

BOAT: Bilateral Local Attention Vision Transformer

Tan Yu¹

tan.yu1503@gmail.com

Gangming Zhao²

gangmingzhao@gmail.com

Ping Li¹

pingli98@gmail.com

Yizhou Yu²

yizhouy@acm.org

¹ Cognitive Computing Lab

Baidu Research

10900 NE 8th St. Bellevue, WA USA

² Department of Computer Science

The University of Hong Kong

Abstract

Vision Transformers achieved outstanding performance in many computer vision tasks. Early Vision Transformers such as ViT and DeiT adopt global self-attention, which is computationally expensive when the number of patches is large. To improve the efficiency, recent Vision Transformers adopt local self-attention mechanisms, where self-attention is computed within local windows. Despite the fact that window-based local self-attention significantly boosts efficiency, it fails to capture the relationships between distant but similar patches in the image plane. To overcome this limitation of image-space local attention, in this paper, we further exploit the locality of patches in the feature space. We group the patches into multiple clusters using their features, and self-attention is computed within every cluster. Such feature-space local attention effectively captures the connections between patches across different local windows but still relevant. We propose a Bilateral Local Attention vision Transformer (BOAT), which integrates feature-space local attention with image-space local attention. We further integrate BOAT with both Swin and CSWin models, and extensive experiments on several benchmark datasets demonstrate that our BOAT-CSWin model clearly and consistently outperforms existing state-of-the-art CNN models and vision Transformers.

1 Introduction

Following the great success of Transformers [26] in natural language processing tasks, researchers have recently proposed vision Transformers [4, 6, 9, 18, 25, 32], which have achieved outstanding performance in many computer vision tasks, including image recognition, detection, and segmentation. As early versions of vision transformers, ViT [9] and DeiT [25] uniformly divide an image into 16×16 patches (tokens) and apply a stack of standard Transformer layers to a sequence of tokens formed using these patches. The original self-attention mechanism is global, i.e., the receptive field of a patch in ViT and DeiT covers all patches of the image, which is vital for modeling long-range interactions among patches. On the other hand, the global nature of self-attention imposes a great challenge in

efficiency. Specifically, the computational complexity of self-attention is quadratic in terms of the number of patches. As the number of patches is inversely proportional to the patch size when the size of the input image is fixed, the computational cost forces ViT and DeiT to adopt medium-size patches, which might not be as effective as smaller patches generating higher-resolution feature maps, especially for dense prediction tasks such as segmentation.

To maintain higher resolution feature maps while achieving high efficiency, some methods [6, 18] exploit image-space local attention. They divide an image into multiple local windows, each of which includes a number of patches. Self-attention operations are only performed on patches within the same local window. This is a reasonable design since a patch is likely to be affiliated with other patches in the same local window but not highly relevant to patches in other windows. Thus, pruning attention between patches from different windows might not significantly deteriorate the performance. Meanwhile, the computational cost of window-based self-attention is much lower than that of the original self-attention over the entire image. Swin Transformer [18] and Twins [6] are such examples. Swin Transformer [18] performs self-attention within local windows. To facilitate communication between patches from different windows, Swin Transformer has two complementary window partitioning schemes, and a window in one scheme overlaps with multiple windows in the second scheme. Twins [6] performs self-attention within local windows and builds connections among different windows by performing (global) self-attention over feature vectors sparsely sampled from the entire image using a regular subsampling pattern.

In this work, we rethink local attention and explore locality from a broader perspective. Specifically, we investigate feature-space local attention apart from its image-space counterpart. Instead of computing local self-attention in the image space, feature-space local attention exploits locality in the feature space. It is based on the fact that patch feature vectors close to each other in the feature space tend to have more influence on each other in the computed self-attention results. This is because the actual contribution of a feature vector to the self-attention result at another feature vector is controlled by the similarity between these two feature vectors. Feature-space local attention computes the self-attention result at a feature vector using its feature-space nearest neighbors only while setting the contribution from feature vectors farther away to zero. This essentially defines a piecewise similarity function, which clamps the similarity between feature vectors far apart to zero. In comparison to the aforementioned image-space local attention, feature-space local attention has been rarely exploited in vision transformers. As shown in Figure 1, feature-space local attention computes attention among relevant patches which might not be close to each other in the image plane. Thus, it is a natural compensation to image-space local attention, which might miss meaningful connections between patches residing in different local windows.

