

# Semantic Equitable Clustering: A Simple and Effective Strategy for Clustering Vision Tokens

Qihang Fan<sup>1,2</sup>, Huaibo Huang<sup>1\*</sup>, Mingrui Chen<sup>1,2</sup>, Ran He<sup>1,2</sup>

<sup>1</sup>MAIS & NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing, China

<sup>2</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

fanqihang.159@gmail.com, huaibo.huang@cripac.ia.ac.cn,  
charmier@hust.edu.cn, rhe@nlpr.ia.ac.cn

## Abstract

The Vision Transformer (ViT) has gained prominence for its superior relational modeling prowess. However, its global attention mechanism’s quadratic complexity poses substantial computational burdens. A common remedy spatially groups tokens for self-attention, reducing computational requirements. Nonetheless, this strategy neglects semantic information in tokens, possibly scattering semantically-linked tokens across distinct groups, thus compromising the efficacy of self-attention intended for modeling inter-token dependencies. Motivated by these insights, we introduce a fast and balanced clustering method, named **Semantic Equitable Clustering (SEC)**. SEC clusters tokens based on their global semantic relevance in an efficient, straightforward manner. In contrast to traditional clustering methods requiring multiple iterations, our method achieves token clustering in a single pass. Additionally, SEC regulates the number of tokens per cluster, ensuring a balanced distribution for effective parallel processing on current computational platforms without necessitating further optimization. Capitalizing on SEC, we propose a versatile vision backbone, **SECViT**. Comprehensive experiments in image classification, object detection, instance segmentation, and semantic segmentation validate the effectiveness of **SECViT**. Moreover, SEC can be conveniently and swiftly applied to multimodal large language models (MLLM), such as LLaVA, to serve as a vision language connector, effectively accelerating the model’s efficiency while maintaining unchanged or better performance.

## 1. Introduction

Since its inception, the Vision Transformer (ViT)[11] has drawn considerable interest from the research community due to its robust modeling prowess. However, the quadratic

\*Huaibo Huang is the corresponding author.

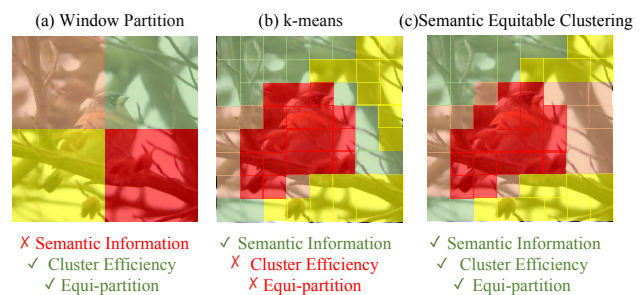


Figure 1. Comparison among Window Partition, Dynamic Group by k-means, and Semantic Equitable Clustering. Our Semantic Equitable Clustering incorporates image semantics while maintaining efficient clustering, eliminating the need for iterative processes such as in k-means. Furthermore, it enables equi-partitioning of tokens, promoting efficient GPU processing without necessitating additional CUDA optimization.

complexity of Self-Attention leads to significant computational overhead, thus constraining the practicality of ViT. A variety of strategies have been devised to alleviate this computational load, the most prevalent of which involves token grouping, thereby constraining the attention span of each token[10, 38, 44, 47].

Specifically, the Swin-Transformer [38] partitions tokens into multiple small windows, restricting token attention within each window. The CSWin-Transformer [10] adopts a cross-shaped grouping, endowing each token with a global receptive field. MaxViT [44] amalgamates window and grid attention, facilitating intra-window tokens to attend to their counterparts in other windows. However, these methods, solely reliant on spatial positioning, neglect token semantics, potentially restricting the self-attention’s capacity to model semantic dependencies. To mitigate this, DGT [37] employs k-means clustering for query grouping, considering the semantic information of tokens for enhanced feature learning. Nonetheless, the iterative nature of k-means clustering and the potential for uneven token counts per cluster can impact the efficiency of parallel attention operations.




Figure 2. **Left:** Top-1 accuracy v.s. FLOPs on ImageNet-1K of recent SOTA models. **Right:** Comparison among different vision language connectors on LLaVA-1.5

Given these considerations, an optimal token partitioning scheme should efficiently segregate tokens, incorporate semantic information, and efficiently utilize computational resources (e.g., GPU). In response, we introduce a simple, fast and equitable clustering approach named Semantic Equitable Clustering (SEC). SEC segments tokens based on their relevance to global semantic information. Specifically, we employ global pooling to generate a global token encapsulating global semantic information. The similarity between this global token and all other tokens is then computed, reflecting global semantic relevance. Upon obtaining the similarity matrix, tokens (excluding the global token) are sorted by similarity scores, and the tokens with similar scores are grouped into clusters, ensuring uniform token distribution across clusters. As depicted in Fig. 1, SEC comprehensively considers token semantics and completes the clustering process in a single iteration, unlike the multi-iteration k-means. The resulting clusters, containing an equal number of tokens, can be processed in parallel by the GPU efficiently.

Building upon Semantic Equitable Clustering (SEC), we introduce the Semantic Equitable Clustering Vision Transformer (SECViT), a versatile vision backbone that is adaptable to a wide spectrum of downstream tasks. As demonstrated in Fig. 2, SECViT exhibits significant performance improvements compared to previous state-of-the-art (SOTA) models. Impressively, SECViT attains an accuracy of **84.3%** utilizing merely **4.6GFLOPs**, without the need for additional training data or supervision. This superior performance is maintained across different model scales. Furthermore, SECViT proves its proficiency in downstream tasks, including but not limited to, object detection, instance segmentation, and semantic segmentation.

Beyond vision tasks, we also apply SEC to multimodal large language models (MLLM) such as LLaVA-1.5 [36] to


serve as an efficient vision language connector. Specifically, we use SEC to cluster the vision tokens, and then merge all the tokens at corresponding positions within each cluster into a single token. Experiments demonstrate that this approach significantly enhances the efficiency of LLaVA-1.5 while improving the model’s performance.

