

Evo-ViT: Slow-Fast Token Evolution for Dynamic Vision Transformer

Yifan Xu^{1,3,4*}, Zhijie Zhang^{2,3*}, Mengdan Zhang³, Kekai Sheng³, Ke Li³, Weiming Dong^{1,4†},
Liqing Zhang², Changsheng Xu^{1,4}, Xing Sun^{3†}

¹NLPR, Institute of Automation, Chinese Academy of Sciences ²Shanghai Jiao Tong University ³Tencent Youtu Lab

⁴School of Artificial Intelligence, University of Chinese Academy of Sciences

{xuyifan2019, weiming.dong, changsheng.xu}@ia.ac.cn, zzj506506@sjtu.edu.cn, zhang-lq@cs.sjtu.edu.cn,
{davinazhang, saulsheng, tristanli}@tencent.com, winfred.sun@gmail.com

Abstract

Vision transformers (ViTs) have attracted considerable research attention recently, but the huge computational cost is still a severe issue. A mainstream paradigm for computation reduction aims to reduce the number of tokens given that the computation complexity of ViT is quadratic with respect to the input sequence length. Existing designs include structured spatial compression that uses a progressive shrinking pyramid to reduce the computations of large feature maps, and unstructured token pruning that dynamically drops redundant tokens. However, limitations of existing token pruning lie in the following aspects: 1) the incomplete spatial structure caused by pruning is incompatible with structured spatial compression that is commonly used in modern deep-narrow transformers; 2) it usually requires a time-consuming pre-training procedure. To address the limitations and expand the applicable scenario of token pruning, we present Evo-ViT, a self-motivated slow-fast token evolution approach for vision transformers. Specifically, we conduct unstructured instance-wise token selection by taking advantage of the simple and effective global class attention that is native to vision transformers. Then, we propose to update the selected informative tokens and uninformative tokens with different computation paths, namely, slow-fast updating. Since slow-fast updating mechanism maintains the spatial structure and information flow, Evo-ViT can accelerate vanilla transformers of both flat and deep-narrow structures from the very beginning of the training process. Experimental results demonstrated that our method significantly reduces the computational cost of vision transformers while maintaining comparable performance on image classification. For example, our method accelerates DeiT-S by over 60% throughput while only sacrificing 0.4% top-1 accuracy on ImageNet-1K, outperforming current token pruning methods on both accuracy and efficiency. Code is available at <https://github.com/YifanXu74/Evo-ViT>.

Introduction

Vision transformers (ViTs) have shown strong power on various computer vision tasks (Xu et al. 2022). The reason of introducing the transformer into computer vision lies in its unique properties that convolution neural networks (CNNs)

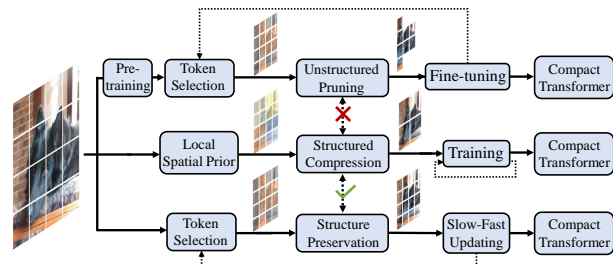


Figure 1: An illustration of technique pipelines for computation reduction via tokens. The dash lines denote iterative training. The first branch: the pipeline of unstructured token pruning (Rao et al. 2021; Tang et al. 2021) based on pre-trained models. The second branch: the pipeline of structured compression (Graham et al. 2021). The third branch: our proposed pipeline that performs unstructured updating while suitable for structured compressed models.

lack, especially the property of modeling long-range dependencies. However, dense modeling of long-range dependencies among image tokens brings computation inefficiency, because images contain large regions of low-level texture and uninformative background.

Existing methods follow two processes to address the inefficiency problem of modeling long-range dependencies among tokens in ViT as shown in the above two branches of Fig. 1. The first process, as shown in the second branch, is to perform structured compression based on local spatial prior, such as local linear projection (Wang et al. 2021a), convolutional projection (Heo et al. 2021), and shift windows (Liu et al. 2021). Most modern transformers with deep-narrow structures are typically within this process. However, the structured compressed models treat the informative object tokens and uninformative background tokens with the same priority. Thus, token pruning, the second process, proposes to identify and drop the uninformative tokens in an unstructured way. (Tang et al. 2021) improves the efficiency of a pre-trained transformer network by developing a top-down layer-by-layer token slimming approach that can identify and remove redundant tokens based on the reconstruction error of the pre-trained network. The trained pruning mask is fixed for all instances. (Rao et al. 2021) proposes to accelerate a pre-trained transformer network by removing redun-

*The first two authors contributed equally.

†The co-corresponding authors.

