the preceding layer and connecting with a small number of subsequent layers. In contrast, the latter establishes a larger number of connections with subsequent layers and encodes inputs from multiple preceding layers. The combination of ganglion and normal layers mimics the combination of RGC layers and non-RGC layers within the human visual system.

Existing literature (D’Souza and Lang 2020) also confirms that RGCs generate diverse intercellular communications including both RGC-RGC and RGC-non-RGC interactions. Inspired by this premise, we incorporate a new Dynamic Multi-layer Channel Aggregator (DMCA) into ganglion layers, aiming to efficiently facilitate adaptive feature aggregation and interaction across multiple layers. Furthermore, to maintain speed and lower memory consumption, we introduce a cross-layer sliding window, which only per-

mits each ganglion layer to connect with other ganglion layers in the same sliding window.

By hierarchically interleaving ganglion and normal layers in a feedforward network, we design a novel Mamba-based vision backbone, namely, Vision Mamba with sparsely cross-connected layers (SparX-Mamba). In addition to inherent advantages of Mamba-based models, our SparX-Mamba achieves an excellent trade-off between performance and computational cost. Notably, our SparX-Mamba differs from conventional vision models that are constructed by simply stacking the same layer design, resulting in similar feature extraction capabilities at every layer. In contrast, the two types of functionally distinct layers of SparX-Mamba facilitate a model to extract more diverse feature representations, thereby enhancing performance. It is worth noting that our SparX has also been successfully applied to hierarchical vision Transformers, such as Swin (Liu et al. 2021), and the resulting models also have the aforementioned advantages over the original Transformer models. For example, our SparX-Swin-T achieves a 1.3% improvement in top-1 accuracy compared to Swin-T.

As shown in Figure 1, the proposed SparX-Mamba demonstrates superior performance over recently proposed VMamba as well as classical CNN and Transformer-based models, with a smaller number of parameters (Params). For instance, when compared with VMamba-T, SparX-Mamba-T achieves a top-1 accuracy of 83.5% on ImageNet-1K, remarkably surpassing the accuracy of VMamba-T (82.5%). When integrated with UperNet for semantic segmentation on the ADE20K dataset, SparX-Mamba-T outperforms VMamba-T by 1.7% in mIoU. Furthermore, when integrated into Mask R-CNN for object detection on the COCO 2017 dataset, SparX-Mamba-T exceeds VMamba-T by 1.3% in AP^b and even outperforms VMamba-S, whose complexity is nearly 1.5 times that of SparX-Mamba-T. Meanwhile, both the small and base versions of SparX-Mamba also exhibit notable performance improvements over VMamba.

In summary, our main contributions are threefold: First, inspired by RGCs in the human visual system, we propose a novel skip-connection mechanism named SparX, which sparsely and dynamically configures cross-layer connections, thereby enabling diverse information flow and improved feature distillation. Second, on the basis of Mamba and Transformer, we propose two versatile vision backbones, SparX-Mamba and SparX-Swin, constructed from two types of functionally distinct layers, one of which is responsible for communications between relatively distant layers. Third, we conduct extensive experiments on image classification, object detection, and semantic segmentation tasks. Results demonstrate that our new network architecture achieves a remarkable trade-off between performance and computational cost.

2 Related Work

CNNs. Convolutional Neural Networks (CNNs) have emerged as the standard deep learning architecture in computer vision (Simonyan and Zisserman 2015; He et al. 2016; Huang et al. 2017). In the realm of modern CNNs, a notable shift has occurred from traditional small convolution

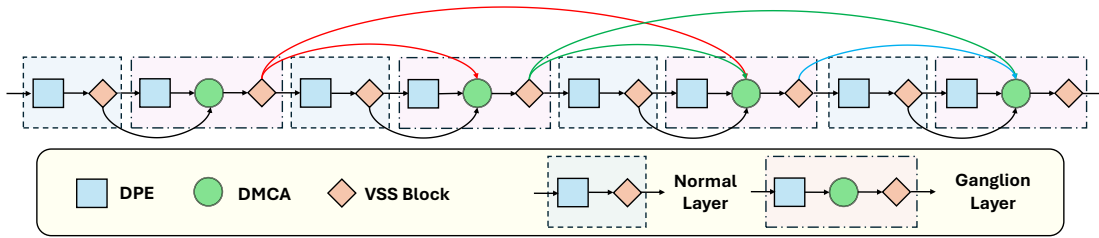


Figure 2: A concrete example of proposed SparX with $S=2$ and $M=2$.

kernels to a design focused on larger kernels. Noteworthy examples of this trend include ConvNeXt (Liu et al. 2022), RepLKNet (Ding et al. 2022b, 2024), and SLaK (Liu et al. 2023). Recently, InceptionNeXt (Yu et al. 2024b) has combined classical Inception networks with ConvNeXt, resulting in notable performance improvements.

Vision Transformers. To adapt Transformers from NLP tasks to computer vision tasks, a vision Transformer (ViT) (Dosovitskiy et al. 2021) splits an image into visual tokens through patch embedding, thereby enabling multi-head self-attention (MHSA) to learn token-to-token dependencies. To further generate hierarchical feature representations and improve model efficiency, many subsequent works have adopted pyramidal architecture designs, including window attention (Liu et al. 2021; Dong et al. 2022; Pan et al. 2023), sparse attention (Yang et al. 2021; Wang et al. 2022; Ren et al. 2022; Wu et al. 2023), and conv-attention hybrid models (Dai et al. 2021; Xiao et al. 2021; Li et al. 2022, 2023; Lou et al. 2023; Yu et al. 2024a).

Vision Mamba. Since Mamba (Gu and Dao 2024) has achieved outstanding performance in NLP tasks, many researchers have transferred Mamba to computer vision tasks. As the core of Mamba, SSM can model long-range dependencies with near-linear complexity and has shown excellent performance in vision tasks. For instance, ViM (Zhu et al. 2024) introduces a bidirectional SSM module and constructs an isotropic architecture like ViT (Dosovitskiy et al. 2021). Likewise, PlainMamba (Yang et al. 2024) also builds an isotropic architecture with continuous 2D scanning. VMamba (Liu et al. 2024) extends the scanning order to include four directions and is an early SSM-based hierarchical architecture. Subsequently, a series of hierarchical vision Mamba models have been proposed, including MambaMixer (Behrouz, Santacatterina, and Zabih 2024), LocalMamba (Huang et al. 2024), EfficientVMamba (Pei, Huang, and Xu 2024), and MSVMamba (Shi, Dong, and Xu 2024).