## 2. Related Works

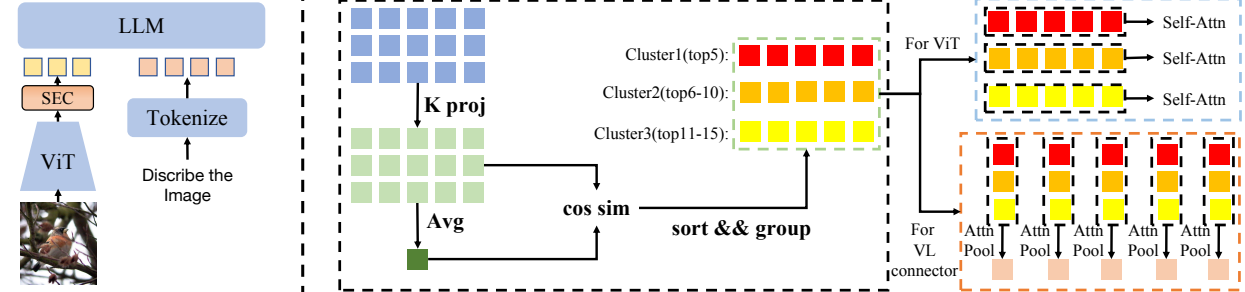
**Vision Transformer.** The Vision Transformer (ViT) [11] is considered a powerful visual architecture. Many works have improved the Vision Transformer, including enhancing its training efficiency and reducing its computational cost [10, 24, 38, 43, 61]. DeiT [43] uses distillation loss and incorporates extensive data augmentation methods into the ViT training process. Hierarchical structures represented by PVT [16, 41, 45, 46, 51] reduce the number of tokens in global attention by downsampling the keys and values (KV), thereby low the computational cost. In addition to them, some methods directly prune tokens based on their importance, retaining important tokens [32, 40]. This reduces the number of tokens and subsequently lowers the computational cost of the model. Another highly representative approach is to group all tokens such that each token can only attend to tokens within its own group [9, 10, 37, 38, 61]. This method also significantly reduces the computational cost of self-attention.

**Grouping-Based Vision Transformer.** Most grouping-based attention mechanisms perform grouping based on spatial structure [9, 10, 37, 38, 44]. Specifically, the Swin-Transformer [38] divides all tokens into equally sized windows based on their spatial positions, where each token can only attend to tokens within its own window. This significantly reduces the model’s computational cost. In addition to dividing tokens into small windows along the spatial dimension, DaViT [9] also splits channels into multiple groups along the channel dimension. Unlike the above methods that only consider positional information for grouping, DGT [37] takes semantic information into account by using k-means clustering to group the queries.

**Vision Language Connector.** The vision language connector is a critical component in MLLMs [3, 23, 36]. It aligns vision tokens with language tokens. Typical vision language connectors include MLP [36], Resampler [1], C-Abstractor [3], and others. Although MLP performs well, it introduces a significant number of vision tokens, which hampers the model’s efficiency. On the other hand, connectors like Resampler improve the model’s efficiency, but at the cost of reduced performance. Unlike these methods, our proposed SEC consider the semantic information of each token and significantly enhances the model’s efficiency while maintaining its performance.



(a) Illustration of SECViT



(b) SEC for vision language connector

(c) Illustration of Semantic Equitable Clustering

Figure 3. (a) Illustration of SECViT (b) Applying SEC to vision language connector. (c) Illustration of Semantic Equitable Clustering for ViT and vision language connector.

### 3. Method

#### 3.1. Overall Architecture

The overall architecture of SECViT is shown in Fig. 3(a). SECViT consists of four stages with downsampling factors of  $\frac{1}{4}$ ,  $\frac{1}{8}$ ,  $\frac{1}{16}$ , and  $\frac{1}{32}$ , respectively. This structural design facilitates downstream tasks, such as object detection, in constructing feature pyramids. A SECViT block is composed of three modules. For each block, the input tensor  $X_{in} \in \mathbb{R}^{C \times H \times W}$  is fed into the CPE to introduce the positional information. Then, The Self-Attention based on the Semantic Equitable Clustering (SEC) is employed to serve as the token mixer. The final FFN is utilized to integrate channel-wise information of tokens.

Beyond the design of the backbone, we also utilize SEC in the design of the vision language connector in MLLM [36]. For the vision tokens output by ViT, we use SEC to cluster the tokens. For each position corresponding to a cluster, we use attentive pooling to merge them into a single token, thereby reducing the number of vision tokens. The process is shown in Fig. 3(b).

#### 3.2. Semantic Equitable Clustering

As previously mentioned, the design objectives of Semantic Equitable Clustering are threefold: **1)** Fully consider the semantic information contained in different tokens during clustering. **2)** Unlike k-means and other clustering methods that require multiple iterations, Semantic Equitable Cluster-

ing can complete clustering in a single step. **3)** Ensure an equal number of tokens in each cluster to facilitate parallel processing on GPUs. In the following paragraphs, we will describe in detail how our Semantic Equitable Clustering achieves these three objectives. And the whole process is illustrated in the Fig. 3(c).

#### Single Clustering Center Related to Semantics.

k-means is relatively complex for two reasons. **First**, it has multiple cluster centers, and each token needs to calculate its distance to each cluster center to determine its cluster membership. **Second**, the determination of each cluster center in k-means is not precise and requires multiple iterations to accurately establish the cluster centers.

To address these two issues, we first discard the use of multiple cluster centers and instead calculate the distance between each token and a single center. Based on each token’s distance to this center, we divide the tokens into different intervals. Then, to ensure that our chosen center contains the most comprehensive semantic information, we directly use the result of average pooling of all tokens as the center token. This is because, in most vision foundation models, the output of the average pool is assumed to contain the richest semantic information and is thus used for classification [6, 10, 12, 38]. Specifically, the process for determining the cluster center is shown in Eq. 1:

$$Q = W_Q X, K = W_K X, V = W_V X, \quad (1)$$

$$k_c = \text{Pool}(K).$$

Where  $W_K$  is a learnable matrix.  $k_c$  is the determined cluster center.  $X$  is the set of input tokens.

**Distance Metric Suitable for ViT.** Unlike the Euclidean distance calculation used in the k-means algorithm for computing the distance between tokens, during the actual computation of Self-Attention, similarity between query and key is computed through dot product. To better adapt to the characteristics of Self-Attention, we also measure the distance between tokens using a method similar to dot product. Specifically, we calculate the cosine similarity between the cluster center and each token, and then sort the tokens according to the magnitude of the computed results. The specific process is shown in Eq. 2:

$$\begin{aligned} sim &= \frac{K \cdot k_c}{\|K\| \cdot \|k_c\|}, \\ idx &= \text{argsort}(sim), \\ Q^* &= Q[idx], K^* = K[idx], V^* = V[idx]. \end{aligned} \quad (2)$$

Where  $sim$  is the similarity matrix between  $K$  and  $k_c$ , the  $\text{argsort}(sim)$  returns the indices of  $sim$  sorted in descending order.  $Q^*, K^*, V^*$  are  $Q, K, V$  rearranged according to  $\text{argsort}(sim)$ .