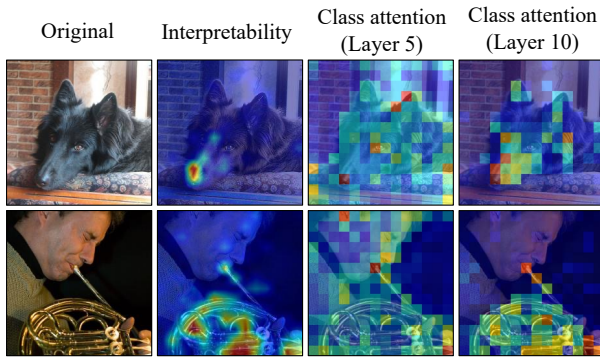


Figure 2: Visualization of class attention in DeiT-T. The *interpretability* comes from (Chefer, Gur, and Wolf 2021).

dant tokens hierarchically, and explores an data-dependent down-sampling strategy via self-distillation. Despite of the significant acceleration, these unstructured token pruning methods are restricted in two aspects due to their incomplete spatial structure and information flow, namely, the inapplicability on structured compressed transformers and inability to train from scratch.

In this paper, as shown in the third branch of Fig. 1, we propose to handle the inefficiency problem in a dynamic data-dependent way from the very beginning of the training process while suitable for structured compression methods. We denote uninformative tokens that contribute little to the final prediction but bring computational cost when bridging redundant long-range dependencies as placeholder tokens. Different from structured compression that reduces local spatial redundancy in (Wang et al. 2021a; Graham et al. 2021), we propose to distinguish the informative tokens from the placeholder tokens for each instance in an unstructured and dynamic way, and update the two types of tokens with different computation paths. Instead of searching for redundancy and pruning in a pre-trained network such as (Tang et al. 2021; Rao et al. 2021), by preserving placeholder tokens, the complete spatial structure and information flow can be maintained. In this way, our method can be a generic plug-in for most ViTs of both flat and deep-narrow structures from the very beginning of training.

Concretely, Evo-ViT, a self-motivated slow-fast token evolution approach for dynamic ViTs is proposed in this work. “Self-motivated” means that transformers can naturally distinguish informative tokens from placeholder tokens for each instance, since they have insights into global dependencies among image tokens. Without loss of generality, we take DeiT (Touvron et al. 2021) in Fig. 2 as example. We find that the class token of DeiT-T estimates importance of each token for dependency modeling and final classification. Especially in deeper layers (*e.g.*, layer 10), the class token usually augments informative tokens with higher attention scores and has a sparse attention response, which is quite consistent to the visualization result provided by (Chefer, Gur, and Wolf 2021) for transformer interpretability. In shallow layers (*e.g.*, layer 5), the attention of the class token is relatively scattered but mainly focused on informative regions. Thus, taking advantage of class to-

kens, informative tokens and placeholder tokens are determined. The preserved placeholder tokens ensure complete information flow in shallow layers of a transformer for modeling accuracy. After the two kinds of tokens are determined, they are updated in a slow-fast approach. Specifically, the placeholder tokens are summarized to a representative token that is evolved via the full transformer encoder simultaneously with the informative tokens in a slow and elaborate way. Then, the evolved representative token is exploited to fast update the placeholder tokens for more representative features.

We evaluate the effectiveness of the proposed Evo-ViT method on two kinds of baseline models, namely, transformers of flat structures such as DeiT (Touvron et al. 2021) and transformers of deep-narrow structures such as LeViT (Graham et al. 2021) on ImageNet (Deng et al. 2009) dataset. Our self-motivated slow-fast token evolution method allows the DeiT model to improve inference throughput by 40%-60% and further accelerates the state-of-the-art efficient transformer LeViT while maintaining comparable performance.

Related Work

Vision Transformer. Recently, a series of transformer models (Han et al. 2020; Xu et al. 2022) are proposed to solve various computer vision tasks. The transformer has achieved promising success in image classification (Dosovitskiy et al. 2021; Touvron et al. 2021; D’Ascoli et al. 2021), object detection (Carion et al. 2020; Liu et al. 2021; Zhu et al. 2020) and instance segmentation (Duke et al. 2021; Zheng et al. 2021) due to its significant capability of modeling long-range dependencies. Vision Transformer (ViT) (Dosovitskiy et al. 2021) is among the pioneering works that achieve state-of-the-art performance with large-scale pre-training. DeiT (Touvron et al. 2021) manages to tackle the data-inefficiency problem in ViT by simply adjusting training strategies and adding an additional token along with the class token for knowledge distillation. To achieve better accuracy-speed trade-offs for general dense prediction, recent works (Yuan et al. 2021; Graham et al. 2021; Wang et al. 2021a) design transformers of deep-narrow structures by adopting sub-sampling operation (*e.g.*, strided down-sampling, local average pooling, convolutional sampling) to reduce the number of tokens in intermediate layers.

Redundancy Reduction. Transformers take high computational cost because the multi-head self-attention (MSA) requires quadratic time complexity and the feed forward network (FFN) increases the dimension of latent features. The existing acceleration methods for transformers can be mainly categorized into sparse attention mechanism, knowledge distillation (Sanh et al. 2019), and pruning. The sparse attention mechanism includes, for example, low rank factorization (Xu et al. 2021; Wang et al. 2020), fixed local patterns (Liu et al. 2021), and learnable patterns (Tay et al. 2020; Beltagy, Peters, and Cohan 2020). (Yue et al. 2021) handles the inefficiency problem by sparse input patch sampling. (He et al. 2020) proposes to add an evolved global attention to the attention matrix in each layer for a better residual mechanism. Motivated by this work, we propose the

evolved global class attention to guide the token selection in each layer. The closest paradigm to this work is token pruning. (Tang et al. 2021) presents a top-down layer-by-layer patch slimming algorithm to reduce the computational cost in pre-trained vision transformers. The patch slimming scheme is conducted under a careful control of the feature reconstruction error, so that the pruned transformer network can maintain the original performance with lower computational cost. (Rao et al. 2021) devises a lightweight prediction module to estimate the importance score of each token given the current features of a pre-trained transformer. The module is plugged into different layers to prune placeholder tokens in a unstructured way and is supervised by a distillation loss based on the predictions of the original pre-trained transformer. Different from these pruning works, we proposed to preserve the placeholder tokens, and update the informative tokens and placeholder tokens with different computation paths; thus our method can achieve better performance and be suitable for various transformers due to the complete spatial structure. In addition, the complete information flow allows us to accelerate transformers with scratch training.