Skip-connections. ResNet (He et al. 2016) introduces residual connections for deep CNNs, alleviating the issues of vanishing and exploding gradients with bypassing shortcuts. To diversify information flow, DenseNet (Huang et al. 2017) further incorporates dense connections, which compute the input of a layer using the collection of outputs from all preceding layers. DPN (Chen et al. 2017) makes use of the skip-connections in both ResNet and DenseNet to build a dual-pathway network. In the domain of dense prediction tasks, U-Net (Ronneberger, Fischer, and Brox 2015) and FPN (Lin et al. 2017) leverage skip-connections to bridge the gap be-

tween low-level details from the encoder and high-level contexts from the decoder. This work introduces a novel skip-connection mechanism that dynamically configures sparse cross-layer connections, resulting in improved performance in comparison to existing methods.

Cross-feature Attention. Cross-feature attention can enhance feature representation by improving interactions among different features. For instance, many works (Chen, Fan, and Panda 2021; Lee et al. 2022; Wang et al. 2023; Ren et al. 2022; Wu et al. 2023) proposed to perform cross-feature attention by generating multi-scale tokens, which are then fed into self-attention or cross-attention layers to model multi-scale interactions. Recently, FcaFormer (Zhang, Hu, and Wang 2023) introduced cross-attention between the spatial tokens of a layer and representative tokens from previous layers to model cross-layer feature interactions. In contrast to existing methods, we introduce a novel channel-wise cross-attention mechanism that dynamically integrates feature channels from preceding layers, thereby facilitating efficient cross-layer feature interactions.

3 Method

3.1 Sparse Cross-Layer Connections

Overview. Inspired by how Retinal Ganglion Cells (RGCs) function in the human visual system, we introduce a novel sparse cross-layer connection mechanism named SparX. The goal is efficiently modeling cross-layer communications, generating diverse information flows, and improving feature reuse within Mamba-based architectures. As shown in Figure 2, SparX has three building components: Dynamic Position Encoding (DPE) (Chu et al. 2022), Mamba block, and the newly-introduced Dynamic Multi-layer Channel Aggregator (DMCA). Among them, DPE leverages a residual 3×3 depthwise convolution (DWConv), which finds widespread adoption in modern vision backbone networks (Li et al. 2023; Chu et al. 2021; Guo et al. 2022). The Mamba block effectively captures long-range dependencies with a near-linear computational complexity. However, the vanilla Mamba block (Gu and Dao 2024) is unsuitable for being directly embedded into vision backbones. In this regard, we utilize the VSS block, which has shown promising performance since its introduction in VMamba (Liu et al. 2024). To enhance channel mixing, a ConvFFN (Wang et al. 2022) is incorporated within VSS following (Shi, Dong, and Xu 2024). The proposed DMCA dynamically aggregates, selects, and encodes cross-layer features, giving rise to powerful and robust feature representa-

tions. On top of these basic components, two types of network layers are constructed in SparX: ganglion layers (“DPE \rightarrow DMCA \rightarrow VSS Block”) and normal layers (“DPE \rightarrow VSS Block”). These two layers are interleaved in a feed-forward network to enable effective modeling of cross-layer feature interactions in a sparse manner.

Connection Rules. In a vanilla dense connection (Huang et al. 2017), each layer is connected to all preceding layers. Despite the relatively low FLOPs of DenseNet, the storage of many preceding feature maps, which need to be repeatedly accessed during a forward pass, gives rise to high memory cost and low speed. Although efforts (Pleiss et al. 2017; Huang et al. 2019) have been made to lower memory cost, dense connections still face challenges when it is necessary to scale models to deeper and wider architectures. Consequently, the vanilla dense connection mechanism is too computationally heavy to use in Mamba-based models whose token mixer is more complex than standard convolutions. To this end, we propose SparX, which can be efficiently integrated into recent Mamba-based architectures. In SparX, we introduce two new rules, including sparse placement of ganglion layers and using a cross-layer sliding window to control the density of cross-layer connections.

Sparse Ganglion Layers aim to designate a subset of evenly spaced layers as ganglion layers while all remaining layers are normal layers. To control the density of ganglion layers, we define a hyperparameter called stride (S), which is one plus the number of normal layers between two nearest ganglion layers. We further define two types of cross-layer connections: (1) Intra-connection between a ganglion layer and a normal layer; (2) Inter-connection between two ganglion layers. To set up sparse cross-layer connections, a ganglion layer only has intra-connections with the normal layers between the closest preceding ganglion layer and itself while it builds inter-connections with multiple preceding ganglion layers. The rationale behind such cross-layer connections is that a ganglion layer can be regarded as an information exchange hub, which gathers information from the closest normal layers and exchanges it with other ganglion layers. To ensure that the final output of a network or stage contains rich semantic information, the last layer of a network or stage is typically a ganglion layer. As an example, if we have an 8-layer network and set the stride (S) to 2, the indices of the normal layers are $\{1, 3, 5, 7\}$ while the indices of the ganglion layers are $\{2, 4, 6, 8\}$.

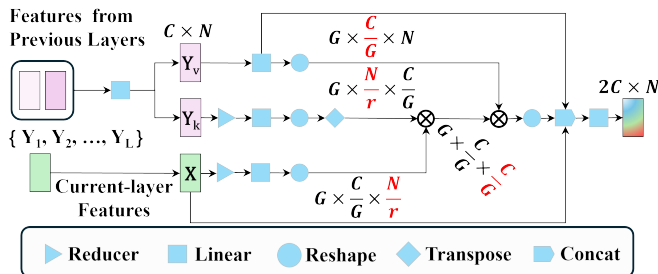


Figure 3: An illustration of the proposed DMCA. The red font indicates the dimensions in matrix multiplications.

Cross-layer Sliding Window is proposed to further improve computational efficiency, inspired by spatial sliding windows. The motivation behind this design is that, despite using the aforementioned sparse connections, deeper networks may still incur high memory costs due to the need to store and access a number of earlier feature maps. To this end, we introduce another hyperparameter M , which restricts a ganglion layer to have inter-connections with the M closest preceding ganglion layers only. Based on these two new rules, even without direct connections, information can still flow from shallower layers to deeper layers quickly through a relatively small number of intra-connections and inter-connections. Figure 2 illustrates an example of SparX with $S=2$ and $M=2$.

