MAIT: INTEGRATING SPATIAL LOCALITY INTO IMAGE TRANSFORMERS WITH ATTENTION MASKS

Anonymous authors

Paper under double-blind review

Abstract

Though image transformers have shown competitive results with convolutional neural networks in computer vision tasks, lacking inductive biases such as locality still poses problems in terms of model efficiency especially for embedded applications. In this work, we address this issue by introducing attention masks to incorporate spatial locality into self-attention heads. Local dependencies are captured with masked attention heads along with global dependencies captured by original unmasked attention heads. With Masked attention image Transformer - MaiT, top-1 accuracy increases by up to 1.0% compared to DeiT, without extra parameters, computation, or external training data. Moreover, attention masks regulate the training of attention maps, which facilitates the convergence and improves the accuracy of deeper transformers. Masked attention heads guide the model to focus on local information in early layers and promote diverse attention maps in latter layers. Deep MaiT improves the top-1 accuracy by up to 1.5%compared to CaiT with fewer parameters and less FLOPs. Encoding locality with attention masks requires no extra parameter or structural change, and thus it can be combined with other techniques for further improvement in vision transformers.

1 INTRODUCTION

Convolutional neural networks (CNNs) (Krizhevsky et al., 2012; He et al., 2016; Tan & Le, 2019) have been the de facto model for computer vision (CV) tasks, which are inherently equipped with inductive biases such as translation equivariance and locality. Recently, vision transformers (Dosovitskiy et al., 2021; Touvron et al., 2021a) are gaining momentum translating the success from transformer-based models in natural language processing (NLP) tasks (Vaswani et al., 2017; Devlin et al., 2019). Self-attention heads excel at capturing the long-range dependencies in sequences but struggle at focusing on local information. Unlike CNNs which have gone through many iterations of optimization, vision transformers are not very efficient and still require huge computing power and a large number of Flops.

Naturally, combining the benefits of both CNNs and vision transformers is promising to further boost performances of CV models. The question remains how to effectively integrate inductive bias such as spatial locality into transformers. One direction is to utilize convolutional blocks to extract spatial information by adapting either the patch-token embedding layer, self-attention module or feed-forward layers, to form CNN-transformer hybrid structures as in Li et al. (2021b); Srinivas et al. (2021); Graham et al. (2021); Wu et al. (2021a); Yuan et al. (2021a); Guo et al. (2021). However, forcefully inserting convolutional operations into transformers may potentially constrain the learning capacity of transformers.

To capture the spatial information without significantly changing the transformer model architecture, Chu et al. (2021b) introduce extra positional encoding. Han et al. (2021); Chen et al. (2021) fuse local and global representations using multiple transformer blocks or branches to simultaneously process images at different scales such as pixel-level, small-patch or large patch. Yuan et al. (2021b) apply a layerwise tokens-to-tokens transformation to capture local structure. These approaches usually come with the cost of extra parameters and model complexity, thus potentially lowering the inference speed.

d'Ascoli et al. (2021); Zhou et al. (2021) improve the self-attention for better representation with gated positional self-attention and learnable transformation matrix respectively. Hu et al. (2019),

Ramachandran et al. (2019) adapt the self-attention module and improve the performance of CNN models.

Different from prior works, we try to incorporate spatial locality without changing its architecture or adding extra parameters/FLOPs. We propose attention masks to guide the attention heads to focus on local information. Masked attention heads extract local dependencies more efficiently by allowing information aggregation only from the closest neighbors. This liberates other unmasked heads to learn global information more effectively. We name the modified model, "Masked attention image Transformer" (MaiT), which is built on top of DeiT (Touvron et al., 2021a). MaiT gathers both local and global information at the same time from different heads. Moreover, the regularization effects from attention masks facilitate the training of deep transformers by guiding the attention map learning and promoting diversity across transformer layers.

We proved that less is more in this specific case with attention masks. MaiT achieves up to 1.0% higher top-1 accuracy on ImageNet (Deng et al., 2009) with the same model architecture as DeiT (Touvron et al., 2021a). Additionally, MaiT outperforms CaiT (Touvron et al., 2021b) by up to 1.5% in top-1 accuracy with fewer parameters and simpler structure in deeper transformers.