In this paper, we propose a novel vision Transformer architecture, Bilateral lOcal Attention vision Transformer (BOAT), to exploit the complementarity between feature-space and image-space local attention. The essential component in our network architecture is the bilateral local attention block, consisting of a feature-space local attention module and an image-space local attention module. The image-space local attention module divides an im-




Figure 1: The image-space local attention versus the feature-space local attention.

age into multiple local windows as Swin [18] and CSWin [8], and self-attention is computed within each local window. In contrast, feature-space local attention groups all the patches into multiple clusters and self-attention is computed within each cluster. Feature-space local attention could be implemented in a straightforward way using K-means clustering. Nevertheless, K-means clustering cannot ensure the generated clusters are evenly sized, thus impedes efficient parallel implementation. In addition, sharing self-attention parameters among unevenly sized clusters may also negatively impact the effectiveness of self-attention. To overcome this obstacle, we propose hierarchical balanced clustering, which groups patches into clusters of equal size.

We conduct experiments on multiple computer vision tasks, including image classification, semantic segmentation, and object detection. Experiments on several public benchmarks demonstrate that our BOAT clearly and consistently improves existing image-space local attention vision Transformers, including Swin [18] and CSWin [8], on these tasks.

2 Related Work

2.1 Vision Transformers

In the past decade, CNN has achieved tremendous successes in numerous computer vision tasks [11, 16]. The natural language processing (NLP) backbone, Transformer, has recently attracted the attention of researchers in the computer vision community. After dividing an image into non-overlapping patches (tokens), Vision Transformer (ViT) [9] applies Transformer for communications among the tokens. Without delicately devised convolution kernels, ViT achieved excellent performance in image recognition in comparison to CNNs using a huge training corpus. DeiT [25] improves data efficiency by exploring advanced training and data augmentation strategies. Recently, many efforts have been devoted to improving the recognition accuracy and efficiency of Vision Transformers.

To boost the recognition accuracy, T2T-ViT [36] proposes a Tokens-to-Token transformation, recursively aggregating neighboring tokens into one token for modeling local structures. TNT [10] also investigates local structure modeling. It additionally builds an inner-level Transformer to model the visual content within each local patch. PVT [28] uses small-scale patches, yielding higher resolution feature maps for dense prediction. Meanwhile, PVT progressively shrinks the feature map size for computation reduction. PiT [13] also decreases spatial dimensions through pooling and increases channel dimensions in deeper layers.

More recently, computing self-attention within local windows [6, 8, 14, 18], has achieved a good trade-off between effectiveness and efficiency. For example, Swin [18] divides an image into multiple local windows and computes self-attention among patches from the same window. To achieve communication across local windows, Swin shifts window configurations in different layers. Twins [8] also exploits local windows for enhancing efficiency. To achieve cross-window communication, it computes additional self-attention over features sampled from the entire image. Similarly, Shuffle Transformer [14] exploits local windows and performs cross-window communication by shuffling patches. CSWin [8] adopts cross-shaped windows, computing self-attention in horizontal and vertical stripes in parallel. The aforementioned local attention models [6, 8, 14, 18] only exploit image-space locality. In contrast, our BOAT exploits not only image-space locality but also feature-space locality.

2.2 Efficient Transformers

High computational costs limit Transformer’s usefulness in practice. Thus, much research [24] has recently been dedicated to improving efficiency. One popularly used strategy for speeding up Transformers enforces sparse attention matrices by limiting the receptive field of each token. Image Transformer [19] and Block-wise Transformer [20] divide a long sequence into local buckets. In this case, the attention matrix has a block-diagonal structure. Only self-attention within each bucket is retained, and cross-bucket attention is pruned. Transformers based on image-space local attention, such as Swin Transformer [18], Twins [6], Shuffle Transformer [24], and CSWin Transformer [8] also adopt buckets (windows) for boosting efficiency. In parallel to bucket-based local attention, strided attention is another approach for achieving sparse attention matrices. Sparse Transformer [5] and LongFormer [9] utilize strided attention, which computes self-attention over features sampled with a sparse grid with a stride larger than one, leading to a sparse attention matrix facilitating faster computation. The global sub-sampling layer in Twins [6] and the shuffle module in Shuffle Transformer [24] can be regarded as strided attention modules. Some recent works exploit pure MLP-based architectures [1, 33, 34, 35] to boost efficiency.