Dynamic Multi-layer Channel Aggregator (DMCA). To selectively retrieve complementary features from previous layers and dynamically model multi-layer interactions, we propose an efficient Dynamic Multi-layer Channel Aggregator (DMCA). As depicted in Figure 3, let $X \in \mathbb{R}^{C \times N}$ represents the features of a ganglion layer, and $\{Y_1, Y_2, \dots, Y_L\} \in \mathbb{R}^{C \times N}$ denote features from preceding ganglion and normal layers, where C and N represent the channel and spatial dimensions, respectively, while L denotes the total number of features from prior layers. First, we concatenate the features from preceding layers and project the channel dimension to $2C$ using a linear layer. Then, we evenly split the projected features into two segments along the channel dimension, Y_k and Y_v . Afterwards, we design a cross-layer grouped channel attention (CGCA) mechanism, where *query*, *key*, and *value* refer to X , Y_k , and Y_v , respectively. Specifically, we divide the channels into multiple groups, and channel-wise cross-attention is computed within each group. The number of groups (G) is always set to 4 in our experiments. Consequently, in our CGCA, the size of the overall attention map A for all groups is $G \times \frac{C}{G} \times \frac{C}{G}$, which avoids high computational and memory overhead when large spatial resolutions are used.

On the other hand, channel-wise cross-attention computes an attention matrix that measures the similarity between every pair of transformed channels from *query* and *key*, and the spatial dimension becomes eliminated. Note that the spatial resolution directly affects the cost of computing this attention matrix because it determines the dimensionality of every channel from *query* and *key*. One key observation is that the size of the attention matrix is independent of the spatial dimensions of *query* and *key*, whose spatially reduced versions can be used for computing an approximation of the original attention matrix. Based on this, we employ a spatial reduction strategy to improve computational efficiency. Specifically, we use a spatial reducer to compress the number of spatial tokens in *query* and *key* from N to N/r . In our experiments, the rule for setting r is making N/r equal to the number of tokens in the final stage of the network, following PVT (Wang et al. 2021, 2022). Although this operation might appear similar to spatial reduction used in PVT, it serves different purposes. In PVT, spatial reduction is exploited to compute spatial token-to-region attention with the reduction applied to *key* and *value*. In our case, we employ spatial reduction to efficiently compute channel-to-channel

attention with the reduction applied to *query* and *key*. Once the attention matrix has been computed, it is used to aggregate the channels of *value*, which maintains the original spatial resolution. The final feature \mathbf{Z} from channel-wise cross-attention can be regarded as the selected feature of *value* under the guidance of *query*. Finally, we concatenate \mathbf{X} , \mathbf{Y}_v , and \mathbf{Z} , followed by a linear layer to project the number of channels down to $2C$. Mathematically, the operations of DMCA are summarized below:

$$\begin{aligned} \mathbf{Y}_k, \mathbf{Y}_v &= \text{Split}(W_1(\text{Cat}(\mathbf{Y}_1, \mathbf{Y}_1, \dots, \mathbf{Y}_L))) \\ Q &= W_2(R_1(\mathbf{X})), K = W_3(R_2(\mathbf{Y}_k)), V = W_4(\mathbf{Y}_v) \\ \mathbf{Z} &= \frac{\text{Softmax}(QK^T)}{\sqrt{N/r}}V \end{aligned} \quad (1)$$

where W_i refers to a linear layer used for channel projection, i.e., $W_1 \in \mathbb{R}^{2C \times LC}$ and $\{W_2, W_3, W_4\} \in \mathbb{R}^{C \times C}$. Meanwhile, R_i denotes a spatial reducer implemented as a strided DWConv, while $\{Q, K\} \in \mathbb{R}^{G \times \frac{C}{r} \times \frac{N}{r}}$ and $V \in \mathbb{R}^{G \times \frac{C}{r} \times N}$.

SparX introduces two flexible hyperparameters, namely \mathbf{S} and \mathbf{M} , making it easier and more efficient to extend the network deeper and wider. In addition, the incorporation of DMCA enables strong cross-layer feature interactions, giving rise to more powerful and robust representations.

3.2 Network Architecture

Based on the proposed sparse cross-layer connections, we have developed a new Mamba-based vision backbone named SparX-Mamba, a 4-stage hierarchical architecture with three different variants: tiny, small, and base. In stage 1, we do not include any ganglion layers due to the high computational overhead when large-resolution inputs are used. That is, stage 1 contains only two normal layers. In stage 2, we designate the second layer as a ganglion layer, which has an intra-connection with its preceding layer. Since Stage 3 has more layers, we set different hyperparameters for the three variants. As listed in Table 1, due to the relatively small depth of the tiny model, we set \mathbf{S} to 2 while $\mathbf{S} = 3$ in both small and base models to avoid high computational cost in deeper models. In stage 4, we make an exception for tiny and small models by setting all layers as ganglion layers, as this stage processes the lowest resolution features and incurs a relatively low computational cost in these models. Regarding more complex base model, only the last layer is set as the ganglion layer to prevent excessive computational costs. Meanwhile, to facilitate communications across different stages, the first ganglion layer of each stage connects to a downsampled version of the final feature from the preceding stage. Furthermore, a cross-layer sliding window is

SparX-Mamba	Channels	Blocks	S, M
Tiny	[96, 192, 320, 512]	[2, 2, 7, 2]	[2, 3]
Small	[96, 192, 328, 544]	[2, 2, 17, 2]	[3, 3]
Base	[120, 240, 396, 636]	[2, 2, 21, 3]	[3, 3]

Table 1: Detailed configurations of the three variants of SparX-Mamba.

Method	F (G)	P (M)	Acc.
ConvNeXt-T	4.5	29	82.1
InceptionNeXt-T	4.2	28	82.3
SLaK-T	5.0	30	82.5
UniRepLKNet-T	4.9	31	83.2
MambaOut-T	4.5	27	82.7
Focal-T	4.9	29	82.2
PVTv2-B2	4.0	25	82.0
Swin-T	4.5	29	81.3
CSWin-T	4.5	23	82.7
UniFormer-S	3.6	22	82.9
MambaOut-T	4.5	27	82.7
ViM-S	-	26	81.6
PlainMamba-L2	8.1	25	81.6
ViM2-T	-	20	82.7
EffVMamba-B	4.0	33	81.8
VMamba-T	4.9	31	82.5
LocalVMamba-T	5.7	26	82.7
MSVMamba-T	4.6	33	82.8
SparX-Mamba-T	5.2	27	83.5
ConvNeXt-S	8.7	50	83.1
InceptionNeXt-S	8.4	49	83.5
SLaK-S	9.8	55	83.8
UniRepLKNet-S	9.1	56	83.9
MambaOut-S	9.0	48	84.1
Focal-S	9.4	51	83.6
PVTv2-B4	10.1	63	83.6
Swin-S	8.7	50	83.0
UniFormer-B	8.3	50	83.9
PlainMamba-L3	14.4	50	82.3
ViM2-S	-	43	83.7
VMamba-S	8.7	50	83.6
LocalVMamba-S	11.4	50	83.7
SparX-Mamba-S	9.3	47	84.2
ConvNeXt-B	15.4	89	83.8
InceptionNeXt-B	14.9	87	84.0
SLaK-B	17.1	95	84.0
MambaOut-B	15.8	85	84.2
Focal-B	16.4	90	84.0
PVTv2-B5	11.8	82	83.8
Swin-B	15.4	88	83.5
CSWin-B	15.0	78	84.2
ViM2-B	-	74	83.9
VMamba-B	15.4	89	83.9
SparX-Mamba-B	15.9	84	84.5

Table 2: Performance comparison on ImageNet-1K with 224×224 input resolution. F and P denote the FLOPs and number of Params of a model, respectively.

only applied within each stage, to prevent high memory consumption due to excessive feature storage.