In summary, we make three major contributions in this work: 1). we propose attention masks to encode spatial locality into self-attention heads without structural change or extra parameter/computation, while improving model efficiency. 2). We present a quick and effective search strategy for the masking scheme exploration. 3). We also reveal the importance of the locality across layers and the impact of attention masks on facilitating the training and convergence of deep transformer models. Note though MaiT is demonstrated with DeiT, the attention mask is applicable to other vision transformers as well.

2 RELATED WORK

Spatial locality is an integral part of the convolutional operation with weight filters attending to local regions of input feature maps. Vision transformer (ViT) (Dosovitskiy et al., 2021) is the first pure transformer-based model on vision tasks, but it requires a large private labeled dataset JFT300M (Sun et al., 2017) to achieve competitive performances. Data-efficient image transformer (DeiT) (Touvron et al., 2021a) improves upon ViT models by introducing stronger data augmentation, regularization, and knowledge distillation. Class-attention in image transformer (CaiT) (Touvron et al., 2021b) extends DeiT by increasing the number of transformer layers. To overcome the difficulties of training deeper transformers, CaiT introduces LayerScale and class-attention layers, which increase the parameters and model complexity.

Tokens-to-Token vision transformer (T2T) (Yuan et al., 2021b) proposes an image transformation by recursive token aggregation to capture local structure. Stand-alone self-attention (Ramachandran et al., 2019) applies local self-attention layer to replace spatial convolution and outperform original ResNet models. Even though sharing value and key spatially is parameter efficient in this approach, content-based information is lost. Transformer-iN-Transformer (TNT) (Han et al., 2021) models both patch-level and pixel-level representations and applies outer and inner transformer blocks to extract global and local information respectively. ConViT (d'Ascoli et al., 2021) proposes the gated positional self-attention to incorporate soft convolutional biases. CrossViT (Chen et al., 2021) proposes a dual-branch transformer architecture for multi-scale feature extraction.

For pixel-level prediction tasks such as semantic segmentation, object detection, Pyramid Vision Transformer (PVT) (Wang et al., 2021a) introduces a progressive shrinking pyramid and spatial-reduction attention with fine-grained image patches. DETR (Carion et al., 2020) adapts transformers for object detection tasks. Swin Transformer (Liu et al., 2021) applies hierarchical transformer with shifted windows of varying sizes. Twins (Chu et al., 2021a) deploys interleaved locally-grouped self-attention and global sub-sample attention layers to improve performances.

To optimize transformer and save computation, Wu et al. (2021b) uses centroid attention to extract and compress input information, Jaegle et al. (2021) iteratively distill inputs into latent space with attention bottlenecks, Wang et al. (2021b) dynamically adjusts the number of tokens with multiple cascading transformers, and Wu et al. (2020) introduced semantic token to replace pixel-based transformers to save computation. There are hybrid architectures fusing convolutional and transformer blocks, such as LocalViT (Li et al., 2021b), BoTNet (Srinivas et al., 2021), LeViT (Graham et al., 2021), BossNet (Li et al., 2021a), CvT (Wu et al., 2021a), CoaT (Xu et al., 2021), CMT (Guo et al., 2021) for higher accuracy and faster inference.

Unlike prior literature, our work explores the intrinsic capability of pure transformer block on incorporating spatial locality and the impact of the locality along the depth direction. This work is also inspired by the emerging graph attention network (Veličković et al., 2018), borrowing the concept of message passing and information aggregation from nearest neighbors.

Note attention masks have been used in NLP tasks as a sparsification method to reduce the computation complexity, as well as to capture local information, as in Guo et al. (2019); Child et al. (2019); Beltagy et al. (2020); Ainslie et al. (2020).

3 VISION TRANSFORMER PRELIMINARIES

The transformer architecture introduced by Vaswani (Vaswani et al., 2017) inspired many model variants with remarkable success in NLP tasks. ViT (Dosovitskiy et al., 2021) extends pure transformer-based architecture into CV applications. Instead of pixel-level processing, ViT splits the original images into a sequence of patches as inputs and transforms them into *patch tokens*, for better computation efficiency. In general, ViT consists of 3 fundamental modules: embedding layer, multi-head self-attention, and feed-forward network.