Preliminaries

ViT (Dosovitskiy et al. 2021) proposes a simple tokenization strategy that handles images by reshaping them into flattened sequential patches and linearly projecting each patch into latent embedding. An extra class token (CLS) is added to the sequence and serves as the global image representation. Moreover, since self-attention in the transformer encoder is position-agnostic and vision applications highly require position information, ViT adds position embedding into each token, including the CLS token. Afterwards, all tokens are passed through stacked transformer encoders and the CLS token is used for final classification.

The transformer is composed of a series of stacked encoders where each encoder consists of two modules, namely, a multi-head self-attention (MSA) module and a feed forward network (FFN) module. The FFN module contains two linear transformations with an activation function. The residual connections are employed around both MSA and FFN modules, followed by layer normalization (LN). Given the input x_0 of ViT, the processing of the k -th encoder can be mathematically expressed as

$$\begin{aligned} x_0 &= [x_{cls} | x_{patch}] + x_{pos}, \\ y_k &= x_{k-1} + MSA(LN(x_{k-1})), \\ x_k &= y_k + FFN(LN(y_k)), \end{aligned} \quad (1)$$

where $x_{cls} \in \mathbb{R}^{1 \times C}$ and $x_{patch} \in \mathbb{R}^{N \times C}$ are CLS and patch tokens respectively and $x_{pos} \in \mathbb{R}^{(1+N) \times C}$ denotes the position embedding. N and C are the number of patch tokens and the dimension of the embedding.

Specifically, a self-attention (SA) module projects the input sequences into query, key, value vectors (*i.e.*, $Q, K, V \in \mathbb{R}^{(1+N) \times C}$) using three learnable linear mapping W_Q, W_K and W_V . Then, a weighted sum over all values in the sequence is computed through:

$$SA(Q, K, V) = Softmax\left(\frac{Q \cdot K^T}{\sqrt{C}}\right)V. \quad (2)$$

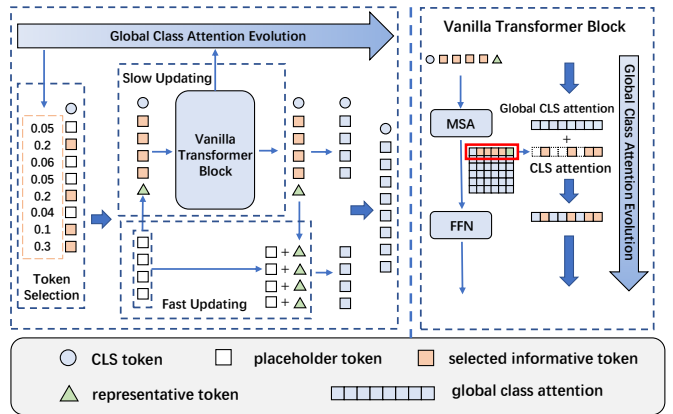


Figure 3: The overall diagram of the proposed method.

MSA is an extension of SA. It splits queries, keys, and values for h times and performs the attention function in parallel, then linearly projects their concatenated outputs.

It is worth noting that one very different design of ViT from CNNs is the CLS token. The CLS token interacts with patch tokens at each encoder and summarizes all the patch tokens for the final representation. We denote the similarity scores between the CLS token and patch tokens as class attention A_{cls} , formulated as:

$$A_{cls} = Softmax\left(\frac{q_{cls} \cdot K^T}{\sqrt{C}}\right), \quad (3)$$

where q_{cls} is the query vector of the CLS token.

Computational complexity. In ViT, the computational cost of the MSA and FFN modules are $O(4NC^2 + 2N^2C)$ and $O(8NC^2)$, respectively. For pruning methods (Rao et al. 2021; Tang et al. 2021), by pruning $\eta\%$ tokens, at least $\eta\%$ FLOPs in the FFN and MSA modules can be reduced. Our method can achieve the same efficiency while suitable for scratch training and versatile downstream applications.

Methodology

Overview

In this paper, we aim to handle the inefficiency modeling issue in each input instance from the very beginning of the training process of a versatile transformer. As shown in Fig 3, the pipeline of Evo-ViT mainly contains two parts: the structure preserving token selection module and the slow-fast token updating module. In the structure preserving token selection module, the informative tokens and the placeholder tokens are determined by the evolved global class attention, so that they can be updated in different manners in the following slow-fast token updating module. Specifically, the placeholder tokens are summarized and updated by a representative token. The long-term dependencies and feature richness of the representative token and the informative tokens are evolved via the MSA and FFN modules.