4 Experiments

In this section, we present comprehensive experimental evaluations on representative vision tasks, starting with image classification. Afterwards, the pre-trained models are transferred to downstream tasks, including object detection and semantic segmentation. All experiments were conducted on 8 NVIDIA H800 GPUs. Due to the page limit, we present only part of the results in this section while more experimental results are provided in the arXiv version¹.

4.1 Image Classification

Setup. We use the ImageNet-1K dataset (Deng et al. 2009) and follow the same experimental setting described in DeiT (Touvron et al. 2021) for a fair comparison. Our method is compared with many representative vision models: CNN-based models (ConvNeXt (Liu et al. 2022), InceptionNeXt (Yu et al. 2024b), SLaK (Liu et al. 2023), UniRepLKNet (Ding et al. 2024), and MambaOut (Yu and Wang 2024)), Transformer-based models (Focal-Transformer (Yang et al. 2021), PVTv2 (Wang et al. 2022), Swin (Liu et al. 2021), CSWin (Dong et al. 2022), and UniFormer (Li et al. 2023)), and Mamba-based models (PlainMamba (Yang et al. 2024), ViM2 (Behrouz, Santacatterina, and Zabih 2024), VMamba (Liu et al. 2024), EffVMamba (Pei, Huang, and Xu 2024), LocalVMamba (Huang et al. 2024), and MSVMamba (Shi, Dong, and Xu 2024)).

Results. Table 2 demonstrates the significant advantages of our models over other CNN-, transformer-, and Mamba-based methods. For example, SparX-Mamba-T outperforms ConvNeXt-T/Swin-T by a large margin of 1.4%/2.2%, respectively, in top-1 accuracy. Compared to the recent VMamba-T, our model achieves a 1.0% higher top-1 accuracy with fewer Params and only a slight increase in FLOPs. Furthermore, even with larger models, SparX-Mamba maintains a clear advantage. Meanwhile, we have also provided a more comprehensive comparison with more advanced vision models in the arXiv version¹. In general, our method demonstrates competitive performance in ImageNet-1K classification and exhibits clear advantages on downstream tasks.

4.2 Object Detection and Instance Segmentation

Setup. To evaluate our network architecture on object detection and instance segmentation, we conduct experiments on the COCO 2017 dataset (Lin et al. 2014). We use the Mask R-CNN framework (He et al. 2017) and adopt the same experimental settings as in Swin (Liu et al. 2021). The backbone networks are initially pre-trained on ImageNet-1K and subsequently fine-tuned for 12 epochs ($1 \times$ schedule) as well as 36 epochs ($3 \times$ schedule + multi-scale training).

Results. As shown in Table 3, our models demonstrate superior performance in both object detection and instance segmentation over other models. Specifically, SparX-Mamba-T achieves a 0.7%/0.4% higher AP^b/AP^m than VMamba-T when fine-tuned using the $1 \times$ schedule, and surpasses VMamba-T by 1.3%/1.0% in AP^b/AP^m when fine-tuned using the $3 \times$ schedule. Notably, SparX-Mamba-T even outperforms VMamba-S, which has significantly more

¹arxiv.org/abs/2409.09649

Backbone	F (G) P (M)		$1 \times$ Schedule		$3 \times$ Schedule	
	AP^b	AP^m	AP^b	AP^m	AP^b	AP^m
ConvNeXt-T	262	48	44.2	40.1	46.2	41.7
MambaOut-T	262	43	45.1	41.0	-	-
Focal-T	291	49	-	-	47.2	42.7
PVTv2-B2	309	45	45.3	41.2	47.8	43.1
Swin-T	267	48	42.7	39.3	46.0	41.6
CSWin-T	279	42	46.7	42.2	49.0	43.6
UniFormer-S	269	41	45.6	41.6	48.2	43.4
PlainMamba-L2	542	53	46.0	40.6	-	-
ViM2-T	-	39	47.1	42.4	-	-
EffVMamba-B	252	53	43.7	40.2	45.0	40.8
LocalVMamba-T	291	45	46.7	42.2	48.7	43.4
MSVMamba-T	252	53	46.9	42.2	48.3	43.2
VMamba-T	270	50	47.4	42.7	48.9	43.7
SparX-Mamba-T	279	47	48.1	43.1	50.2	44.7
ConvNeXt-S	348	70	45.4	41.8	47.9	42.9
MambaOut-S	354	65	47.4	42.7	-	-
Focal-S	401	71	-	-	48.8	43.8
PVTv2-B3	397	65	47.0	42.5	48.4	43.2
Swin-S	354	69	44.8	40.9	48.2	43.2
CSWin-S	342	54	47.9	43.2	50.0	44.5
UniFormer-B	399	69	47.4	43.1	50.3	44.8
PlainMamba-L3	696	79	46.8	41.2	-	-
ViM2-S	-	62	48.5	43.1	-	-
LocalVMamba-S	414	69	48.4	43.2	49.9	44.1
VMamba-S	384	70	48.7	43.7	49.9	44.2
SparX-Mamba-S	361	67	49.4	44.1	51.0	45.2
ConvNeXt-B	486	108	47.0	42.7	48.5	43.5
MambaOut-B	495	100	47.4	42.0	-	-
Focal-B	533	110	45.9	-	49.0	43.7
PVTv2-B5	557	102	47.4	42.5	48.4	42.9
Swin-B	496	107	46.9	42.3	48.6	43.3
CSWin-B	526	97	48.7	43.9	50.8	44.9
VMamba-B	485	108	49.2	43.9	-	-
SparX-Mamba-B	498	103	49.7	44.3	51.8	45.8

Table 3: Comparison of object detection and instance segmentation performance on the COCO dataset using Mask R-CNN framework. FLOPs are calculated for the 800×1280 resolution.