To process images in transformer, the original (224x224) RGB images is flattened into a sequence of N (16x16) patches. Each patch has a fixed size, typically 16x16 pixels. Patches are then transformed into patch embedding with hidden dimensions (D) of 192, 384, and 768 for tiny, small, and base models respectively in ViT/DeiT. In addition to patch tokens, the embedding layer also integrates positional information, classification and knowledge distillation through the positional token, class token and distillation token, respectively.

Positional token is added into the patch embedding with a trainable positional embedding. However, this positional embedding is added only in the embedding layers. The spatial information is largely lost in the transformer layers since all-to-all attention is invariant to the order of the patches.

The class token is another trainable vector (1xD), concatenated to the patch tokens (total N+1). It is used to collect information from the patch tokens to make output predictions, while also spreading information among patches during training.

Distillation token is sometimes added for knowledge transfer from teacher models, such as a CNN model. When training the distilled version of the model, a distillation token is further concatenated to the patch token along with the class token (total N+2).

Multi-head self-attention (MHA) module has multiple parallel attention heads, where each of them comprises three main components: Key, Query, and Value. Key ((N+1)xd) and Query ((N+1)xd) are trained and multiplied to estimate how much weights on each corresponding token in Value ((N+1)xd) for output ((N+1)xd):

$$Attention(K, Q, V) = Softmax\left(\frac{QK^{T}}{\sqrt{d}}\right)V$$
(1)

Where softmax is applied to each row of the input product matrix (QK^T) and \sqrt{d} provides appropriate normalization. Multiple attention heads in MHA attend to different parts of the input simultaneously. Considering H heads in MHA layer, the hidden dimension D is split equally across all heads (D=Hxd).

Feed-forward network (FFN) follows after the MHA module, containing two linear transformation layers separated by Gelu activation. The hidden dimension expands by 4x after the first linear layer from D to 4D, and is reduced back to D in the second linear layer. Both MHA and FFN use skip-connections with layer normalization as the residual operation.

4 MASKED ATTENTION HEAD

Spatial locality plays a crucial role in computer vision tasks. CNN models capture it using the sliding filter of shared weights, typically with a receptive field of 3x3, 5x5, or 7x7. In contrast to CNN models, the locality is not intrinsically encoded in the transformer structure. With attention masks, we can explicitly insert locality into self-attention modules without introducing any extra parameter or computation. The key idea is to apply a mask on the all-to-all attention products (i.e. QK^T) to reinforce the weights from the closest neighbors and allow information aggregation only from tokens selected by the mask.

	224 = 16 x 14													
1	0	1	2	3	4	5	6	7	8	9	10	11	12	13
t = 16 x 14	14	15—	>16 ↓ ↓ ↓ ↓ ↓ ↓	—17	18	19	20	21	22	23	24	25	26	27
	28	29	30	31	32	33	34	35 —	>36 <	—37	38	39	40	41
	Attention mask forPatch16: 3x3 (depth = 1) 5			47	48	49	50	51	52	53	54	55		
22	56	57	58	59	60	61	62	63	64	65	66	67	56	57
							Attention mask for Patch 72: 5x5 (depth = 2)							
16	182	183	184	185	186	187	188	189	190	191	192	193	194	195
	16													

Figure 1: Attention mask with depth of 1 (orange) and 2 (green).

Figure 1 illustrates an example of our proposed masking scheme. The orange box shows a 3×3 mask, where only the direct neighboring patches are selected. Specifically, Patch 16 only gathers information from the closest neighbors of Patch 1, 2, 3, 15, 17, 29, 30, and 31, and ignores the rest of patches. This is different from the typical all-to-all attention module, where Patch 16 attends to all 0-196 patches. We can easily expand the depth of the attention mask beyond the closest neighbors, to second-level neighbors (green box in Figure 1) or more. Note that the class token (and distillation token) still attends to all the patches to collect and spread information during forward and backward passes. Since each attention product selected by the mask is calculated by Q and K, the masked attention head also retains the content-based locality information.