Unlike the above mentioned image-space local attention, several methods determine the scope of local attention in the feature space. Reformer [15] distributes tokens to buckets by feature-space hashing functions. Routing Transformer [22] applies online K-means to cluster tokens. Sinkhorn Sorting Network [23] learns to sort and divide an input sequence into chunks. Our feature-space local attention module also falls into this category. As far as we know, this paper is the first attempt to apply feature-space grouping to vision Transformers.

3 Method




Figure 2: Architecture of Bilateral IOcal Attention Vision Transformer (BOAT).

As visualized in Figure 2, the proposed BOAT architecture consists of a patch embedding module, and a stack of L Bilateral Local Attention blocks. Meanwhile, we exploit a hierarchical pyramid structure. Below we only briefly introduce the patch embedding module and the hierarchical pyramid structure and leave the details of the proposed Bilateral Local Attention block in Section 3.1 and 3.2.

Patch embedding. For an input image with size $H \times W$, we follow Swin [18] and CSWin Transformer [8], and leverage convolutional token embedding (7×7 convolution layer with stride 4) to obtain $\frac{H}{4} \times \frac{W}{4}$ patch tokens, and the dimension of each token is C .

Hierarchical pyramid structure. Similar to Swin [18] and CSWin Transformer [8], we also build a hierarchical pyramid structure. The whole architecture consists of four stages. A convolution layer (3×3 , stride 2) is used between two adjacent stages to merge patches. It reduces the number of tokens and doubles the number of channels. Therefore, in the i -th stage, the feature map contains $\frac{H}{2^{(i+1)}} \times \frac{W}{2^{(i+1)}}$ tokens and $2^{i-1}C$ channels.

3.1 Bilateral Local Attention Block




Figure 3: Architecture of Bilateral Local Attention (BLA) Block.

As shown in Figure 3, a Bilateral Local Attention (BLA) Block consists of an image-space local attention (ISLA) module, a feature-space (content-based) local attention (FSLA) module, an MLP module, and several layer normalization (LN) modules. Let us denote the set of input tokens by $\mathcal{T}_{\text{in}} = \{\mathbf{t}_i\}_{i=1}^N$ where $\mathbf{t}_i \in \mathbb{R}^C$, C is the number of channels and N is the number of tokens. The input tokens go through a normalization layer followed by an image-space local attention (ISLA) module, which has a shortcut connection:

$$\mathcal{T}_{\text{ISLA}} = \mathcal{T}_{\text{in}} + \text{ISLA}(\text{LN}(\mathcal{T}_{\text{in}})). \quad (1)$$

Image-space local attention only computes self-attention among tokens within the same local window. We adopt existing window-based local attention modules, such as those in Swin Transformer [18] and CSWin Transformer [8] as our ISLA module due to their excellent performance. Intuitively, patches within the same local window are likely to be closely related to each other. However, some distant patches in the image space might also reveal important connections, such as similar contents, that could be helpful for visual understanding. Simply throwing away such connections between distant patches in the image space might deteriorate image recognition performance.

To bring back the useful information dropped out by image-space local attention, we develop a feature-space local attention (FSLA) module. The output of the ISLA module, $\mathcal{T}_{\text{ISLA}}$, is fed into another normalization layer followed by a feature-space (content-based) local attention (FSLA) module, which also has a shortcut connection:

$$\mathcal{T}_{\text{FSLA}} = \mathcal{T}_{\text{ISLA}} + \text{FSLA}(\text{LN}(\mathcal{T}_{\text{ISLA}})). \quad (2)$$

The FSLA module computes self-attention among tokens that are close in the feature space, which is complementary to the ISLA module. Meanwhile, by only considering local attention in the feature space, FSLA is more efficient than the original (global) self-attention. We will present the details of FSLA in Section 3.2. Following CSWin [8], we also add locally-enhanced positional encoding to each feature-space local attention layer to model position.

At last, the output of the FSLA module, $\mathcal{T}_{\text{FSLA}}$, is processed by another normalization layer and an MLP module to generate the output of a Bilateral Local Attention Block:

$$\mathcal{T}_{\text{out}} = \mathcal{T}_{\text{FSLA}} + \text{MLP}(\text{LN}(\mathcal{T}_{\text{FSLA}})). \quad (3)$$

Following existing vision Transformers [9, 18], the MLP module consists of two fully-connected layers. The first one increases the feature dimension from C to rC and the second one decreases the feature dimension from rC back to C . By default, we set $r = 4$.