**Equally Partition Tokens based on Distance.** The obtained  $Q^*, K^*$ , and  $V^*$  from the previous step have been sorted based on their distances to the cluster center. **For the design of vision backbone**, we directly group them, so tokens with similar distances to the cluster center are classified into the same cluster. This allows us to directly control an equal number of tokens in each cluster. This process can be clearly illustrated in Fig. 3(c) and denoted as follows:

$$\begin{aligned} Q_m &= Q^*[m \times N : (m + 1) N], \\ K_m &= K^*[m \times N : (m + 1) N], \\ V_m &= V^*[m \times N : (m + 1) N]. \end{aligned} \quad (3)$$

where  $N$  is the basic token number of each cluster for equal partition and  $m$  is the index of the cluster

Based on the above steps, we have completed the clustering process that captures semantic information in the image with minimal sorting cost. Moreover, compared to k-means, we have achieved equi-partitioning of each cluster. After clustering is completed, we apply standard Self-Attention to the tokens within each cluster, thereby completing the interaction of information between tokens:

$$Y_m = \text{Attn}(Q_m, K_m, V_m). \quad (4)$$

**For the design of vision language connector**, we group the tokens according to their similarity, and the tokens

within each group are interleaved, as shown in Eq. 5:

$$\begin{aligned} Q_n &= Q^*[n : N : L], \\ K_n &= K^*[n : N : L], \\ V_n &= V^*[n : N : L]. \end{aligned} \quad (5)$$

in which  $L$  is the token’s sequence length,  $n$  is the index of group tokens.  $N$  is the basic token number of each cluster. After obtaining the token groups, we perform pooling on  $Q$  to effectively reduce the number of tokens input to the LLM, with each group’s output becoming a single token, as shown in Eq 6.

$$Y_n = \text{Attn}(\text{Pool}(Q_n), K_n, V_n). \quad (6)$$

### 3.3. Difference between SEC and EViT.

We use the most representative example, EViT [32], to illustrate the differences between SEC and other methods based on the similarity between the global token and other tokens.

**Pruning v.s. Clustering.** Most similarity-based methods, such as EViT, are pruning methods, where tokens with low similarity to the [cls] token are merged during the forward process, thereby reducing the number of tokens and decreasing computational cost. In contrast, our proposed SECViT employs a clustering-based approach, performing attention operations within each cluster.

**The role of the [cls] token.** In methods like EViT, the [cls] token serves as a measure of importance of a token. Each token computes its similarity to the [cls] token, with higher similarity tokens deemed more important. The less important tokens are abandoned. In contrast, in SEC, the [cls] token (obtained by average pooling over all tokens) measures similarity between tokens. Each token computes its similarity score to the [cls] token; tokens with similar scores are considered to be more similar and grouped into one cluster. Attention is calculated only within the same cluster.

## 4. Experiments

We first make strict comparison with hierarchical/plain baselines. Then we conduct experiments on a wide range of vision tasks for SECViT, including image classification, object detection, instance segmentation, and semantic segmentation. We also verify the role of SEC in MLLM based on LLaVA-1.5 [36]. **More details, experiments, and comparison of models’ efficiency** can be found in the **Appendix**.

### 4.1. SEC for vision models

**Strict Comparison with Baselines.** We select two baselines: hierarchical backbone Swin-Transformer [38] and plain backbone DeiT [43] to make a comparison with our SEC based model. In the comparison models (SEC-Swin

Model	Params (M)	FLOPs (G)	Throughput (imgs/s)	Acc	$AP^b$	$AP^m$	mIoU
DeiT-S [43]	22	4.6	3204	79.8	44.5	40.1	43.0
EViT-DeiT-S (keep-rate=0.9)	22	4.0	3428	79.8	not suit	not suit	not suit
SEC-DeiT-S (num_cluster=4)	22	4.1	3412	80.5 (+0.7)	47.7 (+3.2)	42.7 (+2.6)	47.5 (+4.5)
DeiT-B	86	17.6	1502	81.8	-	-	-
SEC-DeiT-B	86	14.8	1682	82.4 (+0.6)	-	-	-
Swin-T	29	4.5	1723	81.3	43.7	39.8	44.5
SEC-Swin-T	29	4.8	1482	83.8 (+2.5)	48.3 (+4.6)	43.4 (+3.6)	49.3 (+4.8)
Swin-S	50	8.8	1006	83.0	45.7	41.1	47.6
SEC-Swin-S	50	9.2	804	85.0 (+2.0)	50.2 (+4.5)	44.7 (+3.6)	51.3 (+3.7)

Table 1. Comparison with Hierarchy/Plain baselines. Inference speed are measured on the A100 GPU.

Model	Params (M)	FLOPs (G)	Method	Pretrain epoch	Acc(%)
Swin-B [38]	88	15.4	Supervised	-	83.5
ConvNeXt V2-B [50]	88	15.4	Supervised	-	84.3
SEC-Swin-B	88	16.2	Supervised	-	85.3
Swin-B [38]	88	15.4	SimMIM [53]	800	84.0(+0.5)
ConvNeXt V2-B [50]	88	15.4	FCMAE [50]	800	84.6(+0.3)
SEC-Swin-B	88	16.2	SimMIM [53]	800	85.9(+0.6)

Table 2. Comparison with baselines on self-supervised setting.

and SEC-DeiT), we merely substitute the attention mechanism in the original model with our SEC based Self-Attention and without introducing any other modules. As shown in Tab. 1, we conduct experiments on image classification, object detection, instance segmentation and semantic segmentation, the simple replacement of the attention mechanism yields significant advantages in both performance and efficiency.

In addition to the supervised scenario, we also train the model with SimMIM [53] in the self-supervised scenario. As shown in Tab. 2, SEC also performs exceptionally well in the self-supervised scenario.