The attention mask is added before Softmax, regulating the distribution of attention maps to focus more on the closest neighbors:

Masked Attention(K, Q, V) = Softmax
$$\left(\frac{\boldsymbol{M} \odot QK^T}{\sqrt{d}}\right) V$$
 (2)

where $M \in \mathbb{R}^{(N+1)\times(N+1)}$ is a binary attention mask, encoding the spatial locality into the attention head by passing through only the weights from close neighbors and setting the rest to zero.

Note it's important to add the mask before Softmax because it allows the flexibility for the model to learn the importance of the locality. More precisely, unselected patches appear as $e^0 = 1$ in the numerator of the softmax operation. Thus, if the attention product result of the closest neighbors is meaningfully larger than zero (i.e. $M \odot QK^T \gg 0$), it suggests that local information dominates. However, if those results are negative or close to zero, it implies that local information is insignificant and global information is more important. *Therefore, inserting masks before the softmax operation allows models to enforce locality or bypass it.*

4.1 ATTENTION LOCALITY SCORE (ALS)

The softmax operation transfers the QK^T product results into the probability space. As a result, each row of the attention map (A) sums up to 1. For each patch, the probability of focusing on local neighbors equals the sum of their neighbors' attention map weights. If this number approaches one,

the local information is crucial; whereas if it is close to zero, it means global information matters. For Patch n, we define ALS_n as the *Attention Locality Score (ALS)* and the average from all patches (N+1) as ALS for each attention head, as follow:

$$ALS = \frac{\sum_{n} ALS_{n}}{(N+1)}, \text{ where } ALS_{n} = \sum_{i} (\boldsymbol{M} \odot \boldsymbol{A})_{n,i}$$
(3)

where M is the same attention mask defined earlier, and $A = Softmax \left(Mask \odot QK^T / \sqrt{d} \right)$ is the attention map, or $A = Softmax \left(QK^T / \sqrt{d} \right)$ for unmasked heads. We later use the ALS metric to get some insights about the locality behavior of different attention heads in our models.

4.2 MASKING STRATEGY

With total H attention heads, h' number of attention heads can be allocated to focus on local information. The rest of the unaltered attention heads (H - h') capture global dependencies. Therefore, we can extract the local information through masked attention heads and global information through original attention heads at the same time, as shown in Figure 2.



Figure 2: Masked attention heads and original attention heads for local and global dependencies.

Besides the depth of the attention mask, the number of masked attention heads in MHA and the position (which layer to insert mask) are also hyper-parameters. Though masks encode locality into the attention heads, the regularization from masking (pruning attention map) can also limit the learning capacity. Consequently, where to add attention masks requires careful consideration. The complexity to search for the best masking strategy is up to 2^{36} for a 12-layer model with 3 heads, and thus we leave it for future work. Instead, we leverage the insights from ALS study to guide the mask placement, by using a masking strategy as follows:

- 1. Add mask to only one attention head (Head 0) for all layers
- 2. Compute $ALS^{0,i}$ for every layer i
 - (a) If $ALS^{0,i}$ is close to 1: add masks to all attention heads of layer i, proceed to 3.
 - (b) If $ALS^{0,i}$ is close to 0.5: keep the mask, add no more
 - (c) If $ALS^{0,i}$ is close to 0: remove the mask
- 3. Following 2.a), compute $ALS^{h,i}$ for each head and remove the mask if $ALS^{h,i}$ is nearly 0

4.3 CROSS-LAYER SIMILARITY

To evaluate the impact of attention masks on the diversity of the attention maps across all layers, we apply a cross-layer similarity metric similarly defined in DeepViT (Zhou et al., 2021):

$$M_{h,t}^{i,j} = \frac{\mathbf{A}_{h,t,:}^{i} \mathbf{A}_{h,t,:}^{j\top}}{\|\mathbf{A}_{h,t,:}^{i}\|\|\mathbf{A}_{h,t,:}^{i}\|}$$
(4)

where $M_{h,t}^{i,j}$ is the attention map cosine similarity between layer i and layer j for attention head h and token t. $\mathbf{A}_{h,t,:}$ measures the weight contribution for each token in input (Value) for output token T.

Therefore, it reflects the similarity of attention maps on information aggregation across all T input tokens. For example, if $M_{h,t}^{i,j}$ reaches 1, output token t at head h attends to all N+1 tokens in Value with the same probability for both layer i and layer j.

