

Which Transformer to Favor: A Comparative Analysis of Efficiency in Vision Transformers

Tobias Christian Nauen^{1,2}

tobias.christian.nauen@dfki.de

Sebastian Palacio³

sebastian.palacio@de.abb.com

Federico Raue²

federico.raue@dfki.de

Andreas Dengel^{1,2}

andreas.dengel@dfki.de

¹University of Kaiserslautern-Landau, Gottlieb-Daimler-Straße, Kaiserslautern²German Research Center for Artificial Intelligence (DFKI), Trippstadter Straße 122, Kaiserslautern³ABB AG, Kallstadter Str. 1, 68309 Mannheim

Abstract

Self-attention in Transformers comes with a high computational cost because of their quadratic computational complexity, but their effectiveness in addressing problems in language and vision has sparked extensive research aimed at enhancing their efficiency. However, diverse experimental conditions, spanning multiple input domains, prevent a fair comparison based solely on reported results, posing challenges for model selection. To address this gap in comparability, we perform a large-scale benchmark of more than 45 models for image classification, evaluating key efficiency aspects, including accuracy, speed, and memory usage. Our benchmark provides a standardized baseline for efficiency-oriented transformers. We analyze the results based on the Pareto front – the boundary of optimal models. Surprisingly, despite claims of other models being more efficient, ViT remains Pareto optimal across multiple metrics. We observe that hybrid attention-CNN models exhibit remarkable inference memory- and parameter-efficiency. Moreover, our benchmark shows that using a larger model in general is more efficient than using higher resolution images. Thanks to our holistic evaluation, we provide a centralized resource for practitioners and researchers, facilitating informed decisions when selecting or developing efficient transformers. ¹

1. Introduction

The introduction of the Transformer [63] prompted its widespread adoption in both language and vision communities. In particular for image classification, the Vision Trans-

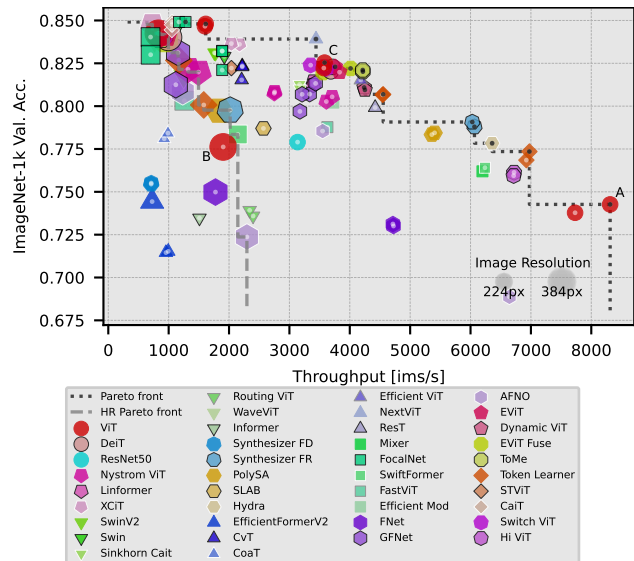


Figure 1. Pareto front (dotted line) showing the throughput-accuracy trade-off for efficient Vision Transformers. Markers (varying in shape and hue) represent different efficiency strategies, with black dots highlighting Pareto optimal models. Marker size indicates fine-tuning resolution. Baselines ViT-Ti@224 (A), ViT-Ti@384 (B), and ViT-S@224 (C) are shown, along with the Pareto front for high-resolution images (dashed line).

former (ViT) [16] has positioned itself as one of the best known applications of the original architecture, outperforming traditional CNN architectures on benchmarks like ImageNet [75, 76]. However, a major challenge in working with Transformer models is dealing with the computational complexity of the self-attention mechanism. This mechanism enables the Transformer to capture global dependencies between pairs of tokens, but it has a computational complexity

¹<https://github.com/tobna/WhatTransformerToFavor>

of $\mathcal{O}(N^2)$ in the input length N , making it impractical for long sequences and high-resolution images.