Params and FLOPs, using the $3 \times$ schedule. Meanwhile, our method also demonstrates superior performance compared to other excellent vision backbones.

4.3 Semantic Segmentation

Setup. Experiments on semantic segmentation are conducted using the ADE20K dataset (Zhou et al. 2017). We employ two segmentation frameworks, Semantic FPN (SFPN) (Kirillov et al. 2019) and UperNet (Xiao et al. 2018). To ensure fair comparisons, we initialize all backbone net-

Backbone	S-FPN 80K			UperNet 160K		
	F (G)	P (M)	mIoU	F (G)	P (M)	mIoU
ConvNeXt-T	-	-	-	939	60	46.0/46.7
InceptionNeXt-T	-	28	43.1	933	56	47.9/-
SLaK-T	-	-	-	936	65	47.6/-
UniRepLKNet-T	-	-	-	946	61	48.6/49.1
MambaOut-T	-	-	-	938	54	47.4/48.6
PVTv2-B2	167	29	45.2	-	-	-
Swin-T	182	32	41.5	945	60	44.5/45.8
CSWin-T	202	26	48.2	959	59	49.3/50.7
UniFormer-S	247	25	46.6	1008	52	47.6/48.5
PlainMamba	-	-	-	419	81	49.1/-
ViM2-T	-	-	-	-	51	48.6/49.9
EffVMamba-B	-	-	-	930	65	46.5/47.3
LocalVMamba-T	-	-	-	970	57	47.9/49.1
VMamba-T	189	34	47.2 [†]	948	62	48.3/48.6
MSVMamba-T	-	-	-	942	65	47.6/48.5
SparX-Mamba-T	197	31	49.5	954	57	50.0/50.8
ConvNeXt-S	-	-	-	1027	82	48.7/49.6
InceptionNeXt-S	-	50	45.6	1020	78	50.0/-
SLaK-S	-	-	-	1028	91	49.4/-
UniRepLKNet-S	-	-	-	1036	86	50.5/51.0
MambaOut-S	-	-	-	1032	76	49.5/50.6
PVTv2-B4	291	66	47.9	-	-	-
Swin-S	274	53	45.2	1038	81	47.6/49.5
CSWin-S	271	39	49.2	1027	65	50.4/51.5
UniFormer-B	471	54	48.0	1227	80	50.0/50.8
ViM2-S	-	-	-	-	75	50.2/51.4
LocalVMamba-S	-	-	-	1095	81	50.0/51.0
VMamba-S	269	54	49.4 [†]	1038	82	50.6/51.2
SparX-Mamba-S	281	51	50.5	1039	77	51.3/52.5
ConvNeXt-B	-	-	-	1170	122	49.1/49.9
InceptionNeXt-B	-	85	46.4	1159	115	50.6/-
SLaK-B	-	-	-	1172	135	50.2/-
MambaOut-B	-	-	-	1178	112	49.6/51.0
PVTv2-B5	324	91	48.7	-	-	-
Swin-B	422	91	46.0	1188	121	48.1/49.7
CSWin-B	464	81	49.9	1222	109	51.1/52.2
VMamba-B	409	92	49.8 [†]	1170	122	51.0/51.6
SparX-Mamba-B	422	87	51.9	1181	115	52.3/53.4

Table 4: Comparison of semantic segmentation performance on the ADE20K dataset. FLOPs are calculated for the 512×2048 resolution. †: baselines implemented by ourselves. For UperNet, we report both single-scale and multi-scale mIoU.

works with ImageNet-1K pre-trained weights. Furthermore, we strictly adhere to the same training settings as outlined in previous work (Liu et al. 2021; Li et al. 2023).

Results. As shown in Table 4, our SparX-Mamba

Method	ImageNet-1K					ADE20K
	Mem (MB)	T (imgs/s)	F (G)	P (M)	Acc.	mIoU
VMamba-T	6784	1613	4.9	31.0	82.5	47.2
DGC-Mamba-T	7523	860	5.5	28.8	83.4	49.3
DSN-Mamba-T	7641	830	6.0	30.9	83.4	49.5
SparX-Mamba-T	7066	1370	5.2	27.1	83.5	49.5

Table 5: Comparison with DenseNet-like networks. The throughput (T) and memory (Mem) usage are evaluated on a single H800 GPU using a batch size of 128. Due to space constraints, we omit the complexity of the segmentation network. However, it is noteworthy that its complexity is closely aligned with that of the classification network.

demonstrates significant advantages when integrated into both S-FPN and UperNet. For instance, when UperNet is used, SparX-Mamba increases mIoU by 1.7%/2.2% over VMamba-T. Performance improvements achieved with small/base models are also evident. Meanwhile, when using S-FPN, SparX-Mamba-T improves VMamba-T by a notable 2.3% mIoU and outperforms CSWin-T by 1.3% mIoU. These advantages are also evident in small and base models, effectively demonstrating the robustness of our model in dense prediction tasks.

4.4 Ablation Studies

Setup. To assess the effectiveness of individual components in SparX, we conduct comprehensive ablation studies on image classification and semantic segmentation. We first train each model variant on the ImageNet-1K dataset using the same training setting as in Section 4.1. Then, we fine-tune the pre-trained models on the ADE20K dataset, using the S-FPN framework and maintaining an identical training configuration as described in Section 4.3.

Comparison with DenseNet-like networks. Based on SparX-Mamba-T, we conduct the following experiments: (1) To validate our sparse cross-layer connections, we maintain the stride (S) = 2 in SparX, and set up cross-layer connections by removing the sliding window and connecting each ganglion layer with all preceding ganglion and normal layers. This design is denoted as ‘‘Dense Ganglion Connections (DGC-Mamba-T)’’. (2) To verify the superiority of our model over DenseNet (Huang et al. 2017), we set $S=1$ and remove the cross-layer sliding window, thereby designing a network similar to DenseNet. This model is referred to as ‘‘DenseNet-style Network (DSN-Mamba-T)’’.

As shown in Table 5, we observe that both types of dense cross-layer connections (DGC and DSN) incur obvious computational overhead compared to VMamba, particularly in terms of noticeably reduced throughput and increased memory usage, which may hinder scaling to larger models. In contrast, our SparX achieves more pronounced improvements while only slightly reducing throughput and marginally increasing memory usage, demonstrating its high

Method	ImageNet-1K					ADE20K
	Mem (MB)	T (imgs/s)	F (G)	P (M)	Acc.	mIoU
Concat	6861	1413	5.5	27.9	82.8	48.3
w/o CGCA	6853	1425	5.3	26.4	82.7	48.1
CSA	7053	600	5.2	26.0	83.1	49.0
SRA	9333	997	5.3	27.0	83.2	49.2
SCA	12453	1275	5.2	27.9	83.3	49.2
DMCA	7066	1370	5.2	27.1	83.5	49.5

Table 6: Comparison with different cross-layer interaction methods.

efficiency and effectiveness in Mamba-based models.