**Image Classification.** We compare our SECViT with numerous state-of-the-art models, the results are shown in Tab.3. We adopt the training strategy proposed in DeiT [43], with the only supervision is cross entropy loss. All of our models are trained from scratch for 300 epochs with the input resolution of  $224 \times 224$ . SECViT consistently outperforms preceding models across all scales. Notably, SECViT-S attains a Top1-accuracy of **84.3%** with a mere **27M** parameters and **4.6G** FLOPs. The comparison of the models’ efficiency can be found in Appendix.

**Object Detection and Instance Segmentation.** We utilize MMDetection [4] to implement Mask-RCNN [21], Cascade Mask R-CNN [2], and RetinaNet [33] to evaluate the performance of the SECViT. Tab. 5 and Tab. 4 show the

Cost	Model	Params (M)	FLOPs (G)	Top1-acc (%)
tiny model $\sim 2.5G$	PVTv2-b1 [46]	13	2.1	78.7
	TCFormer-light [59]	14	3.8	79.4
	QuadTree-B-b1 [42]	14	2.3	80.0
	MPViT-XS [28]	11	2.9	80.9
	BiFormer-T [61]	13	2.2	81.4
	CrossFormer-T [47]	28	2.9	81.5
	FAT-B2 [12]	14	2.0	81.9
	GC-ViT-XT [20]	20	2.6	82.0
	SMT-T [34]	12	2.4	82.2
	RMT-T [13]	14	2.5	82.4
	SECViT-T	15	2.5	<b>82.7</b>
	small model $\sim 4.5G$	PS-ViT-B14 [58]	21	5.4
DVT-T2T-ViT-19 [49]		39	6.2	81.9
ConvNeXt-T [39]		29	4.5	82.1
TCFormer [59]		26	5.8	82.3
SG-Former-S [15]		23	4.8	83.2
StructViT-S-8-1 [25]		24	5.4	83.3
InternImage-T [48]		30	5.0	83.5
MLLA-T [18]		25	4.2	83.5
MaxViT-T [44]		31	5.6	83.6
FAT-B3 [12]		29	4.4	83.6
SMT-S [34]		20	4.8	83.7
BiFormer-S [61]		26	4.5	83.8
SECViT-S	27	4.6	<b>84.3</b>	
base model $\sim 9.0G$	ConvNeXt-S [39]	50	8.7	83.1
	NAT-S [19]	51	7.8	83.7
	Quadtree-B-b4 [42]	64	11.5	84.0
	MOAT-1 [54]	42	9.1	84.2
	InternImage-S [48]	50	8.0	84.2
	GC-ViT-S [20]	51	8.5	84.3
	BiFormer-B [61]	57	9.8	84.3
	iFormer-B [41]	48	9.4	84.6
	FAT-B4 [12]	52	9.3	84.8
	SECViT-B	57	9.8	<b>85.2</b>
large model $\sim 18.0G$	CrossFormer-L [47]	92	16.1	84.0
	SMT-L [34]	81	17.7	84.6
	DaViT-B [9]	88	15.5	84.6
	SG-Former-B [15]	78	15.6	84.7
	iFormer-L [41]	87	14.0	84.8
	InterImage-B [48]	97	16.0	84.9
	GC-ViT-B [20]	90	14.8	85.0
	RMT-L [13]	95	18.2	85.5
SECViT-L	101	18.2	<b>85.7</b>	

Table 3. Comparison with the state-of-the-art on ImageNet-1K classification.

results of SECViT with different detection frameworks. The results show that SECViT performs better than its counterparts in all comparisons.

**Semantic Segmentation.** We utilize Semantic FPN [26] and UperNet [52] to validate our SECViT’s performance, implementing these frameworks via MMSegmentation [7]. The results of semantic segmentation can be found in the Tab. 6. All the FLOPs are measured with the input resolution of  $512 \times 2048$ . SECViT achieves the best performance in all settings.

## 4.2. SEC for MLLM

SEC can greatly facilitate the design of vision language connectors in MLLMs. First, we conduct a rigorous compari-

Backbone	Params (M)	FLOPs (G)	Mask R-CNN 1×						Params (M)	FLOPs (G)	RetinaNet 1×					
			$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP^m$	$AP_{50}^m$	$AP_{75}^m$			$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP_S^b$	$AP_M^b$	$AP_L^b$
PVTv2-B1 [46]	33	243	41.8	54.3	45.9	38.8	61.2	41.6	23	225	41.2	61.9	43.9	25.4	44.5	54.3
FAT-B2 [12]	33	215	45.2	67.9	49.0	41.3	64.6	44.0	23	196	44.0	65.2	47.2	27.5	47.7	58.8
RMT-T [13]	33	218	47.1	68.8	51.7	42.6	65.8	45.9	23	199	45.1	66.2	48.1	28.8	48.9	61.1
SECViT-T	34	221	<b>47.8</b>	<b>69.5</b>	<b>52.5</b>	<b>43.0</b>	<b>66.7</b>	<b>46.3</b>	24	202	<b>45.8</b>	<b>66.8</b>	<b>49.2</b>	<b>29.1</b>	<b>49.8</b>	<b>60.9</b>
MPViT-S [28]	43	268	46.4	68.6	51.2	42.4	65.6	45.7	32	248	45.7	57.3	48.8	28.7	49.7	59.2
MLLA-T [18]	44	255	46.8	69.5	51.5	42.1	66.4	45.0	-	-	-	-	-	-	-	-
STViT-S [22]	44	252	47.6	70.0	52.3	43.1	66.8	46.5	-	-	-	-	-	-	-	-
RMT-S [13]	46	262	49.0	70.8	53.9	43.9	67.8	47.4	36	244	47.8	69.1	51.8	32.1	51.8	63.5
SECViT-S	45	262	<b>49.9</b>	<b>70.9</b>	<b>54.7</b>	<b>44.6</b>	<b>68.3</b>	<b>47.7</b>	35	240	<b>48.4</b>	<b>69.4</b>	<b>52.0</b>	<b>31.3</b>	<b>53.3</b>	<b>63.8</b>
ScalableViT-B [56]	95	349	46.8	68.7	51.5	42.5	65.8	45.9	85	330	45.8	67.3	49.2	29.9	49.5	61.0
InternImage-S [48]	69	340	47.8	69.8	52.8	43.3	67.1	46.7	-	-	-	-	-	-	-	-
MLLA-S [18]	63	319	49.2	71.5	53.9	44.2	68.5	47.2	-	-	-	-	-	-	-	-
STViT-B [22]	70	359	49.7	71.7	54.7	44.8	68.9	48.7	-	-	-	-	-	-	-	-
SECViT-B	76	371	<b>51.5</b>	<b>72.9</b>	<b>56.7</b>	<b>45.4</b>	<b>69.9</b>	<b>48.7</b>	63	349	<b>49.3</b>	<b>70.3</b>	<b>52.9</b>	<b>32.0</b>	<b>53.8</b>	<b>64.8</b>
Focal-B [55]	110	533	47.8	70.2	52.5	43.2	67.3	46.5	101	514	46.3	68.0	49.8	31.7	50.4	60.8
CSwin-B [10]	97	526	48.7	70.4	53.9	43.9	67.8	47.3	-	-	-	-	-	-	-	-
InternImage-B [48]	115	501	48.8	70.9	54.0	44.0	67.8	47.4	-	-	-	-	-	-	-	-
MLLA-B [18]	115	502	50.5	72.0	55.4	45.0	69.3	48.6	-	-	-	-	-	-	-	-
SECViT-L	119	550	<b>52.0</b>	<b>73.5</b>	<b>57.3</b>	<b>46.3</b>	<b>70.6</b>	<b>49.8</b>	105	527	<b>50.2</b>	<b>71.4</b>	<b>53.9</b>	<b>33.2</b>	<b>54.5</b>	<b>66.3</b>