We first elaborate on the proposed structure preserving token selection module. Then, we introduce how to update the informative tokens and the placeholder tokens in a slow-fast approach. Finally, the training details, such as the loss and other training strategies, are introduced.

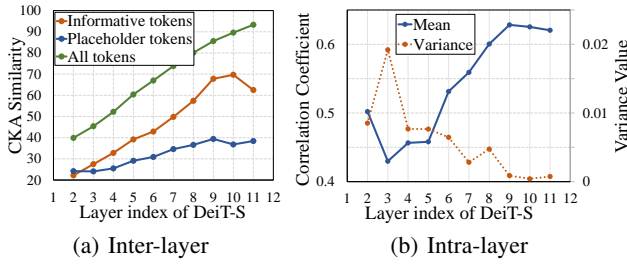


Figure 4: Two folds that illustrate the difficulty of pruning the shallow layers. (a) The CKA similarity between the final CLS token and token features in each layer. (b) The Pearson correlation coefficient of the token features in each layer.

Structure preserving token selection

In this work, we propose to preserve all the tokens and dynamically distinguish informative tokens and placeholder tokens for complete information flow. The reason is that it is not trivial to prune tokens in shallow and middle layers of a vision transformer, especially in the beginning of the training process. We explain this problem in both inter-layer and intra-layer ways. First, shallow and middle layers usually present fast growing capability of feature representation. Pruning tokens brings severe information loss. Following Refiner (Zhou et al. 2021), we use centered kernel alignment (CKA) similarity (Kornblith et al. 2019) to measure similarity of the intermediate token features in each layer and the final CLS token, assuming that the final CLS token is strongly correlated with classification. As shown in Fig. 4(a), the token features of DeiT-T keep evolving fast when the model goes deeper and the final CLS token feature is quite different from token features in shallow layers. It indicates that the representations in shallow or middle layers are insufficiently encoded, which makes token pruning quite difficult. Second, tokens have low correlation with each other in the shallow layers. We evaluate the Pearson correlation coefficient (PCC) among different patch token queries with respect to the network depth in the DeiT-S model to show redundancy. As shown in Fig. 4(b), the lower correlation with larger variance in the shallow layers also proves the difficulty to distinguish redundancy in shallow features.

The attention weight is the easiest and most popular approach (Abnar and Zuidema 2020; Wang et al. 2021b) to interpret a model’s decisions and to gain insights about the propagation of information among tokens. The class attention weight described in Eqn. 3 reflects the information collection and broadcast processes of the CLS token. We find that our proposed evolved global class attention is able to be a simple measure to help dynamically distinguish informative tokens and placeholder tokens in a vision transformer. In Fig. 4(a), the distinguished informative tokens have high CKA correlation with the final CLS token, while the placeholder tokens have low CKA correlation. As shown in Fig. 2, the global class attention is able to focus on the object tokens, which is consistent to the visualization results of (Chefer, Gur, and Wolf 2021). In the following part of this section, detailed introduction of our structure preserving token selection method is provided.

As discussed in Preliminaries Section, the class attention A_{cls} is calculated by Eqn. 3. We select k tokens whose scores in the class attention are among the top k as the informative tokens. The remaining $N - k$ tokens are recognized as placeholder tokens that contain less information. Different from token pruning, the placeholder tokens are kept and fast-updated rather than dropped.

For better capability of capturing the underlying information among tokens in different layers, we propose a global class attention that augments the class attention by evolving it across layers as shown in Fig. 3. Specifically, a residual connection between class attention of different layers is designed to facilitate the attention information flow with some regularization effects. Mathematically,

$$A_{cls,g}^k = \alpha \cdot A_{cls,g}^{k-1} + (1 - \alpha) \cdot A_{cls}^k, \quad (4)$$

where $A_{cls,g}^k$ is the global class attention in the k -th layer, and A_{cls}^k is the class attention in the k -th layer. We use $A_{cls,g}^k$ for the token selection in the $(k+1)$ -th layer for stability and efficiency. For each layer with token selection, only the global class attention scores of the selected informative tokens are updated.

Slow-fast token updating

Once the informative tokens and the placeholder tokens are determined by the global class attention, we propose to update tokens in a slow-fast way instead of harshly dropping placeholder tokens as (Tang et al. 2021; Rao et al. 2021). As shown in Fig. 3, informative tokens are carefully evolved via MSA and FFN modules, while placeholder tokens are coarsely summarized and updated via a representative token. We introduce our slow-fast token updating strategy mathematically as follows.

For N patch tokens x_{patch} , we first split them into k informative tokens $x_{inf} \in \mathbb{R}^{k \times C}$ and $N - k$ placeholder tokens $x_{ph} \in \mathbb{R}^{(N-k) \times C}$ by the above-mentioned token selection strategy. Then, the placeholder tokens x_{ph} are aggregated into a representative token $x_{rep} \in \mathbb{R}^{1 \times C}$, as follows:

$$x_{rep} = \phi_{agg}(x_{ph}), \quad (5)$$

where $\phi_{agg} : \mathbb{R}^{(N-k) \times C} \rightarrow \mathbb{R}^{1 \times C}$ denotes an aggregating function, such as weighted sum or transposed linear projection (Tolstikhin et al. 2021). Here we use weighted sum based on the corresponding global attention score in Eqn. 4.