3.2 Feature-Space Local Attention

Different from image-space local attention which groups tokens according to their spatial locations in the image plane, feature-space local attention seeks to group tokens according to

their content, *i.e.*, features. We could simply perform K-means clustering on token features to achieve this goal. Nevertheless, K-means clustering cannot ensure that the generated clusters are equally sized, which makes it difficult to have efficient parallel implementation on GPU platforms, and may also negatively impact the overall effectiveness of self-attention.

Balanced hierarchical clustering. To overcome the imbalance problem of K-means clustering, we propose a balanced hierarchical clustering, which performs K levels of clustering. At each level, it conducts balanced binary clustering, which equally splits a set of tokens into two clusters. Let us denote the set of input tokens by $\mathcal{T} = \{\mathbf{t}_i\}_{i=1}^N$. In the first level, it splits N tokens in \mathcal{T} into two subsets with $N/2$ tokens each. At the k -th level, it splits $N/2^{k-1}$ tokens assigned to the same subset in the upper level into two smaller subsets of $N/2^k$ size. At the end, we obtain 2^K evenly sized subsets in the final level, $\{\mathcal{T}_i\}_{i=1}^{2^K}$, and the size of each subset $|\mathcal{T}_i|$ is equal to $N/2^K$.

Here, we require the condition that N is divisible by 2^K , which can be easily satisfied in existing vision Transformers. We visualize the process of balanced hierarchical clustering in Figure 4. The core operation in balanced hierarchical clustering is our devised balanced binary clustering, which we elaborate below.

Balanced binary clustering. Given a set of $2m$ tokens $\{\mathbf{t}_i\}_{i=1}^{2m}$, balanced binary clustering divides them into two groups and the size of each group is m . Similar to K-means clustering, our balanced binary clustering relies on cluster centroids. To determine the cluster membership of each sample, K-means clustering only considers the distance between the sample and all centroids. In contrast, our balanced binary clustering further requires that the two resulting clusters have equal size. Let us denote the two cluster centroids as \mathbf{c}_1 and \mathbf{c}_2 . For each token \mathbf{t}_i , we compute distance ratio, r_i , as a metric to determine its cluster membership:

$$r_i = \frac{s(\mathbf{t}_i, \mathbf{c}_1)}{s(\mathbf{t}_i, \mathbf{c}_2)}, \forall i \in [1, 2m], \quad (4)$$

where $s(\mathbf{x}, \mathbf{y})$ denotes the cosine similarity between \mathbf{x} and \mathbf{y} . The $2m$ tokens $\{\mathbf{t}_i\}_{i=1}^{2m}$ are sorted in the decreasing order of their distance ratios $\{r_i\}_{i=1}^{2m}$. We assign the tokens in the first half of the sorted list to the first cluster \mathcal{C}_1 and those in the second half of the sorted list to the second cluster \mathcal{C}_2 , where the size of both \mathcal{C}_1 and \mathcal{C}_2 is m . The mean of the tokens from each cluster is used to update the cluster centroid. Similar to K-means, our balanced binary clustering updates cluster centroids and the cluster membership of every sample in an iterative manner. Note that cluster centroids are always computed on the fly, and are not learnable parameters. The detailed steps of the proposed balanced binary clustering are given in Algorithm 1.

In the aforementioned balanced binary clustering, two resulting clusters have no shared tokens, *i.e.*, $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$. One main drawback of the non-overlapping setting is that, a token in the middle portion of the sorted list has some of its feature-space neighbors in one cluster while the other neighbors in the other cluster. No matter which cluster this token is finally assigned to, the connection between the token and part of its feature-space neighbors will be cut off. For example, the token at the m -th location of the sorted list cannot communicate




Figure 4: Example of balanced hierarchical clustering. In this example, the number of hierarchical levels is 3. There are $2^3 = 8$ clusters in the bottom level.