5 **EXPERIMENTS**

In this section, we report the results for MaiT, implemented by applying attention masks on top of DeiT. We consider different hyper parameters related to the masking scheme including mask depth, number of masked heads per layer, and number of transformer layers. We also evaluate the importance of locality in the depth direction and the impact of masks for deep transformers. In our evaluation, we merely focus on models with fewer than 50M parameters, which are more applicable to embedded systems and mobile devices.

5.1 Setup

We follow the same training procedure as DeiT, unless specified otherwise. The implementation is based on the Timm library (Wightman, 2019). Models are trained with ILSVRC-2012 ImageNet dataset (Deng et al., 2009) with the default batch size of 1024 on 4 GPUs for 300 epochs for 12-layer models and 400 epochs for models with more than 12 layers. The model parameters we adopted is summarized in Table 1.

Table 1: Details of MaiT model parameters

Model	Extra tiny (XT)	Tiny (T)	Extra small (XS)	Small (S)
Hidden dim.	144	192	288	384
Heads	3	3	3	6

5.2 IMPACT OF MASK DEPTH

Intuitively, the field of view is larger with an attention mask of larger depth, to gather broader neighborhood information. However, in the extreme case with mask depth of N+1 (standard self-attention), all the patch token information is merged and positional information is lost. On the other hand, attention masks can be considered as a structured pruning on attention maps as well. A smaller depth for the attention mask potentially translates to fewer FLOPs, which might be useful for some hardware devices such as embedded CPUs. For example, with 3x3 mask, instead of calculating all 197 attention weights, we only compute 9 of them, a 95% reduction for attention map computation.

To study the impact of the mask depth, we mask only one attention head for every MHA module, with the field of view as 3x3, 3x5, and 5x5 on DeiT-small model. The results are summarized in Table 2. There is 0.1 - 0.3 % improvement compared to DeiT small. Various mask sizes only lead to marginal differences in accuracy. This is probably because each patch (original 16x16 pixels) is a processed sub-image and already contains some locality information, and thus further away patch tokens barely provide much extra information. Therefore, 3x3 mask is the default setting for MaiT in the subsequent sections.

Table 2: Top-1 accuracy on ImageNet for DeiT and MaiT (one masked head for all 12 layers) with various attention mask depth. MaiT* has two masked heads for all 12 layers and is trained for 400 epochs.

DeiT	MaiT(3x3)	MaiT(3x5)	MaiT(5x5)	MaiT*(3x3)
79.8	79.9(+0.1)	80.1(+0.3)	80.0(+0.2)	80.8(+1.0)

In contrast to DeiT, whose accuracy saturates beyond 300 epochs as reported in (Touvron et al., 2021a), MaiT-small sees another 0.9% improvement when trained for 400 epochs. Accordingly, we

suspect a higher learning rate is beneficial for MaiT since attention masks provide some guidance for attention map training. If we change the batch size from 1024 to 3072 for MaiT-tiny, essentially increasing the learning rate ($lr = 0.0005 \times batchsize/512$) by 3x, the top-1 accuracy for MaiT-tiny increases by 0.4% to 72.6 after 300 epochs. Therefore, the number of training epochs and learning rate for MaiT is subject to further tuning for performance improvement.

5.3 IMPORTANCE OF SPATIAL LOCALITY

To explore the importance of locality across transformer layers, we start by masking one head in each layer and compute the ALS for all the heads in all layers. Figure 3 a) illustrates the result for 12-layer MaiT-small. ALS of the masked head is significantly higher for the first 10 out of 12 layers, than the unmasked heads. This indicates the model learns to focus on local information with the masked attention head and global information with unmasked heads. We observe three stages in the ALS across all layers for the masked head: 1. high ALS, where more than 70% of attention weights are allocated to closest neighbors in the first 4 layers; 2. mid ALS, where the probability to focus on locality decreases to around 50% for Layer 4 to Layer 9; 3. low ALS, where the chance of attending to local patches are less than 0.5% for the last 2 layers. Notice that in the third layer, even for unmasked heads, two heads focus heavily on closest neighbors with ALS larger than 0.7. Similarly, among the unmasked heads, the first 4 layers show higher ALS, which decreases in the following 6 layers and down to an average of 0.0051 (1/(N+1) = 1/197) if the attention weights are distributed uniformly.