Table 4. Comparison to other backbones using "1×" schedule on COCO.

Backbone	Params (M)	FLOPs (G)	$AP^b$	$AP_{50}^b$	$AP_{75}^b$	$AP^m$	$AP_{50}^m$	$AP_{75}^m$
Mask R-CNN 3×+MS								
GC-ViT-T [20]	48	291	47.9	70.1	52.8	43.2	67.0	46.7
MLLA-T [18]	44	255	48.8	71.0	53.6	43.8	68.0	46.8
SMT-S [34]	40	265	49.0	70.1	53.4	43.4	67.3	46.7
InternImage-T [48]	49	270	49.1	70.4	54.1	43.7	67.3	47.3
RMT-S [13]	46	262	50.7	71.9	55.6	44.9	69.1	48.4
SECViT-S	45	262	<b>51.6</b>	<b>72.5</b>	<b>55.9</b>	<b>45.6</b>	<b>69.9</b>	<b>48.8</b>
NAT-S [19]	70	330	48.4	69.8	53.2	43.2	66.9	46.4
InternImage-S [48]	69	340	49.7	71.1	54.5	44.5	68.5	47.8
SMT-B [34]	52	328	49.8	71.0	54.4	44.0	68.0	47.3
MLLA-S [18]	63	319	50.5	71.8	55.2	44.9	69.1	48.2
RMT-B [13]	73	373	52.2	72.9	57.0	46.1	70.4	49.9
SECViT-B	75	371	<b>52.8</b>	<b>73.6</b>	<b>57.7</b>	<b>46.4</b>	<b>70.8</b>	<b>49.9</b>
Cascade Mask R-CNN 3×+MS								
GC-ViT-T [20]	85	770	51.6	70.4	56.1	44.6	67.8	48.3
SMT-S [34]	78	744	51.9	70.5	56.3	44.7	67.8	48.6
UniFormer-S [30]	79	747	52.1	71.1	56.6	45.2	68.3	48.9
RMT-S [13]	83	741	53.2	72.0	57.8	46.1	69.8	49.8
SECViT-S	83	741	<b>54.1</b>	<b>72.8</b>	<b>58.6</b>	<b>47.0</b>	<b>70.3</b>	<b>51.0</b>
NAT-S [19]	108	809	51.9	70.4	56.2	44.9	68.2	48.6
GC-ViT-S [20]	108	866	52.4	71.0	57.1	45.4	68.5	49.3
CSwin-S [10]	92	820	53.7	72.2	58.4	46.4	69.6	50.6
UniFormer-B [30]	107	878	53.8	72.8	58.5	46.4	69.9	50.4
RMT-B [13]	111	852	54.5	72.8	59.0	47.2	70.5	51.4
SECViT-B	114	849	<b>55.4</b>	<b>74.1</b>	<b>59.9</b>	<b>47.8</b>	<b>71.7</b>	<b>51.7</b>

Table 5. Comparison with other backbones using "3×+MS" schedule on COCO.

son between SEC and various baseline vision language connectors based on LLaVA-1.5. Then, we compare LLaVA-1.5+SEC with several popular contemporary MLLMs.

**Strict Comparison with Baselines.** In Tab. 7, we strictly compare various commonly used vision language connectors, including MLP, Resampler [1], Pooling, and EViT [32], which has achieved success in the design of ViT. Among these, MLP is the original design in LLaVA-1.5 [36], capable of achieving good results. However, it incurs significant computational cost due to the excessive vision tokens. To address this issue, some connectors at-

Model	Semantic FPN 80K			Upernet 160K		
	Params (M)	FLOPs (G)	mIoU (%)	Params (M)	FLOPs (G)	mIoU <sub>ss</sub> (%)
PVTv2-B1 [46]	18	136	42.5	-	-	-
VAN-B1 [17]	18	140	42.9	-	-	-
RMT-T [13]	17	136	46.4	-	-	-
SECViT-T	18	136	<b>47.2</b>	44	894	<b>48.8</b>
StructViT-S [25]	26	271	46.9	-	-	-
MogaNet-S [31]	29	189	47.7	55	946	49.2
SMT-S [34]	-	-	-	50	935	49.2
SGFormer-S [15]	25	205	49.0	52.5	989	49.9
RMT-S [13]	30	180	49.4	56	937	49.8
SECViT-S	30	180	<b>49.6</b>	56	936	<b>50.6</b>
MogaNet-B [31]	-	-	-	74	1050	50.1
InternImage-S [48]	-	-	-	80	1017	50.2
StructViT-B [25]	54	529	48.5	-	-	-
RMT-B [13]	57	294	50.4	83	1051	52.0
SECViT-B	60	291	<b>50.7</b>	86	1048	<b>52.2</b>
MogaNet-L [31]	-	-	-	113	1176	50.9
MLLA-B [18]	-	-	-	128	1183	51.9
SGFormer-B [15]	81	475	50.6	109	1304	52.0
RMT-L [13]	98	482	51.4	125	1241	52.8
SECViT-L	103	475	<b>52.2</b>	131	1256	<b>53.8</b>

Table 6. Comparison with the state-of-the-art on ADE20K.

tempt to use fewer vision tokens to accelerate LLaVA-1.5. Nonetheless, these adjustments inevitably lead to performance degradation. The results in Tab. 7 show that using SEC can effectively accelerate the inference of LLaVA-1.5 without causing performance degradation, and can even improve the performance of LLaVA-1.5 to a certain extent.