Algorithm 1: Balanced Binary Clustering.**Input:** Tokens $\{\mathbf{t}_i\}_{i=1}^{2m}$ and the iteration number, T .**Output:** Two clusters, \mathcal{C}_1 and \mathcal{C}_2 .

```

1 Initialize centroids  $\mathbf{c}_1 = \frac{\sum_{i=1}^m \mathbf{t}_i}{m}$ ,  $\mathbf{c}_2 = \frac{\sum_{i=m+1}^{2m} \mathbf{t}_i}{m}$ 
2 while  $n\_iter \in [1, T]$  do
3   for  $i \in [1, 2m]$  do
4      $r_i = \frac{s(\mathbf{t}_i, \mathbf{c}_1)}{s(\mathbf{t}_i, \mathbf{c}_2)}$ 
5    $[i_1, \dots, i_{2m}] = \text{argsort}([r_1, \dots, r_{2m}])$ 
6    $\mathcal{C}_1 = \{\mathbf{t}_{i_j}\}_{j=1}^m$ ,  $\mathcal{C}_2 = \{\mathbf{t}_{i_j}\}_{j=m+1}^{2m}$ 
7    $\mathbf{c}_1 = \frac{\sum_{\mathbf{c} \in \mathcal{C}_1} \mathbf{c}}{m}$ ,  $\mathbf{c}_2 = \frac{\sum_{\mathbf{c} \in \mathcal{C}_2} \mathbf{c}}{m}$ 

```

with the token at the $m + 1$ -st location during attention calculation because they are assigned to different clusters. Overlapping balanced binary clustering overcomes this drawback by assigning the first $m + n$ tokens in the sorted list to the first cluster, *i.e.*, $\hat{\mathcal{C}}_1 = \{\mathbf{t}_{j_i}\}_{i=1}^{m+n}$, and the last $m + n$ tokens in the sorted list to the second cluster, *i.e.*, $\hat{\mathcal{C}}_2 = \{\mathbf{t}_{j_i}\}_{i=m-n+1}^{2m}$. Thus, the two resulting clusters have $2n$ tokens in common, *i.e.*, $\hat{\mathcal{C}}_1 \cap \hat{\mathcal{C}}_2 = \{\mathbf{t}_{j_i}\}_{i=m-n+1}^{m+n}$. By default, we only adopt overlapping binary clustering at the last level of the proposed balanced hierarchical clustering and use the non-overlapping version at the other levels. We set $n = 20$ in all experiments for overlapping binary clustering.

Local attention within cluster. Through the above introduced balanced hierarchical clustering, the set of tokens, \mathcal{T} , are grouped into 2^K subsets $\{\mathcal{T}_{ij}\}_{i=1}^{2^K}$, where $|\mathcal{T}_{ij}| = \frac{N}{2^K}$. The standard self-attention (SA) is performed within each subset:

$$\hat{\mathcal{T}}_k = \text{SA}(\mathcal{T}_k), \forall k \in [1, 2^K]. \quad (5)$$

The output, $\hat{\mathcal{T}}$, is the union of all attended subsets:

$$\hat{\mathcal{T}} = \bigcup_{k \in [1, K]} \hat{\mathcal{T}}_k. \quad (6)$$

Following the multi-head configuration in Transformer, we also devise multi-head feature-space local attention. Note that, in our multi-head feature-space local attention, we implement multiple heads not only for computing self-attention in Eq. (5) as a standard Transformer, but also for performing balanced hierarchical clustering. That is, balanced hierarchical clustering is performed independently in each head. Thus, for a specific token, in different heads, it might pay feature-based local attention to different tokens. This configuration is more flexible than Swin [18], where multiple heads share the same local window.

4 Experiments

To demonstrate the effectiveness of our BOAT as a general vision backbone, we conduct experiments on image classification, semantic segmentation and object detection. We build BOAT on top of two recent local attention vision Transformers, Swin [18] and CSWin [8]. We term the BOAT built upon Swin as BOAT-Swin. In BOAT-Swin, the image-space local

attention (ISLA) module adopts shifted window attention in Swin. In contrast, the ISLA module in BOAT-CSWin uses cross-shape window attention in CSWin. We provide the detailed specifications of BOAT-Swin and BOAT-CSWin in Section 1 of the supplementary materials. Meanwhile, we present main experimental results in the following sections. More ablation studies are presented in Section 2 of the supplementary materials.