Then, both the informative tokens x_{inf} and the representative token x_{rep} are fed into MSA and FFN modules, and their residuals are recorded as $x_{inf}^{(*)}$ and $x_{rep}^{(*)}$ for skip connections, which can be denoted by:

$$\begin{aligned} x_{inf}^{(1)}, x_{rep}^{(1)} &= MSA(x_{inf}, x_{rep}), \\ x_{inf} &\leftarrow x_{inf} + x_{inf}^{(1)}, x_{rep} \leftarrow x_{rep} + x_{rep}^{(1)}, \\ x_{inf}^{(2)}, x_{rep}^{(2)} &= FFN(x_{inf}, x_{rep}), \\ x_{inf} &\leftarrow x_{inf} + x_{inf}^{(2)}. \end{aligned} \quad (6)$$

Thus, the informative tokens x_{inf} and the representative token x_{rep} are updated in a slow and elaborate way.

Finally, the placeholder tokens x_{ph} are updated in a fast way by the residuals of x_{rep} :

$$x_{ph} \leftarrow x_{ph} + \phi_{exp}(x_{rep}^{(1)}) + \phi_{exp}(x_{rep}^{(2)}), \quad (7)$$

where $\phi_{exp} : \mathbb{R}^{1 \times C} \rightarrow \mathbb{R}^{(N-k) \times C}$ denotes an expanding function, such as simple copy in our method.

It is worth noting that the placeholder tokens are fast updated by the residuals of x_{rep} rather than the output features. In fact, the fast updating serves as a skip connection for the placeholder tokens. By utilizing residuals, we can ensure the output features of the slow updating and fast updating modules within the same order of magnitude.

Training Strategies

Layer-to-stage training schedule. Our proposed token selection mechanism becomes increasingly stable and consistent as the training process. Fig. 5 shows that the token selection results of a well-trained transformer turn to be consistent across different layers; thereby indicating that the transformer tends to augment informative tokens with computing resource as much as possible, namely the full transformer networks. Thus, we propose a layer-to-stage training strategy for further efficiency. Specifically, we conduct the token selection and slow-fast token updating layer by layer at the first 200 training epochs. During the remaining 100 epochs, we only conduct token selection at the beginning of each stage, and then slow-fast updating is normally performed in each layer. For transformers with flat structure such as DeiT, we manually arrange four layers as one stage.

Assisted CLS token loss. Although many state-of-the-art vision transformers (Wang et al. 2021a; Graham et al. 2021) remove the CLS token and use the final average pooled features for classification, it is not difficult to add a CLS token in their models for our token selection strategy. We empirically find that the ability of distinguishing two types of tokens of the CLS token as illustrated in Fig. 2 is kept in these models even without supervision on the CLS token. For better stability, we calculate classification losses based on the CLS token together with the final average pooled features during training. Mathematically,

$$\begin{aligned} \hat{y}_{cls}, \hat{y} &= m(x_{cls}, x_{patch}), \\ \mathcal{L} &= \phi(\hat{y}_{cls}, y) + \phi(Avg(\hat{y}), y), \end{aligned} \quad (8)$$

where y is the ground-truth of x_{cls} and x_{patch} ; m denotes the transformer model; ϕ is the classification metric function, usually realized by the cross-entropy loss. During inference, the final average pooled features are used for classification and the CLS token is only used for token selection.

Experiments

Setup

In this section, we demonstrate the superiority of the proposed Evo-ViT approach through extensive experiments on the ImageNet-1k (Deng et al. 2009) classification dataset. To demonstrate the generalization of our method, we conduct experiments on vision transformers of both flat and deep-narrow structures, *i.e.*, DeiT (Touvron et al. 2021) and

Table 1: Comparison with existing token pruning methods on DeiT. The image resolution is 224×224 unless specified. * denotes that the image resolution is 384×384 .

Method	Top-1 Acc. (%)	Throughput (img/s)	(%)
DeiT-T			
Baseline (Touvron et al. 2021)	72.2	2536	-
PS-ViT (Tang et al. 2021)	72.0	3563	40.5
DynamicViT (Rao et al. 2021)	71.2	3890	53.4
SViTE (Chen et al. 2021)	70.1	2836	11.8
Evo-ViT (ours)	72.0	4027	58.8
DeiT-S			
Baseline (Touvron et al. 2021)	79.8	940	-
PS-ViT (Tang et al. 2021)	79.4	1308	43.6
DynamicViT (Rao et al. 2021)	79.3	1479	57.3
SViTE (Chen et al. 2021)	79.2	1215	29.3
IA-RED ² (Pan et al. 2021)	79.1	1360	44.7
Evo-ViT (ours)	79.4	1510	60.6
DeiT-B			
Baseline (Touvron et al. 2021)	81.8	299	-
Baseline* (Touvron et al. 2021)	82.8	87	-
PS-ViT (Tang et al. 2021)	81.5	445	48.8
DynamicViT (Rao et al. 2021)	80.8	454	51.8
SViTE (Chen et al. 2021)	82.2	421	40.8
IA-RED ² (Pan et al. 2021)	80.9	453	42.9
IA-RED ^{2*} (Pan et al. 2021)	81.9	129	51.5
Evo-ViT (ours)	81.3	462	54.5
Evo-ViT* (ours)	82.0	139	59.8

LeViT (Graham et al. 2021). Following (Graham et al. 2021), we train LeViT with distillation and without batch normalization fusion. We apply the position embedding in (Wang et al. 2021a) to LeViT for better efficiency. For overall comparisons with the state-of-the-art methods (Rao et al. 2021; Tang et al. 2021; Chen et al. 2021; Pan et al. 2021), we conduct the token selection and slow-fast token updating from the fifth layer of DeiT and the third layer (excluding the convolution layers) of LeViT, respectively. The selection ratios of informative tokens in all selected layers of both DeiT and LeViT are set to 0.5. The global CLS attention trade-off α in Eqn. 4 are set to 0.5 for all layers. For fair comparisons, all the models are trained for 300 epochs.