Comparison with different interaction methods. Based on SparX-Mamba-T, we replace our DMCA with other fusion methods for multi-layer feature interactions: (1) Remove the entire DMCA module and directly concatenate features from the current layer and selected preceding layers, followed by a Conv-FFN layer (Wang et al. 2022) for feature fusion. Using Conv-FFN for feature fusion maintains a comparable level of complexity with the baseline model. This design is referred to as ‘‘Concat’’. (2) Remove the cross-layer grouped channel attention (CGCA) operation in DMCA. That means there is no need to perform spatial reduction and compute the cross-attention matrix between spatially reduced *query* and *key*, but *value* still needs to be computed. Then *query* and *value* are directly concatenated and fused using a Conv-FFN layer. This design is denoted as ‘‘w/o CGCA’’. (3) Replace the attention computation in DMCA with Spatial Reduction Attention (SRA) (Wang et al. 2021, 2022) and Cross-Scale Attention (CSA) (Shang et al. 2023) to compare the performance of channel mixing and spatial mixing. (4) Replace our DMCA with spatial cross-attention (SCA) proposed in FcaFormer (Zhang, Hu, and Wang 2023) to compare the performance of different multi-layer feature aggregation mechanisms. This version is denoted as ‘‘SCA’’.

As listed in Table 6, our DMCA demonstrates better cross-layer feature aggregation and selection thanks to dynamic feature retrieval from preceding layers. Specifically, both ‘‘Concat’’ and ‘‘w/o CGCA’’ result in obvious performance degradation compared with our DMCA. Meanwhile, DMCA outperforms both SRA and SCA possibly because DMCA offers dynamic channel mixing, which is complementary to spatial mixing performed with SSMs. It is noteworthy that some previous works have demonstrated that incorporating both channel and spatial dynamics can lead to improved performance (Fu et al. 2019; Ding et al. 2022a). Conversely, both SRA and SCA function as a spatial mixer similar to SSM, compromising the model’s ability to effectively represent multidimensional features. More importantly, compared with simple feature concatenation, both SRA and SCA incur higher memory costs and lower throughput due to quadratic complexity, whereas our

Method	ImageNet-1K			ADE20K		
	F (G)	P (M)	Acc.	F (G)	P (M)	mIoU
Baseline	5.2	27.1	83.5	197	30.8	49.5
PlainNet	5.4	27.4	82.4 _(-1.1)	199	31.1	47.7 _(-1.8)
S=1	5.7	29.9	83.5 _(+0.0)	206	33.7	49.7 _(+0.2)
S=3	5.1	26.3	83.2 _(-0.3)	194	30.1	49.0 _(-0.5)
S=4	5.1	25.8	82.9 _(-0.6)	193	29.6	48.7 _(-0.8)
M=2	5.2	26.8	83.0 _(-0.5)	197	30.6	48.6 _(-0.9)
M=4	5.5	27.5	83.5 _(+0.0)	198	31.2	49.7 _(+0.2)

Table 7: Ablation study on hyperparameters of the SparX.

Method	ImageNet-1K			ADE20K		
	F (G)	P (M)	Acc.	F (G)	P (M)	mIoU
Baseline	5.2	27.1	83.5	197	30.8	49.5
w/o DPE	5.2	27.0	83.3 _(-0.2)	196	30.7	49.0 _(-0.5)
w/o Skip	5.2	27.0	83.1 _(-0.4)	200	30.8	48.9 _(-0.6)
w/o SR	5.4	27.0	83.4 _(-0.1)	196	30.8	49.3 _(-0.2)

Table 8: Ablation study on DMCA and DPE.

DMCA only marginally increases computational overhead, indicating greater efficiency.

Hyperparameter settings in SparX. We conduct an in-depth exploration of hyperparameter settings in the proposed sparse cross-layer connection mechanism. By using the tiny model as the baseline (i.e., $\mathbf{S}=2$, $\mathbf{M}=3$), we conceive the following model variants: (1) All ganglion layers are replaced with normal layers, i.e., the network no longer has any cross-layer connections. This design is denoted as ‘‘PlainNet’’. For a fair comparison, we make models considered in this comparison have comparable complexity by adjusting the number of channels in stage 3 and stage 4 of PlainNet to 384 and 576, respectively. (2) Keeping $\mathbf{M}=3$ fixed, we vary the stride hyperparameter (\mathbf{S}) controlling the sparsity of ganglion layers to 1, 3, and 4, respectively, to assess the impact of stride on final performance. (3) Keeping $\mathbf{S}=2$ fixed, we investigate the impact of the cross-layer sliding window size by setting \mathbf{M} to 2 and 4, respectively.

Results. As shown in Table 7, having dense ganglion layers (i.e., $\mathbf{S}=1$) does not result in an obvious performance improvement even though the model complexity increases. This indicates that our sparse placement of ganglion layers already facilitates cross-layer communication. However, as \mathbf{S} further increases, model performance gradually declines, suggesting that larger gaps between ganglion layers adversely affects the performance. Note that although $\mathbf{S}=3$ is the second-best option, considering the increased complexity of deeper models, we set $\mathbf{S}=3$ in our small and base versions for higher computational efficiency. Furthermore, experiments demonstrate that the optimal number of preceding ganglion layers in the cross-layer sliding window is 3, as further increasing \mathbf{M} to 4 does not yield any obvious performance gains.

Method	ImageNet-1K			UperNet 160K			Mask R-CNN					
	F (G)	P (M)	Acc.	F (G)	P (M)	mIoU	F (G)	P (M)	$AP_{1\times}^b$	$AP_{1\times}^m$	$AP_{3\times}^b$	$AP_{3\times}^m$
Swin-T	4.5	29	81.3	945	60	44.5	48	267	42.7	39.3	46.0	41.6
SparX-Swin-T	4.7	26	82.6 _(+1.3)	948	57	45.5 _(+1.0)	45	270	44.1 _(+1.4)	40.8 _(+1.5)	47.6 _(+1.6)	43.0 _(+1.4)
Swin-B	15.4	88	83.5	1188	121	48.1	107	496	46.9	42.3	48.6	43.3
SparX-Swin-B	16.2	88	83.9 _(+0.4)	1201	118	48.6 _(+0.5)	107	518	48.1 _(+1.2)	42.9 _(+0.6)	49.1 _(+0.5)	43.7 _(+0.4)

Table 9: Performance Comparison when SparX is applied to Vision Transformers.