4.1 Image Classification

Method	size	#para.	FLOPs	Top-1	Method	size	#para.	FLOPs	Top-1
ReGNetY-4G [10]	224	21M	4.0G	80.0					
PVTv2-B2 [10]	224	25M	4.0G	82.0	Focal-T [10]	224	29M	4.9G	82.2
Swin-T [10]	224	29M	4.5G	81.3	BOAT-Swin-T (ours)	224	31M	5.2G	82.3
CSWin-T [8]	224	23M	4.3G	82.7	BOAT-CSWin-T (ours)	224	27M	5.1G	83.7
ReGNetY-8G [10]	224	39M	8.0	81.7	PVTv2-B4 [10]	224	62M	10.1G	83.6
Twins-B	224	56M	8.3G	83.2	Shuffle-S [10]	224	50M	8.9G	83.5
NesT-S [8],	224	38M	10.4G	83.3	Focal-S [8]	224	51M	9.1G	83.5
Swin-S [10]	224	50M	8.7G	83.0	BOAT-Swin-S (ours)	224	56M	10.1G	83.6
CSWin-S [8]	224	35M	6.9G	83.6	BOAT-CSWin-S (ours)	224	41M	8.0G	84.1
ReGNetY-16G [10]	224	84M	16.0G	82.9	ViT-B/16T [8]	384	86M	55.4G	77.9
DeiT-B [10]	224	86M	17.5G	81.8	T2T-24 [10]	224	64M	14.1G	82.3
TNT-B [10]	224	66M	14.1G	82.8	PiT-B [10]	224	74M	12.5G	82.0
PVTv2-B5 [10]	224	82M	11.8G	83.8	Twins-L	224	99M	14.8G	83.7
Shuffle-B [10]	224	88M	15.4G	84.0	NesT-B [8],	224	68M	17.9G	83.8
Focal-B [10]	224	90M	16.0G	83.8	CrossFormer-L [10]	224	92M	16.1G	84.0
Swin-B [10]	224	88M	15.4G	83.5	BOAT-Swin-B (ours)	224	98M	17.8G	83.8
CSWin-B [8]	224	78M	15.0G	84.2	BOAT-CSWin-B (ours)	224	90M	17.5G	84.7

Table 1: Comparison of image classification performance on the ImageNet-1K dataset.

We follow the same training strategies as other vision Transformers. We train our models using the training split of ImageNet-1K [10] with 224×224 input resolution and without external data. Specifically, both Swin and BOAT-Swin are trained for 300 epochs, and both CSWin and BOAT-CSWin are trained for 310 epochs. Table 1 compares the performance of the proposed BOAT models with the state-of-the-art vision backbones. As shown in the table, with a slight increase in the number of parameters and FLOPs, our BOAT-Swin model consistently improves the vanilla Swin model under the tiny, small and base settings. Meanwhile, our BOAT-CSWin model also improves the vanilla CSWin model by a similar degree under the tiny, small and base settings. Such improvements over Swin and CSWin models demonstrate the effectiveness of feature-space local attention.

Comparisons with Reformer and K-means. Reformer [10] also exploits feature-space local attention. It divides tokens into multiple groups using Locality Sensitivity Hashing (LSH), based on sign random projections [2, 10], which is independent of specific input data, and might be sub-optimal for different input data. K-means clustering is another choice for dividing tokens into multiple groups for exploiting feature-space local attention. Nevertheless, K-means clustering cannot ensure that the generated clusters are equally sized, which makes it difficult to have efficient parallel implementation on GPU platforms. To enforce the clusters from K-means clustering to be equally sized, we can sort the tokens according to their cluster index and then equally divide the sorted tokens into multiple groups, as visualized in Figure 5. However, this sort-and-divide process might divide tokens from a large cluster into multiple groups and also merge tokens from small clusters into the same group. This would negatively impact the overall effectiveness of feature-space local attention.




Figure 5: The process of enforcing the clusters from K-means to be equally sized.

Method	Reformer	K-means	Ours
Top-1 Accuracy	81.7	81.8	82.3

Table 2: Comparison of image classification accuracy with Reformer and K-means.

We compare the performance of BOAT-Swin-Tiny against the performance of a model where our balanced hierarchical clustering is replaced with LSH in Reformer or K-means clustering. We keep the other layers unchanged. As shown in Table 2, our BOAT-Swin-Tiny clearly outperforms Reformer and K-means clustering.

The effectiveness of FSLA. To directly demonstrate the effectiveness of feature-space local attention (FSLA), we replace all FSLA blocks in BOAT-Swin-T with image-space local attention (ISLA) blocks. As shown in Table 3, the accuracy drops from 82.3% to 81.5%.

Model	BOAT-Swin-T (with FSLA)	Baseline (with ISLA)
Accuracy	82.3	81.5

Table 3: Ablation study on FSLA by replacing FSLA with ISLA.

The effectiveness of overlapping balanced hierarchical clustering. We compare the performance of overlapping balanced hierarchical clustering with its non-overlapping counterpart on the ImageNet-1K dataset. As shown in Table 4, the overlapping setting achieves consistently higher classification accuracy in BOAT-CSwin-Tiny, Small and Base models. Higher accuracy is expected since the overlapping setting gives rise to larger receptive fields.