Main Results

Comparisons with existing pruning methods. In Table 1, we compare our method with existing token pruning methods (Rao et al. 2021; Pan et al. 2021; Tang et al. 2021; Chen et al. 2021). Since token pruning methods are unable to recover the 2D structure and are usually designed for transformers with flat structures, we comprehensively conduct the comparisons based on DeiT (Touvron et al. 2021) on ImageNet dataset. We report the top-1 accuracy and throughput for performance evaluation. The throughput is measured on a single NVIDIA V100 GPU with batch size fixed to 256, which is the same as the setting of DeiT. Results indicate that our method outperforms previous token pruning methods on both accuracy and efficiency. Our method accelerates

Table 2: Comparison with state-of-the-art vision transformers. The input image resolution is 224×224 unless specified. * denotes that the image resolution is 384×384 .

Model	Param (M)	Throughput (img/s)	Top-1 Acc. (%)
LeViT-128S	7.8	8755	74.5
LeViT-128	9.2	6109	76.2
LeViT-192	10.9	4705	78.4
PVTv2-B1	14.0	1225	78.7
CoaT-Lite Tiny	5.7	1083	76.6
PiT-Ti	4.9	3030	73.0
Evo-LeViT-128S	7.8	10135	73.0
Evo-LeViT-128	9.2	8323	74.4
Evo-LeViT-192	11.0	6148	76.8
LeViT-256	18.9	3357	80.1
LeViT-256*	19.0	906	81.8
PVTv2-B2	25.4	687	82.0
PiT-S	23.5	1266	80.9
Swin-T	29.4	755	81.3
CoaT-Lite Small	20.0	550	81.9
Evo-LeViT-256	19.0	4277	78.8
Evo-LeViT-256*	19.2	1285	81.1
LeViT-384	39.1	1838	81.6
LeViT-384*	39.2	523	82.8
PVTv2-B3	45.2	457	83.2
PiT-B	73.8	348	82.0
Evo-LeViT-384	39.3	2412	80.7
Evo-LeViT-384*	39.6	712	82.2

the inference throughput by over 60% with negligible accuracy drop (-0.4%) on DeiT-S.

Comparisons with state-of-the-art ViTs. Owing to the preserved placeholder tokens, our method guarantees the spatial structure that is indispensable for most existing modern ViT networks. Thus, we further apply our method to the state-of-the-art efficient transformer LeViT (Graham et al. 2021), which presents a deep-narrow architecture. As shown in Table 2, our method can further accelerate the deep-narrow transformer such as LeViT. We have observed larger accuracy degradation of our method on LeViT than on DeiT. The reason lies that the deeper layers of LeViT have few tokens and therefore have less redundancy due to the shrinking pyramid structure. With dense input, such as the image resolution of 384×384 , our method accelerates LeViT with less accuracy degradation and more acceleration ratio, which indicates the effectiveness of our method on dense input.

Ablation Analysis

Effectiveness of each module. To evaluate the effectiveness of each sub-method, we add the following improvements step by step in Tab. 3 on transformers of both flat and deep-narrow structures, namely DeiT and LeViT: *i) Naive selection.* Simply drop the placeholder tokens based on the original class attention in each layer; *ii) Structure preservation.* Preserve the placeholder tokens but not fast update them; *iii) Global attention.* Utilize the proposed global class attention

Table 3: Method ablation on DeiT and LeViT.

Strategy	DeiT-T		LeViT 128S	
	Acc. (%)	Throughput (img/s)	Acc. (%)	Throughput (img/s)
baseline	72.2	2536	74.5	8755
+ naive selection	70.8	3824	-	-
+ structure preservation	71.6	3802	72.1	9892
+ global attention	72.0	3730	72.5	9452
+ fast updating	72.0	3610	73.0	9360
+ layer-to-stage	72.0	4027	73.0	10135

Table 4: Different token selection strategies on DeiT-T. We conduct all sub-sampling methods at the seventh layer and conduct token selection strategies from the fifth layer.

Method	Acc. (%)	Throughput (img/s)
average pooling	69.5	3703
max pooling	69.8	3698
convolution	70.2	3688
random selection	66.4	3760
last class attention	69.7	1694
attention column mean	71.2	3596
global class attention	72.0	3730

instead of vanilla class attention for token selection; *iv) Fast updating.* Augment the preserved placeholder tokens with fast updating; *v) Layer-to-stage.* Apply the proposed layer-to-stage training strategy to further accelerate inference.