This indicates the locality is more important in the early layers (first 4 layers) and becomes trivial in the last two layers, which is consistent with d'Ascoli et al. (2021), Raghu et al. (2021) When the number of transformer blocks increases, the importance of locality expands to deeper layers as shown in Figure 7. The trends remain for deeper MaiT of 24 or 36 layers.

Accordingly in Figure 3 b), we mask all heads for the first one-third of total layers, only one masked head for 8 layers in the middle (8-15), and no mask for the rest of the layers. For fully masked layers, ALS of most heads are larger than 0.5. Though one head in every other layer of the first 8 blocks shows lower than 0.1 ALS. This suggests at least one head is trying to gather global information even in the first 8 layers. Thus it's beneficial to keep one global attention head even for the initial few layers.



Figure 3: Attention Locality Scores (ALS) for MaiT-small: a) 1 masked head for all 12 layers; b) out of 24 layers, 6 masked heads for the first 8 layers, 1 masked head for the middle 8 layers, and no masking for the last 8 layers.

5.4 IMPACT OF ATTENTION MASK ON DEEPER TRANSFORMERS

Besides lacking inductive bias, another issue with transformers is the difficulty to train deeper models. Naively stacking transformer layers fails to deliver the expected performance gain as shown in Touvron et al. (2021b). One reason for the performance degradation is attention collapse, i.e. the attention maps are more alike among deeper layers (Zhou et al., 2021). We find the attention mask is able to break the attention map homogeneity and facilitates the convergence of the deep transformers. The mixed masking scheme breaks the structural repetition in deep transformers and promotes diversity in later layers.



Figure 4: Top-1 accuracy on ImageNet for DeiT (red), MaiT (green), CaiT (yellow) and randomly masked DeiT (blue) with various numbers of transformer blocks. Green square MaiT applies one mask heads for all layers while green star MaiT masks all three heads for the first 12 layers and only one masked attention head for the rest 24 layers.

When increasing transformer blocks from 24 to 36, even though the model capacity is largely increased, top-1 accuracy only increases by 0.25% for DeiT-tiny (Figure 4). With MaiT changing one head to focus on local information, top-1 accuracy is 0.8% higher over DeiT of the same layers up to 32 layers. However, at 36 layers, the performance of MaiT degrades to the same level as DeiT. Beyond 32 layers, a mixed masking strategy can be adopted to extend the performance gain. If the 36-layer MaiT applies fully 3 masked heads in the first 12 layers and only one masked head for the rest 24 layers, the top-1 accuracy increases by 1.2% (from 78.6% to 79.8%).

Interestingly, we noticed that for 24-layer transformers, CaiT shows even worse performances than naively stacked DeiT-24, suggesting that the hyper-parameters for 24-layer CaiT-tiny are probably not optimal.

To rule out the impact of pure pruning effects on attention maps from masking, we also apply a random mask of the same drop-out rate to one head of all 24 layers. The randomly masked model leads to 0.5% accuracy drop as compared to 24-layer DeiT. This provides further proof that the performance gain with attention masks is indeed from better locality information aggregation instead of pruning on attention maps.

To better understand the impact of attention masks in deep MaiT, we compare the cross layer similarity heat map among 36 layers of DeiT, MaiT (one masked head), and MaiT (mixed masking) in Figure 8, Figure 9, and Figure 5 respectively. Each point in the heat map denotes the averaged similarity between the y-axis indexed layer and the layer from x-axis.

DeiT shows the highest cross-layer similarity at the last 8 stages, with an average of 0.75 (Figure 8). This homogeneity in attention maps at the late stage limits the model learning capability. Adding mask to one head for all 36 layers does not alleviate this problem since the same structure is repeated 36 times in the model. As a result, the average cross-layer similarity for the last 8 layers with unmasked heads is also 0.75 (Figure 9). Alternatively, if we mask all heads in the first 8 layers, they are naturally different from the rest of the partially masked layers, and the same structure is only repeated 8 or 24 times. As shown in Figure 5, cross-layer similarity between the first 8 layers and the following 24 layer are mostly well below 0.3 for Head 1 and 2. Additionally, a mixed masking scheme lower the cross-layer similarity among the last 8 layers by 0.1 to 0.65.