Ablations of DMCA and DPE. By using the tiny model as the baseline, we further investigate the effect of DMCA and DPE (Chu et al. 2022) by evaluating the performance of the following alternative model designs. (1) Remove the DPE module, and the resulting model is denoted as “w/o DPE”. (2) Investigate the necessity of concatenating the output of GCCA (Z) with *query* and *value* by only using the output of GCCA without concatenating it with X and Y_v . This design is denoted as “w/o Skip”. (3) Eliminate the spatial reduction (SR) operations, and use the original spatial resolution of *query* and *key* when computing GCCA. This design is denoted as “w/o SR”. According to the results given in Table 8, the use of DPE boosted the top-1 accuracy by 0.2% and mIoU by 0.5%. Then, we can find that “w/o Skip” also gives rise to a performance drop, reflecting the importance of the shortcut connection. Moreover, although spatial reduction (SR) compresses spatial tokens, it actually slightly improves performance. One possible explanation is that utilizing compressed information is sufficient when estimating image-level feature similarities.

4.5 Versatility Analysis

Setup. To demonstrate the versatility of our cross-layer connectivity pattern, we have conducted experiments by incorporating SparX into networks with different token mixer architectures, including self-attention, vanilla SSM in Mamba, and Bi-SSM in ViM. To ensure a fair comparison with other hierarchical networks, we follow the conventional design, where a token mixer is followed by an FFN, with the token

Method	ImageNet-1K			UperNet 160K		
	F (G)	P (M)	Acc.	F (G)	P (M)	mIoU
VMamba-T (SS2D)	4.9	31	82.5	948	62	48.3
SparX-Mamba-T (SSM)	4.8	25	82.9 (+0.4)	946	55	48.1 (-0.2)
SparX-Mamba-T (Bi-SSM)	5.1	25	83.1 (+0.6)	951	55	48.7 (+0.4)
SparX-Mamba-T (SS2D)	5.2	27	83.5 (+1.0)	954	57	50.0 (+1.7)

Table 10: Comparison of image classification and semantic segmentation performance when SparX is applied to various SSM modules.

mixer having the aforementioned multiple options. Specifically, taking SparX-Mamba-T as a basic network, we replace the SS2D block with shifted window attention proposed in Swin Transformer (Liu et al. 2021) to obtain a network denoted as SparX-Swin-T (i.e., Shifted Window Attention \rightarrow FFN), and increase the number of channels in stage 4 from 512 to 576 to align its computational cost with Swin-T. Next, we scale up the model size of SparX-Swin-T by increasing the depths to [2, 2, 27, 3] and channels to [120, 240, 396, 616] to obtain a larger model, denoted as SparX-Swin-B, which has a model complexity comparable to that of Swin-B. Moreover, we also replace the SS2D module in VSS with the vanilla SSM module in original Mamba (Gu and Dao 2024) (i.e., SSM \rightarrow FFN) and Bi-SSM module in ViM (Zhu et al. 2024) (i.e., Bi-SSM \rightarrow FFN). For fair comparisons, we have strictly maintained the similarity of other micro-blocks, such as removing DPE and using identical patch embedding layers.

Results. As shown in Table 9, SparX-Swin-T/B models outperform their vanilla counterparts on both image classification and dense prediction tasks while maintaining similar computational costs. This demonstrates that our method is general and equally applicable to Transformer-based models. Meanwhile, as shown in Table 10, although vanilla SSM and Bi-SSM do not utilize the DWConv employed in SS2D of VMamba, our SparX enables both vanilla SSM and Bi-SSM to outperform VMamba-T in terms of accuracy and computational efficiency. On the semantic segmentation task, the performance of SparX-Mamba-T (SSM) marginally lags behind VMamba-T. This is because causal modeling in vanilla SSM fails to capture sufficient contextual information. When the Bi-SSM module in ViM is used instead, the mIoU improves significantly by 0.4%, demonstrating better performance than VMamba-T.

5 Conclusion

In this work, we propose a new skip-connection strategy named SparX, drawing inspiration from the Retinal Ganglion Cell (RGC) layer in the human visual system. SparX aims to create sparse cross-layer connections to enhance information flow and promote feature distillation and reuse in vision backbone networks. In addition, we propose a Dynamic Multi-layer Channel Aggregator (DMCA) that facilitates dynamic feature aggregation and interaction across layers. Based on SparX, we further propose SparX-Mamba and SparX-Swin, both of which demonstrate superior performance across a range of challenging vision tasks.

Acknowledgements

This work is supported in part by the Hong Kong Research Grants Council under Collaborative Research Fund (Project No. HKU C7004-22G).