**Comparison with Popular MLLMs.** In Tab. 8 and Tab. 9, we compare LLaVA-1.5 equipped with SEC as a vision-language connector with other MLLMs. It is evident that SEC not only enhances the performance of MLLMs across various benchmarks but also significantly improves the efficiency of the models. This fully demonstrates the

Model	Connector	V-T Num	Time	Speed	TextVQA	GQA	VQAv2	POPE	MME
LLaVA-1.5	MLP	576+1	194s	1.0×	58.2	62.0	78.5	86.1	1510.7
LLaVA-1.5+Resampler	Resampler	288+1	126s	1.5×	52.1	56.8	76.0	83.1	1393.2
LLaVA-1.5+EViT	MLP+EViT	288+1	126s	1.5×	54.6	60.0	77.9	84.3	1483.2
LLaVA-1.5+SEC	MLP+SEC	288+1	126s	1.5×	<b>60.1</b>	<b>63.5</b>	<b>78.9</b>	<b>87.7</b>	<b>1510.7</b>
LLaVA-1.5+Resampler	Resampler	256+1	116s	1.7×	51.6	56.0	75.2	82.7	1387.2
LLaVA-1.5+Pool	MLP+Pool	256+1	116s	1.7×	52.4	57.6	76.4	83.3	1415.5
LLaVA-1.5+EViT	MLP+EViT	256+1	116s	1.7×	52.8	59.6	77.1	83.7	1443.7
LLaVA-1.5+SEC	MLP+SEC	256+1	116s	1.7×	<b>59.6</b>	<b>63.2</b>	<b>78.6</b>	<b>87.1</b>	<b>1505.2</b>
LLaVA-1.5+Resampler	Resampler	192+1	102s	1.9×	50.1	55.2	74.3	82.7	1337.6
LLaVA-1.5+EViT	MLP+EViT	192+1	102s	1.9×	51.6	58.6	76.3	83.1	1427.6
LLaVA-1.5+SEC	MLP+SEC	192+1	102s	1.9×	<b>57.7</b>	<b>62.7</b>	<b>78.4</b>	<b>86.7</b>	<b>1500.1</b>
LLaVA-1.5+Resampler	Resampler	144+1	94s	2.1×	47.6	54.6	72.0	81.9	1293.7
LLaVA-1.5+Pool	MLP+Pool	144+1	94s	2.1×	50.0	56.2	73.6	81.9	1310.7
LLaVA-1.5+EViT	MLP+EViT	144+1	94s	2.1×	51.2	58.0	76.0	83.1	1393.6
LLaVA-1.5+SEC	MLP+SEC	144+1	94s	2.1×	<b>56.8</b>	<b>62.0</b>	<b>78.0</b>	<b>86.1</b>	<b>1487.1</b>

Table 7. Comparison of different vision language connectors on LLaVA-1.5. “V-T Num” denotes the quantity of visual tokens. The computation expense is impacted by V-T Num, with larger values resulting in higher costs. “Speed” refers to the comparative inference velocity relative to LLaVA-1.5. “Time” is the average inference time. Inference speed are measured on the A100.

Model	LLM	Connector	V-T Num	Res	TextVQA	GQA	VQAv2	VisWiz	SQA <sub>img</sub>	Speed (†)
7B LLM										
Shikra [5]	Vicuna-7B	MLP	257	224	-	-	77.4	-	-	-
IDEFICS-9B [27]	LLaMA-7B	Cross Attn	257	224	-	38.4	50.9	35.5	-	-
Qwen-VL [1]	Qwen-7B	Resampler	256	448	-	59.3	78.8	35.2	67.1	-
Qwen-VL-Chat [1]	Qwen-7B	Resampler	256	448	-	57.5	78.2	38.9	68.2	-
LLaVA-1.5 [35]	Vicuna-7B	MLP	577	336	58.2	62.0	78.5	50.0	66.8	1.0×
LLaVA-1.5+SEC (ours)	Vicuna-7B	MLP+SEC	257	336	<b>59.6</b>	<b>63.2</b>	<b>78.9</b>	<b>52.8</b>	<b>69.6</b>	<b>1.7×</b>
13B LLM										
InstructBLIP [8]	Vicuna-13B	Q-Former	32	224	-	49.5	-	33.4	63.1	-
BLIP-2 [29]	Vicuna-13B	Q-Former	32	224	-	41.0	41.0	19.5	61.0	-
LLaVA-1.5 [35]	Vicuna-13B	MLP	577	336	61.2	63.3	<b>80.0</b>	53.6	71.6	1.0×
LLaVA1.5+SEC (ours)	Vicuna-13B	MLP+SEC	257	336	<b>62.3</b>	<b>64.3</b>	<b>80.0</b>	<b>54.7</b>	<b>72.0</b>	<b>1.8×</b>

Table 8. Results on General VQA tasks.

effectiveness of SEC in extracting visual information.

### 4.3. Ablation Study

In this section, we present some of the ablation study results for SEC, and more results can be found in the [Appendix](#).

**Number of Vision Tokens in Each Clusters.** The number of vision tokens has a significant impact on the performance and speed of the model. We thoroughly investigate the effect of the number of vision tokens on SECViT. As shown in Tab. 10, the number of vision tokens in each cluster greatly influences the model’s performance. Specifically, in downstream dense prediction tasks, having too few tokens in each cluster leads to substantial performance degradation. When the number of tokens in each cluster is too large, the model’s performance does not see a significant improvement, but its speed decreases.

### 4.4. Visualization of SEC

To further understand the working mechanism of SEC, we visualize some clustering results for SECViT. As shown in Fig. 4, the left side presents the clustering results of vision tokens at different stages of the model. From the clustering results, we analyze that in the shallow layers, the model

distinguishes fine-grained features well, while in the deeper layers, it captures global semantic features effectively. The right side shows the Grad-CAM diagrams at different stages of the model, from which we can draw similar conclusions to the clustering results. More visualization results can be found in [Appendix](#).