Figure 5: Similarity across 36 layers for MaiT-tiny with mixed masking scheme

5.5 COMPARISON WITH DEIT AND CAIT

In Table 3, we find the attention mask is more effective with small hidden dimensions on deeper transformer models. To improve DeiT-T, we introduce MaiT-XT (24 layers) with similar parameters and FLOPs, achieving 1.5% higher top-1 accuracy. 24-layer MaiT-T outperforms CaiT-T by 1.5% even with fewer parameters and less FLOPs. However, with increased hidden dimension, MaiT-XS and MaiT-S are 0.8% and 0.7% lower than CaiT-XS and CaiT-S, respectively. Note that CaiT applies multiple optimization techniques to improve accuracy. In comparison, by simply adding attention masks, MaiT-S improves the accuracy by 1.1% and thus bridges the gap between CaiT-S and naively stacked DeiT-S. Moreover, if we add LayerScale to MaiT-S, its accuracy surpasses CaiT-S by 0.1% with fewer parameters and FLOPs. It's promising for more gain when combining other optimization techniques with MaiT.

Model	Layers	Params(M)	GFLOPs	Top-1 Acc.
DeiT-T	12	5.7	1.3	72.2
MaiT-XT	24	6.3	1.5	73.7
CaiT-T	24+2	12	2.5	77.6
MaiT-T	24	11	2.5	79.1
DeiT-S	12	22	4.6	79.8
CaiT-XS	24+2	26.6	5.4	81.8
MaiT-XS	24	24.5	5.3	81.0
CaiT-S	24+2	46.9	9.4	82.7
DeiT-S ⁺	24	43.3	9.1	81.0
MaiT-S	24	43.3	9.1	82.1
MaiT-S*	24	43.3	9.1	82.8

Table 3: Top-1 accuracy for DeiT, MaiT, and CaiT. Hyper parameters for MaiT are listed in the Appendix. *Added LayerScale on MaiT. The accuracy for DeiT-S⁺ is from Touvron et al. (2021b).

6 CONCLUSION

In this work, we incorporate spatial locality into vision transformers by inserting masks to attention heads. Masked heads are able to focus on local information and liberate other unmasked heads to extract global information more effectively. We also introduce attention locality score to guide the masking strategy search and rapidly evaluate various masking schemes. Attention masking is a simple and effective technique, especially for smaller models (i.e., < 6M parameters). We observe that attention masking also serves as a regularizer to guide the training of attention maps for deeper transformers. Moreover, it is promising to further boost performances by combining this attention mask technique with other optimization approaches.

REFERENCES

- Joshua Ainslie, Santiago Ontanon, Chris Alberti, Vaclav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. Etc: Encoding long and structured inputs in transformers, 2020.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer, 2020.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers, 2020.
- Chun-Fu Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification, 2021.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers, 2019.
- 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, 2021a.
- Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers, 2021b.
- Stéphane d'Ascoli, Hugo Touvron, Matthew Leavitt, Ari Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases, 2021.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pp. 248–255. Ieee, 2009.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. NAACL, 2019.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- Ben Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet's clothing for faster inference, 2021.
- Jianyuan Guo, Kai Han, Han Wu, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Cmt: Convolutional neural networks meet vision transformers, 2021.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Yunfan Shao, Xiangyang Xue, and Zheng Zhang. Startransformer, 2019.
- Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer, 2021.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, pp. 770–778, 2016.
- Han Hu, Zheng Zhang, Zhenda Xie, and Stephen Lin. Local relation networks for image recognition, 2019.
- Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver: General perception with iterative attention, 2021.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. NIPS, 25:1097–1105, 2012.
- Changlin Li, Tao Tang, Guangrun Wang, Jiefeng Peng, Bing Wang, Xiaodan Liang, and Xiaojun Chang. Bossnas: Exploring hybrid cnn-transformers with block-wisely self-supervised neural architecture search, 2021a.