Results on DeiT indicate that our structure preserving strategy further improves the selection performance due to its capacity of preserving complete information flow. The evolved global class attention enhances the consistency of token selection across layers and achieves better performance. The fast updating strategy has less effect on DeiT than on LeViT. We claim that the performance of DeiT turns to be saturated based on structure preservation and global class attention, while LeViT still has some space for improvement. LeViT exploits spatial pooling for token reduction, which makes unstructured token reduction in each stage more difficult. By using the fast updating strategy, it is possible to collect some extra cues from placeholder tokens for accurate and augmented feature representations. We also evaluate the layer-to-stage strategy. Results indicate that it further accelerates inference while maintaining accuracy.

Different Token Selection Strategy. We compare our global-attention-based token selection strategy with several common token selection strategies and sub-sampling methods in Tab. 4 to evaluate the effectiveness of our method. All token selection strategies are conducted under our structure preserving strategy without layer-to-stage training schedule. The token selection strategies include: randomly selecting the informative tokens (*random selection*); Utilizing the class attention of the last layer for selection in all layers via twice inference (*last class attention*); taking the column mean of the attention matrix as the score of each token as proposed in (Kim et al. 2021) (*attention column mean*).

Results in Tab. 4 indicate that our evolved global class at-




Figure 5: Token selection results on DeiT-T. The left, middle, and right three columns denote the selection results on a well-trained Evo-ViT, the fifth layer at different training epochs, and Evo-ViT with the proposed layer-to-stage strategy, respectively.

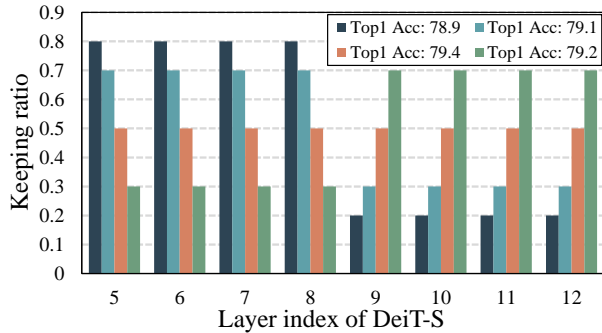


Figure 6: Different architecture of the accelerated DeiT-S via our method. We start our token selection from the fifth layer.

attention outperforms the other selection strategies and common sub-sampling methods on both accuracy and efficiency. We have observed obvious performance degradation with last class attention, although the attention in deeper layers is more focused on objects in Fig. 2. A possible reason is that the networks require some background information to assist classification, while restricting all layers to only focus on objects during the entire training process leads to underfitting on the background features.

Visualization. We visualize the token selection in Fig. 5 to demonstrate performance of our method during both training and testing stages. The visualized models in the left and middle three columns are trained without the layer-to-stage training strategy. The left three columns demonstrate results on different layers of a well-trained DeiT-T model. Results show that our token selection method mainly focuses on objects instead of backgrounds, thereby indicating that our method can effectively discriminate the informative tokens from placeholder tokens. The selection results tend to be consistent across layers, which proves the feasibility of our layer-to-stage training strategy. Another interesting finding is that some missed tokens in the shallow layers are retrieved in the deep layers owing to our structure preserving strategy. Take the baseball images as an example, tokens of the bat are gradually picked up as the layer goes deeper. This phenomenon is more obvious under our layer-to-stage training strategy in the right three columns. We also investigate how the token selection evolves during the training

stage in the middle three columns. Results demonstrate that some informative tokens, such as the fish tail, are determined as placeholder tokens at the early epochs. With more training epochs, our method gradually turns to be stable for discriminative token selection.

Consistent keeping ratio. We set different keeping ratio of tokens in each layer to investigate the best acceleration architecture of Evo-ViT. The keeping ratio determines how many tokens are kept as informative tokens. Previous token pruning works (Rao et al. 2021; Tang et al. 2021) present a gradual shrinking architecture, in which more tokens are recognized as placeholder tokens in deeper layers. They are restricted in this type of architecture due to direct pruning. Our method allows more flexible token selection owing to the structure preserving slow-fast token evolution. As shown in Fig. 6, we maintain the sum of the number of placeholder tokens in all layers and adjust the keeping ratio in each layer. Results demonstrate that the best performance is reached with a consistent keeping ratio across all layers. We explain the reason as follows. In the above visualization, we find that the token selection results tend to be consistent across layers, indicating that the transformer tends to augment informative tokens with computing resource as much as possible. In Fig. 6, at most 50% tokens are passed through the full transformer network when the keeping ratios in all layers are set to 0.5, thereby augmenting the most number of informative tokens with the best computing resource, namely, the full transformer network.

Conclusions

In this work, we investigate the efficiency of vision transformers by developing a self-motivated slow-fast token evolution (Evo-ViT) method. We propose the structure preserving token selection and slow-fast updating strategies to solve the limitation of token pruning on modern structured compressed transformers and scratch training. Extensive experiments on two popular ViT architectures, *i.e.*, DeiT and LeViT, indicate that our Evo-ViT approach is able to accelerate various transformers significantly while maintaining comparable classification performance.

As for future work, an interesting and worthwhile direction is to extend our method to dense prediction tasks, such as object detection and instance segmentation.

Acknowledgements

This work was supported by National Key R&D Program of China under no. 2020AAA0106200, by Shanghai Municipal Science and Technology Key Project under no. 20511100300, by National Natural Science Foundation of China under nos. 62076162, U20B2070, 61832016 and 62036012, and by CASIA-Tencent YouTu joint research project.