Overlap	BOAT-CSWin-T	BOAT-CSWin-S	BOAT-CSWin-B
No	83.3%	84.0%	84.5%
Yes	83.7%	84.1%	84.7%

Table 4: Comparison of image classification accuracy between overlapping balanced hierarchical clustering and the non-overlapping version.

4.2 Semantic Segmentation

We further investigate the effectiveness of our BOAT for semantic segmentation on the ADE20K dataset [83]. Here, we employ UperNet [80] as the basic framework. For a fair comparison, we follow previous work and train UperNet 160K iterations with batch size 16 using 8 GPUs. In Table 5, we compare the semantic segmentation performance of our BOAT with other vision Transformer models including Swin [18], Twins [6], Shuffle Transformer [14], Focal Transformer [51], and CSWin [8]. As shown in the table, with a slight increase in the number of parameters and FLOPs, our BOAT-Swin model consistently improves the semantic segmentation performance of the Swin model under the tiny, small and base settings. Meantime, our BOAT-CSWin also constantly obtains higher segmentation mIoUs than the CSWin model under the tiny, small and base settings.

Method	#para.(M)	FLOPs(G)	mIoU(%)	Method	#para.(M)	FLOPs(G)	mIoU(%)
TwinsP-S [8]	55	919	46.2	Twins-S [8]	54	901	46.2
Shuffle-T [10]	60	949	46.6	Focal-T [10]	62	998	45.8
Swin-T [13]	60	945	44.5	BOAT-Swin-T (ours)	62	986	46.0
CSWin-T [8]	60	959	49.3	BOAT-CSWin-T (ours)	64	1012	50.5
TwinsP-B [8]	74	977	47.1	Twins-B [8]	89	1020	47.7
Shuffle-S [10]	81	1044	48.4	Focal-S [10]	85	1130	48.0
Swin-S [13]	81	1038	47.6	BOAT-Swin-S (ours)	87	1113	48.4
CSWin-S [8]	65	1027	50.0	BOAT-CSWin-S (ours)	70	1101	50.6
TwinsP-L [8]	92	1041	48.6	Twins-L [8]	133	1164	48.8
Shuffle-B [10]	121	1196	49.0	Focal-B [10]	126	1354	49.0
Swin-B [13]	121	1188	48.1	BOAT-Swin-B (ours)	131	1299	48.7
CSWin-B [8]	109	1222	50.8	BOAT-CSWin-B (ours)	121	1349	50.9

Table 5: Performance of semantic segmentation on ADE20K. FLOPs are obtained at 512×2048 resolution. mIoU is for the single-scale setting. Testing image size is 512×512 .

4.3 Object Detection

We also evaluate the proposed BOAT on object detection. Experiments are conducted on the MS-COCO dataset using the Mask R-CNN [12] framework. Since CSWin has not released codes for object detection, we only implement BOAT-Swin for this task. We adopt the $3 \times$ learning rate schedule, which is the same as Swin. We compare the performance of our BOAT-Swin and the original Swin in Table 6. The evaluation is on the MSCOCO val2017 split. Since Swin only reports the performance of Swin-Tiny and Swin-Small models when using the Mask R-CNN framework, we also report the performance of our BOAT-Swin-Tiny and BOAT-Swin-Small only. As shown in Table 6, with a slight increase in the number of parameters and FLOPs, our BOAT-Swin consistently outperforms the original Swin.

Method	#para.(M)	FLOPs(G)	mAP ^{Box}	mAP ^{Mask}
Swin-T	48	267	46.0	41.6
BOAT-Swin-T (ours)	50	306	47.5	42.8
Swin-S	69	359	48.5	43.3
BOAT-Swin-S (ours)	75	431	49.0	43.8

Table 6: Performance of object detection on the MS-COCO dataset. FLOPs are obtained at 800×1280 resolution.

5 Conclusion

In this paper, we have presented a new Vision Transformer architecture named Bilateral Local Attention Transformer (BOAT), which performs multi-head local self-attention in both feature and image spaces. To compute feature-space local attention, we propose a hierarchical balanced clustering approach to group patches into multiple evenly sized clusters, and self-attention is computed within each cluster. We have applied BOAT to multiple computer vision tasks including image classification, semantic segmentation and object detection. Our systematic experiments on several benchmark datasets have demonstrated that BOAT can clearly and consistently improve the performance of existing image-space local attention vision Transformers, including Swin [13] and CSWin [8], on these tasks.