- Yawei Li, Kai Zhang, Jiezhang Cao, Radu Timofte, and Luc Van Gool. Localvit: Bringing locality to vision transformers, 2021b.
- 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, 2021.
- Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks?, 2021.
- Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Aravind Srinivas, Tsung-Yi Lin, Niki Parmar, Jonathon Shlens, Pieter Abbeel, and Ashish Vaswani. Bottleneck transformers for visual recognition, 2021.
- Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era, 2017.
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pp. 6105–6114. PMLR, 2019.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *ICML*, volume 139, pp. 10347–10357, July 2021a.
- Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239*, 2021b.
- 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, pp. 5998–6008, 2017.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks, 2018.
- 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, 2021a.
- 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, 2021b.
- Ross Wightman. Pytorch image models. https://github.com/rwightman/ pytorch-image-models, 2019.
- Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph Gonzalez, Kurt Keutzer, and Peter Vajda. Visual transformers: Token-based image representation and processing for computer vision, 2020.
- Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers, 2021a.
- Lemeng Wu, Xingchao Liu, and Qiang Liu. Centroid transformers: Learning to abstract with attention, 2021b.
- Weijian Xu, Yifan Xu, Tyler Chang, and Zhuowen Tu. Co-scale conv-attentional image transformers, 2021.
- Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating convolution designs into visual transformers, 2021a.
- Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zihang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet, 2021b.
- Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer, 2021.

A APPENDIX

A.1 ALS OF DEIT AND DEEP MAIT

When trained without restriction from attention mask, DeiT-small emphasizes the locality by allocating more than 50% of attention weights on closest neighbors for 3 heads at second and third layer, even close to 90% at the second layer as shown in Figure 6. In comparison, in Figure 3, MaiT utilizes masked Head 0 to extract local information from the first layer to the ninth layer. Another notable change between MaiT and DeiT is the ALS from Layer 4 to 9, masked head 0 elevates ALS above other heads to focus on local information. Head 1 to Head 5 in MaiT show lower ALS compared to DeiT, and focus more on global information.



Figure 6: ALS for DeiT-small for all 6 attention heads.



Figure 7: ALS of deeper MaiT-tiny with 12, 24 and 36 layers.

When the number of transformer blocks increases, the importance of locality expands to deeper layers as shown in Figure 7. For the masked Head 0 in MaiT-tiny, 12-layer model shows more than 0.8 ALS for 6 layers out of the first 7 layers, the following 4 layers give a ALS around 0.5, and the ALS of the last layer declines below 0.05. Likewise, for 24-layer and 36 layer models, 10 layers out of the first 12 layers show ALS higher than 0.8, and decreases to 0.5 or 0.6 in the following layers before the last layer.

A.2 HYPER PARAMETERS FOR MAIT

Table 4: Hyper parameters for 24-layer MaiT in Table 3. Drop rate refers to linear stochastic depth drop rate, the same as in Touvron et al. (2021a). *denotes the constant drop out for all layers.

Model	Masked Head/Layer	Drop rate	Batch size
MaiT-XT	3/0-7 + 1/8-21	0.0	1024
MaiT-T	1/0-23	0.1	2048
MaiT-XS	1/0-23	0.1	1024
MaiT-S	5/0-7 + 1/8-19	0.1*	1024

A.3 CROSS LAYER SIMILARITY HEAT MAP

In Figure 8, 36-layer DeiT-tiny shows the highest similarity at the last 8 stages across all three heads, with an average similarity of 0.75, whereas the similarity is quite small among the first 16 layers and between the first 16 layers and the last 20 layers. This is agreement with Zhou et al. (2021), indicating an attention collapse at later stage of deep transformers.



Figure 8: Similarity across 36 layers for DeiT-tiny



Figure 9: Similarity across 36 layers for MaiT-tiny

In comparison, for MaiT with one masked head across all layers, the masked attention Head 0 shows a more uniform distribution with a lower peak intensity as compared to the the non-masked or partially masked (in early layers) Head 1 and Head 2 (Figure 9. For unmasked Head 1 and Head 2, later stage especially the last 8 layers show highest similarity among each other, the average of them is also 0.75. This suggests one masked head for all transformer blocks is not as effective in

36-layer model as in 24-layer model. This indicates repeating the same structure for 36 times in the depth direction is causing the attention collapse. The diversity among the heads within the same layer doesn't alleviate this issue.