**Number of Vision Tokens Outputs by SEC.** MLLM is quite sensitive to the number of vision tokens. We conduct a detailed exploration based on LLaVA-1.5 regarding the number of vision tokens output by SEC, as shown in Tab. 11. The first row represents the speed and performance of the original LLaVA-1.5 without using SEC. Compared to LLaVA-1.5, employing SEC effectively reduces the number of vision tokens and improves training efficiency. As the number of vision tokens decreases, the model’s performance shows a slight decline, but its efficiency is further enhanced.

## 5. Conclusion

We propose a simple and straightforward clustering method for vision tokens—Semantic Equitable Clustering (SEC). This method assigns each token to a cluster by calculating the similarity between each token and a global token, and

Model	LLM	Connector	V-T Num	Res	POPE	MMB	MM-Vet	Speed (↑)
7B LLM								
MiniGPT-4 [60]	Vicuna-7B	Resampler	32	224	72.2	24.3	22.1	-
mPLUG-Owl2 [57]	LLaMA2-7B	Resampler	32	224	-	49.4	-	-
LLaMA-AdapterV2 [14]								
Shikra [5]	LLaMA2-7B	LLaMA-Adapter	257	224	-	41.0	31.4	-
Qwen-VL [1]	Vicuna-7B	MLP	257	224	-	58.8	-	-
Qwen-VL-Chat [1]	Qwen-7B	Resampler	256	448	-	38.2	-	-
	Qwen-7B	Resampler	256	448	-	60.6	-	-
LLaVA-1.5 [35]	Vicuna-7B	MLP	577	336	86.1	64.3	31.1	1.0×
LLaVA1.5+SEC (ours)	Vicuna-7B	MLP+SEC	145	336	<b>86.1</b>	<b>68.4</b>	<b>31.7</b>	<b>2.1×</b>
13B LLM								
MiniGPT-4 [60]	Vicuna-13B	Resampler	32	224	-	-	24.4	-
BLIP-2 [29]	Vicuna-13B	Q-Former	32	224	85.3	-	22.4	-
LLaVA-1.5 [35]	Vicuna-13B	MLP	577	336	86.2	67.7	36.1	1.0×
LLaVA-1.5+SEC (ours)	Vicuna-13B	MLP+SEC	145	336	<b>86.4</b>	<b>69.2</b>	<b>37.3</b>	<b>2.2×</b>

Table 9. Results on benchmark designed for MLLMs.

V-T num	Params (M)	FLOPs (G)	Throughput (imgs/s)	Acc	$AP^b$	$AP^m$	mIoU
98	15	2.5	2004	82.7	47.8	43.0	47.2
196	15	3.1	1722	83.0 (+0.3)	48.2 (+0.4)	43.4 (+0.4)	47.5 (+0.3)
64	15	2.5	1946	82.7 (+0.0)	47.8 (+0.0)	42.8 (-0.2)	46.9 (-0.3)
49	15	2.4	2102	82.6 (-0.1)	47.5 (-0.3)	42.7 (-0.3)	47.7 (-0.5)
24	15	2.3	2186	82.0 (-0.7)	45.9 (-1.9)	40.6 (-2.4)	44.6 (-2.6)

Table 10. Effect of the number of vision tokens in each cluster. “V-T num” means the number of vision tokens in each cluster. The experiments are conducted based on SECViT-T.

V-T num	Time	Speed	TextVQA	GQA	VQAv2	POPE	MM-Vet
576+1	21h	1.0×	58.2	62.0	78.5	86.1	31.1
288+1	14h	1.5×	60.1(+1.9)	63.5(+1.5)	78.9(+0.4)	87.7(+1.6)	33.2(+2.1)
256+1	13h	1.6×	59.6(+1.4)	63.2(+0.3)	78.6(+0.1)	87.1(+1.0)	32.7(+1.6)
192+1	11h	1.9×	57.7(-0.5)	62.7(+0.7)	78.4(-0.1)	86.7(+0.6)	32.1(+1.0)
144+1	10h	2.1×	56.8(-1.4)	62.0(+0.0)	78.0(-0.5)	86.1(+0.0)	31.7(+0.6)

Table 11. Effect of the number of vision tokens outputs by SEC. “V-T num” means the number of vision tokens output by SEC. The experiments are conducted based on LLaVA-1.5 [36].

completes the whole clustering process in only one step. Our clustering method takes into account the semantic information contained in the tokens, and ensures an equal number of tokens in each cluster, facilitating efficient parallel processing on modern GPUs. Based on Semantic Equitable Clustering, we designed SECViT, a versatile vision backbone that achieves impressive results across various vision tasks, including image classification, object detection, instance segmentation, and semantic segmentation. Besides, SEC can also be conveniently applied to multimodal large language models (MLLM) to serve as a vision language connector and benefits the model’s efficiency.

## 6. Acknowledgements

This work is partially funded by Beijing Natural Science Foundation (4252054), Youth Innovation Promotion




Figure 4. Visualization for SEC.

Association CAS(Grant No.2022132), Beijing Nova Program(20230484276), and CCF-Kuaishou Large Model Explorer Fund (NO. CCF-KuaiShou 2024005).