References

- Behrouz, A.; Santacatterina, M.; and Zabih, R. 2024. Mambamixer: Efficient selective state space models with dual token and channel selection. *arXiv preprint arXiv:2403.19888*.
- Chen, C.-F. R.; Fan, Q.; and Panda, R. 2021. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 357–366.
- Chen, Y.; Li, J.; Xiao, H.; Jin, X.; Yan, S.; and Feng, J. 2017. Dual path networks. *Advances in Neural Information Processing Systems*, 30.
- Chu, X.; Tian, Z.; Wang, Y.; Zhang, B.; Ren, H.; Wei, X.; Xia, H.; and Shen, C. 2021. Twins: Revisiting the design of spatial attention in vision transformers. *Advances in Neural Information Processing Systems*, 34: 9355–9366.
- Chu, X.; Tian, Z.; Zhang, B.; Wang, X.; and Shen, C. 2022. Conditional Positional Encodings for Vision Transformers. In *International Conference on Learning Representations*.
- Curcio, C. A.; Sloan, K. R.; Kalina, R. E.; and Hendrickson, A. E. 1990. Human photoreceptor topography. *Journal of Comparative Neurology*, 292(4): 497–523.
- Dai, Z.; Liu, H.; Le, Q. V.; and Tan, M. 2021. Coatnet: Marrying convolution and attention for all data sizes. *Advances in Neural Information Processing Systems*, 34: 3965–3977.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE conference on computer vision and pattern recognition*, 248–255.
- Ding, M.; Xiao, B.; Codella, N.; Luo, P.; Wang, J.; and Yuan, L. 2022a. Davit: Dual attention vision transformers. In *European Conference on Computer Vision*, 74–92. Springer.
- Ding, X.; Zhang, X.; Han, J.; and Ding, G. 2022b. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11963–11975.
- Ding, X.; Zhang, Y.; Ge, Y.; Zhao, S.; Song, L.; Yue, X.; and Shan, Y. 2024. Unireplknet: A universal perception large-kernel convnet for audio, video, point cloud, time-series and image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Dong, X.; Bao, J.; Chen, D.; Zhang, W.; Yu, N.; Yuan, L.; Chen, D.; and Guo, B. 2022. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition*, 12124–12134.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- D’Souza, S.; and Lang, R. A. 2020. Retinal ganglion cell interactions shape the developing mammalian visual system. *Development*, 147(23): dev196535.
- Fu, J.; Liu, J.; Tian, H.; Li, Y.; Bao, Y.; Fang, Z.; and Lu, H. 2019. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3146–3154.
- Gu, A.; and Dao, T. 2024. Mamba: Linear-time sequence modeling with selective state spaces. In *Conference on Language Modeling*.
- Guo, J.; Han, K.; Wu, H.; Tang, Y.; Chen, X.; Wang, Y.; and Xu, C. 2022. Cmt: Convolutional neural networks meet vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12175–12185.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Huang, G.; Liu, Z.; Pleiss, G.; Van Der Maaten, L.; and Weinberger, K. Q. 2019. Convolutional networks with dense connectivity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12): 8704–8716.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4700–4708.
- Huang, T.; Pei, X.; You, S.; Wang, F.; Qian, C.; and Xu, C. 2024. Localmamba: Visual state space model with windowed selective scan. *arXiv preprint arXiv:2403.09338*.
- Kim, U. S.; Mahroo, O. A.; Mollon, J. D.; and Yu-Wai-Man, P. 2021. Retinal ganglion cells—diversity of cell types and clinical relevance. *Frontiers in Neurology*, 12: 661938.
- Kirillov, A.; Girshick, R.; He, K.; and Dollár, P. 2019. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6399–6408.
- Lee, Y.; Kim, J.; Willette, J.; and Hwang, S. J. 2022. Mpvit: Multi-path vision transformer for dense prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7287–7296.
- Li, K.; Wang, Y.; Zhang, J.; Gao, P.; Song, G.; Liu, Y.; Li, H.; and Qiao, Y. 2023. Uniformal: Unifying convolution and self-attention for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Li, Y.; Yuan, G.; Wen, Y.; Hu, J.; Evangelidis, G.; Tulyakov, S.; Wang, Y.; and Ren, J. 2022. Efficientformer: Vision transformers at mobilenet speed. *Advances in Neural Information Processing Systems*, 35: 12934–12949.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2117–2125.

- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 740–755. Springer.
- Liu, S.; Chen, T.; Chen, X.; Chen, X.; Xiao, Q.; Wu, B.; Pechenizkiy, M.; Mocanu, D.; and Wang, Z. 2023. More convnets in the 2020s: Scaling up kernels beyond 51x51 using sparsity. In *International Conference on Learning Representations*.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024. Vmamba: Visual state space model. *arXiv preprint arXiv:2401.10166*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11976–11986.
- Lou, M.; Zhou, H.-Y.; Yang, S.; and Yu, Y. 2023. TransXNet: Learning both global and local dynamics with a dual dynamic token mixer for visual recognition. *arXiv preprint arXiv:2310.19380*.
- Pan, X.; Ye, T.; Xia, Z.; Song, S.; and Huang, G. 2023. Slide-Transformer: Hierarchical Vision Transformer with Local Self-Attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Pei, X.; Huang, T.; and Xu, C. 2024. Efficientvmamba: Atrous selective scan for light weight visual mamba. *arXiv preprint arXiv:2403.09977*.
- Pleiss, G.; Chen, D.; Huang, G.; Li, T.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Memory-efficient implementation of densenets. *arXiv preprint arXiv:1707.06990*.
- Ren, S.; Zhou, D.; He, S.; Feng, J.; and Wang, X. 2022. Shunted self-attention via multi-scale token aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10853–10862.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer.
- Shang, L.; Liu, Y.; Lou, Z.; Quan, S.; Adluru, N.; Guan, B.; and Sethares, W. A. 2023. Vision Backbone Enhancement via Multi-Stage Cross-Scale Attention. *arXiv preprint arXiv:2308.05872*.
- Shi, Y.; Dong, M.; and Xu, C. 2024. Multi-Scale VMamba: Hierarchy in Hierarchy Visual State Space Model. *arXiv preprint arXiv:2405.14174*.
- Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, 10347–10357. PMLR.
- Wang, W.; Chen, W.; Qiu, Q.; Chen, L.; Wu, B.; Lin, B.; He, X.; and Liu, W. 2023. Crossformer++: A versatile vision transformer hinging on cross-scale attention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision*, 568–578.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2022. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3): 415–424.
- Watson, A. B. 2014. A formula for human retinal ganglion cell receptive field density as a function of visual field location. *Journal of Vision*, 14(7): 15–15.
- Wu, Y.-H.; Liu, Y.; Zhan, X.; and Cheng, M.-M. 2023. P2T: Pyramid pooling transformer for scene understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *European Conference on Computer Vision*, 418–434.
- Xiao, T.; Singh, M.; Mintun, E.; Darrell, T.; Dollár, P.; and Girshick, R. 2021. Early convolutions help transformers see better. *Advances in Neural Information Processing Systems*, 34: 30392–30400.
- Yang, C.; Chen, Z.; Espinosa, M.; Ericsson, L.; Wang, Z.; Liu, J.; and Crowley, E. J. 2024. Plainmamba: Improving non-hierarchical mamba in visual recognition. *arXiv preprint arXiv:2403.17695*.
- Yang, J.; Li, C.; Zhang, P.; Dai, X.; Xiao, B.; Yuan, L.; and Gao, J. 2021. Focal attention for long-range interactions in vision transformers. *Advances in Neural Information Processing Systems*, 34: 30008–30022.
- Yu, W.; Si, C.; Zhou, P.; Luo, M.; Zhou, Y.; Feng, J.; Yan, S.; and Wang, X. 2024a. MetaFormer Baselines for Vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(2): 896–912.
- Yu, W.; and Wang, X. 2024. MambaOut: Do We Really Need Mamba for Vision? *arXiv preprint arXiv:2405.07992*.
- Yu, W.; Zhou, P.; Yan, S.; and Wang, X. 2024b. Inception-next: When inception meets convnext. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhang, H.; Hu, W.; and Wang, X. 2023. Fcaformer: Forward cross attention in hybrid vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6060–6069.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 633–641.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. In *International Conference on Machine Learning*.