References

- [1] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [2] Moses S Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing (STOC)*, pages 380–388, Montreal, Canada, 2002.
- [3] Shuo Chen, Tan Yu, and Ping Li. MVT: multi-view vision transformer for 3D object recognition. In *Proceedings of the 32nd British Machine Vision Conference 2021 (BMVC)*, page 349, Online, 2021.
- [4] Shuo Chen, Tan Yu, and Ping Li. R²-MLP: Round-roll MLP architecture for multi-view 3D object recognition. *Technical Report*, 2022.
- [5] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [6] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9355–9366, virtual, 2021.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, Miami, FL, 2009.
- [8] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. CSWin transformer: A general vision transformer backbone with cross-shaped windows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12114–12124, New Orleans, LA, 2022.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, Virtual Event, Austria, 2021.
- [10] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 15908–15919, virtual, 2021.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, Las Vegas, NV, 2016.
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, Venice, Italy, 2017.

- [13] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11916–11925, Montreal, Canada, 2021.
- [14] Zilong Huang, Youcheng Ben, Guozhong Luo, Pei Cheng, Gang Yu, and Bin Fu. Shuffle transformer: Rethinking spatial shuffle for vision transformer. *arXiv preprint arXiv:2106.03650*, 2021.
- [15] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, 2020.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, Lake Tahoe, CA, 2012.
- [17] Ping Li. Sign-full random projections. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*, pages 4205–4212, Honolulu, HI, 2019.
- [18] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, Montreal, Canada, 2021.
- [19] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 4052–4061, Stockholmsmäs-san, Stockholm, Sweden, 2018.
- [20] Jiezhong Qiu, Hao Ma, Omer Levy, Wen-tau Yih, Sinong Wang, and Jie Tang. Block-wise self-attention for long document understanding. In *Findings of the Association for Computational Linguistics (EMNLP Findings)*, pages 2555–2565, Online Event, 2020.
- [21] Ilija Radosavovic, Raj Prateek Kosaraju, Ross B. Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10425–10433, Seattle, WA, 2020.
- [22] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Trans. Assoc. Comput. Linguistics*, 9: 53–68, 2021.
- [23] Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. Sparse sinkhorn attention. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 9438–9447, Virtual Event, 2020.
- [24] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*, 2020.

- [25] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 10347–10357, Virtual Event, 2021.
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*, pages 5998–6008, Long Beach, CA, 2017.
- [27] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. PVTv2: Improved baselines with pyramid vision transformer. *arXiv preprint arXiv:2106.13797*, 2021.
- [28] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 548–558, Montreal, Canada, 2021.
- [29] 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. *arXiv preprint arXiv:2108.00154*, 2021.
- [30] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the 15th European Conference on Computer Vision (ECCV), Part V*, pages 432–448, Munich, Germany, 2018.
- [31] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal attention for long-range interactions in vision transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 30008–30022, virtual, 2021.
- [32] Tan Yu and Ping Li. Degenerate Swin to win: Plain window-based transformer without sophisticated operations. *Technical Report*, 2022.
- [33] Tan Yu, Xu Li, Yunfeng Cai, Mingming Sun, and Ping Li. Rethinking token-mixing MLP for mlp-based vision backbone. In *Proceedings of the 32nd British Machine Vision Conference (BMVC)*, page 139, Online, 2021.
- [34] Tan Yu, Xu Li, Yunfeng Cai, Mingming Sun, and Ping Li. S²-MLPv2: Improved spatial-shift mlp architecture for vision. *arXiv preprint arXiv:2108.01072*, 2021.
- [35] Tan Yu, Xu Li, Yunfeng Cai, Mingming Sun, and Ping Li. S²-MLP: Spatial-shift MLP architecture for vision. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3615–3624, Waikoloa, HI, 2022.
- [36] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zihang Jiang, Francis E. H. Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 538–547, Montreal, Canada, 2021.

- [37] Zizhao Zhang, Han Zhang, Long Zhao, Ting Chen, Sercan Ö. Arik, and Tomas Pfister. Nested hierarchical transformer: Towards accurate, data-efficient and interpretable visual understanding. In *Proceedings of the thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*, pages 3417–3425, Virtual Event, 2022.
- [38] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20K dataset. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5122–5130, Honolulu, HI, 2017.