## References

- [1] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023. 2, 6, 7, 8
- [2] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 2018. 5
- [3] Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. In *CVPR*, 2024. 2
- [4] Kai Chen, Jiaqi Wang, Jiangmiao Pang, et al. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5
- [5] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic, 2023. 7, 8
- [6] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, and Chunhua Shen. Conditional positional encodings for vision transformers. In *ICLR*, 2023. 3
- [7] MMSegmentation Contributors. Mmsegmentation, an open source semantic segmentation toolbox, 2020. 5
- [8] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. 7
- [9] Mingyu Ding, Bin Xiao, Noel Codella, et al. Davit: Dual attention vision transformers. In *ECCV*, 2022. 2, 5
- [10] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, et al. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *CVPR*, 2022. 1, 2, 3, 6
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 2
- [12] Qihang Fan, Huaibo Huang, Xiaoqiang Zhou, and Ran He. Lightweight vision transformer with bidirectional interaction. In *NeurIPS*, 2023. 3, 5, 6
- [13] Qihang Fan, Huaibo Huang, Mingrui Chen, Hongmin Liu, and Ran He. Rmt: Retentive networks meet vision transformers. In *CVPR*, 2024. 5, 6
- [14] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, Hongsheng Li, and Yu Qiao. Llama-adapter v2: Parameter-efficient visual instruction model, 2023. 8
- [15] SG-Former: Self guided Transformer with Evolving Token Reallocation. Sucheng ren, xingyi yang, songhua liu, xinzhao wang. In *ICCV*, 2023. 5, 6
- [16] Jianyuan Guo, Kai Han, Han Wu, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Cmt: Convolutional neural networks meet vision transformers. In *CVPR*, 2022. 2
- [17] Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual attention network. *arXiv preprint arXiv:2202.09741*, 2022. 6
- [18] Dongchen Han, Ziyi Wang, Zhuofan Xia, Yizeng Han, Yifan Pu, Chunjiang Ge, Jun Song, Shiji Song, Bo Zheng, and Gao Huang. Demystify mamba in vision: A linear attention perspective. In *NeurIPS*, 2024. 5, 6
- [19] Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer. In *CVPR*, 2023. 5, 6
- [20] Ali Hatamizadeh, Hongxu Yin, Greg Heinrich, Jan Kautz, and Pavlo Molchanov. Global context vision transformers. In *ICML*, 2023. 5, 6
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask r-cnn. In *ICCV*, 2017. 5
- [22] Huaibo Huang, Xiaoqiang Zhou, Jie Cao, Ran He, and Tieniu Tan. Vision transformer with super token sampling. In *CVPR*, 2023. 6
- [23] Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver: General perception with iterative attention, 2021. 2
- [24] Zi-Hang Jiang, Qibin Hou, Li Yuan, Daquan Zhou, Yujun Shi, Xiaojie Jin, Anran Wang, and Jiashi Feng. All tokens matter: Token labeling for training better vision transformers. In *NeurIPS*, 2021. 2
- [25] Manjin Kim, Paul Hongsuck Seo, Cordelia Schmid, and Minsu Cho. Learning correlation structures for vision transformers. In *CVPR*, 2024. 5, 6
- [26] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *CVPR*, 2019. 5
- [27] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander Rush, Douwe Kiela, et al. Obelics: An open web-scale filtered dataset of interleaved image-text documents. In *NeurIPS*, 2024. 7
- [28] Youngwan Lee, Jonghee Kim, Jeffrey Willette, and Sung Ju Hwang. Mpvit: Multi-path vision transformer for dense prediction. In *CVPR*, 2022. 5, 6
- [29] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 7, 8
- [30] Kunchang Li, Yali Wang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unified transformer for efficient spatiotemporal representation learning, 2022. 6
- [31] 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*, 2024. 6
- [32] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *International Conference on Learning Representations*, 2022. 2, 4, 6
- [33] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, and Kaiming He and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 5
- [34] Weifeng Lin, Ziheng Wu, Jiayu Chen, Jun Huang, and Lianwen Jin. Scale-aware modulation meet transformer. In *ICCV*, 2023. 5, 6

- [35] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. [7](#), [8](#)
- [36] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. [2](#), [3](#), [4](#), [6](#), [8](#)
- [37] Kai Liu, Tianyi Wu, Cong Liu, and Guodong Guo. Dynamic group transformer: A general vision transformer backbone with dynamic group attention. In *IJCAI*, 2022. [1](#), [2](#)
- [38] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. [1](#), [2](#), [3](#), [4](#), [5](#)
- [39] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, et al. A convnet for the 2020s. In *CVPR*, 2022. [5](#)
- [40] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, 2021. [2](#)
- [41] Chenyang Si, Weihao Yu, Pan Zhou, Yichen Zhou, Xinchao Wang, and Shuicheng YAN. Inception transformer. In *NeurIPS*, 2022. [2](#), [5](#)
- [42] Shitao Tang, Jiahui Zhang, Siyu Zhu, et al. Quadtree attention for vision transformers. In *ICLR*, 2022. [5](#)
- [43] Hugo Touvron, Matthieu Cord, Matthijs Douze, et al. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. [2](#), [4](#), [5](#)
- [44] Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxvit: Multi-axis vision transformer. In *ECCV*, 2022. [1](#), [2](#), [5](#)
- [45] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021. [2](#)
- [46] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):1–10, 2022. [2](#), [5](#), [6](#)
- [47] Wenxiao Wang, Lu Yao, Long Chen, Binbin Lin, Deng Cai, Xiaofei He, and Wei Liu. Crossformer: A versatile vision transformer hinging on cross-scale attention. In *ICLR*, 2022. [1](#), [5](#)
- [48] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *CVPR*, 2023. [5](#), [6](#)
- [49] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic vision transformers with adaptive sequence length. In *NeurIPS*, 2021. [5](#)
- [50] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. *arXiv preprint arXiv:2301.00808*, 2023. [5](#)
- [51] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *CVPR*, 2022. [2](#)
- [52] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. [5](#)
- [53] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simsim: A simple framework for masked image modeling. In *CVPR*, 2022. [5](#)
- [54] Chenglin Yang, Siyuan Qiao, Qihang Yu, et al. Moat: Alternating mobile convolution and attention brings strong vision models. In *ICLR*, 2023. [5](#)
- [55] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. In *NeurIPS*, 2021. [6](#)
- [56] Rui Yang, Hailong Ma, Jie Wu, Yansong Tang, Xuefeng Xiao, Min Zheng, and Xiu Li. Scalablevit: Rethinking the context-oriented generalization of vision transformer. In *ECCV*, 2022. [6](#)
- [57] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. mplug-owl: Modularization empowers large language models with multimodality, 2024. [8](#)
- [58] Xiaoyu Yue, Shuyang Sun, Zhanghui Kuang, Meng Wei, Philip HS Torr, Wayne Zhang, and Dahua Lin. Vision transformer with progressive sampling. In *ICCV*, 2021. [5](#)
- [59] Wang Zeng, Sheng Jin, Wentao Liu, Chen Qian, Ping Luo, Wanli Ouyang, and Xiaogang Wang. Not all tokens are equal: Human-centric visual analysis via token clustering transformer. In *CVPR*, 2022. [5](#)
- [60] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models, 2023. [8](#)
- [61] Lei Zhu, Xinjiang Wang, Zhanghan Ke, Wayne Zhang, and Rynson Lau. Biformer: Vision transformer with bi-level routing attention. In *CVPR*, 2023. [2](#), [5](#)