References

- Abnar, S.; and Zuidema, W. 2020. Quantifying attention flow in transformers. *arXiv preprint arXiv:2005.00928*.
- Beltagy, I.; Peters, M. E.; and Cohan, A. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*. Springer.
- Chefer, H.; Gur, S.; and Wolf, L. 2021. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Chen, T.; Cheng, Y.; Gan, Z.; Yuan, L.; Zhang, L.; and Wang, Z. 2021. Chasing Sparsity in Vision Transformers: An End-to-End Exploration. *arXiv preprint arXiv:2106.04533*.
- D’Ascoli, S.; Touvron, H.; Leavitt, M. L.; Morcos, A. S.; Biroli, G.; and Sagun, L. 2021. ConViT: Improving Vision Transformers with Soft Convolutional Inductive Biases. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Duke, B.; Ahmed, A.; Wolf, C.; Aarabi, P.; and Taylor, G. W. 2021. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Graham, B.; El-Nouby, A.; Touvron, H.; Stock, P.; Joulin, A.; Jégou, H.; and Douze, M. 2021. LeViT: a Vision Transformer in ConvNet’s Clothing for Faster Inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Han, K.; Wang, Y.; Chen, H.; Chen, X.; Guo, J.; Liu, Z.; Tang, Y.; Xiao, A.; Xu, C.; Xu, Y.; et al. 2020. A survey on visual transformer. *arXiv preprint arXiv:2012.12556*.
- He, R.; Ravula, A.; Kanagal, B.; and Ainslie, J. 2020. Realformer: Transformer likes residual attention. *arXiv preprint arXiv:2012.11747*.
- Heo, B.; Yun, S.; Han, D.; Chun, S.; Choe, J.; and Oh, S. J. 2021. Rethinking Spatial Dimensions of Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Kim, S.; Shen, S.; Thorsley, D.; Gholami, A.; Hassoun, J.; and Keutzer, K. 2021. Learned Token Pruning for Transformers. *arXiv preprint arXiv:2107.00910*.
- Kornblith, S.; Norouzi, M.; Lee, H.; and Hinton, G. 2019. Similarity of neural network representations revisited. In *International Conference on Machine Learning*. PMLR.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Pan, B.; Jiang, Y.; Panda, R.; Wang, Z.; Feris, R.; and Oliva, A. 2021. IA-RED²: Interpretability-Aware Redundancy Reduction for Vision Transformers. *arXiv preprint arXiv:2106.12620*.
- Rao, Y.; Zhao, W.; Liu, B.; Lu, J.; Zhou, J.; and Hsieh, C.-J. 2021. DynamicViT: Efficient Vision Transformers with Dynamic Token Sparsification. *arXiv preprint arXiv:2106.02034*.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Tang, Y.; Han, K.; Wang, Y.; Xu, C.; Guo, J.; Xu, C.; and Tao, D. 2021. Patch Slimming for Efficient Vision Transformers. *arXiv preprint arXiv:2106.02852*.
- Tay, Y.; Bahri, D.; Metzler, D.; Juan, D.-C.; Zhao, Z.; and Zheng, C. 2020. Synthesizer: Rethinking self-attention in transformer models. *arXiv preprint arXiv:2005.00743*.
- Tolstikhin, I.; Houlsby, N.; Kolesnikov, A.; Beyer, L.; Zhai, X.; Unterthiner, T.; Yung, J.; Keysers, D.; Uszkoreit, J.; Lucic, M.; et al. 2021. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021. Training data-efficient image transformers & distillation through attention. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR.
- Wang, S.; Li, B. Z.; Khabsa, M.; Fang, H.; and Ma, H. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021a. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Wang, Y.; Yang, Y.; Bai, J.; Zhang, M.; Bai, J.; Yu, J.; Zhang, C.; Huang, G.; and Tong, Y. 2021b. Evolving attention with residual convolutions. *arXiv preprint arXiv:2102.12895*.
- Xu, W.; Xu, Y.; Chang, T.; and Tu, Z. 2021. Co-Scale Conv-Attentional Image Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

Xu, Y.; Wei, H.; Lin, M.; Deng, Y.; Sheng, K.; Zhang, M.; Tang, F.; Dong, W.; Huang, F.; and Xu, C. 2022. Transformers in computational visual media: A survey. *Computational Visual Media*, 8(1).

Yuan, L.; Chen, Y.; Wang, T.; Yu, W.; Shi, Y.; Jiang, Z.; Tay, F. E.; Feng, J.; and Yan, S. 2021. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

Yue, X.; Sun, S.; Kuang, Z.; Wei, M.; Torr, P. H.; Zhang, W.; and Lin, D. 2021. Vision transformer with progressive sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

Zheng, S.; Lu, J.; Zhao, H.; Zhu, X.; Luo, Z.; Wang, Y.; Fu, Y.; Feng, J.; Xiang, T.; Torr, P. H.; et al. 2021. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Zhou, D.; Shi, Y.; Kang, B.; Yu, W.; Jiang, Z.; Li, Y.; Jin, X.; Hou, Q.; and Feng, J. 2021. Refiner: Refining Self-attention for Vision Transformers. *arXiv preprint arXiv:2106.03714*.

Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.