Efforts have been made to reduce the computational load of Transformer models, particularly for resource-constrained settings, like embedded systems [51, 64]. Researchers have explored numerous strategies, like implementing sparse, local, or kernelized attention mechanisms. However, selecting the most efficient model that meets certain performance standards remains a challenging task, especially since “*efficiency*” can refer to different concepts, like training resources, inference speed, or memory requirements. Moreover, difficulties identifying the most efficient, yet highly performant vision transformers are accentuated by different training and evaluation conditions reported in the literature. It is even more unclear how tradeoffs in efficiency for models proposed for NLP carry over to vision tasks.

To address these problems of fairly comparing efficient ViTs, we design an impartial test bed, training models from scratch on the same data, using the same setup. While many papers at most report inference speed [37, 38, 41, 61], our benchmark enables us to fairly provide empirical values across multiple dimensions of efficiency. We identify common strategies employed to enhance model efficiency and utilize them to correlate changes in efficiency metrics with the respective strategies. Through our benchmark, we provide a centralized resource of comprehensive baselines in the domain of image classification, conducting a thorough review of the current state of research on efficient Transformers. In particular, we provide baselines for Transformers proposed for NLP, which can be trivially adapted to CV.

We analyze model speed and memory requirements both at training and inference time, as well as parameter efficiency. This reveals that not all efficiency claims are realized. Our analysis is based on the Pareto front, the set of models that provide an optimal tradeoff between model performance and one aspect of efficiency. It lets us analyze the complex multidimensional tradeoffs involved in judging efficiency. We plot inference throughput against accuracy in Figure 1 to visualize the Pareto front. Here, a model is Pareto optimal (black dot) if and only if there is no model that is both more accurate and faster at the same time. Through this analysis, we show that ViT remains Pareto optimal across three of the four empirical metrics we track, and also find other efficiency strategies that are Pareto optimal for different metrics. This benchmark, based on more than 200 experiments, provides researchers and practitioners with a valuable resource, aiding them in choosing the most efficient and effective model architecture for their specific use case.

Main Results

- A fair **large-scale benchmark** for efficiency in transformer models run on **45+ models**, closing the comparability gap due to different experimental conditions.

- We find that a well-trained ViT still is Pareto optimal for three out of four metrics, **challenging claims** of other models being more efficient [10, 43, 68].
- Our analysis uncovers **novel findings** like that in most cases, scaling up the model size is more efficient than scaling up the image resolution.

2. Related Work

This section provides an overview of relevant work on transformer surveys and approaches to measure efficiency, while Section 3.3 delves into specific efficient architectures.

Surveys on Efficient Transformers Efficiency has become a critical aspect of transformers, leading to the creation of multiple surveys on efficient transformers providing valuable taxonomies and insights. [55] focuses on efficiency gains primarily in NLP, while [19] and [81] collect general approaches for enhancing model efficiency, including those applicable to transformers. [35, 41] present an extensive list of efficient ViTs, classified based on aspects such as computational complexity, robustness, and transparency, comparing the efficiency in terms of parameters, ImageNet accuracy, and inference speed, using information from the original papers. Surveys on ViT-like models [22, 26, 72] focus on categorizing models according to different vision tasks, while [82] specializes in dense prediction. Specialized surveys investigate the application of transformers in specific domains, like action recognition [60], image restoration [2], medical imaging [23, 31, 40, 50], remote sensing [1], speech recognition [29], language processing [9], time series [66], or multimodal tasks [69].

We build on this by analyzing efficiency gains in terms of broader strategies for efficiency. In contrast to previous work, we close the gap of different training strategies and enable a fair comparison of efficient architectures under the same conditions based on more than new 200 experiments.

Evaluation of Efficiency Efficiency evaluation in deep learning models is another area of investigation. [4, 14] provide an overview of efficiency aspects and metrics, along with measurement methodologies, highlighting the potential pitfalls of relying solely on theoretical metrics. [8] conducts a highly regarded survey on the efficiency of CNNs, and [34] compares the efficiency of their novel architectures with older models using the Pareto front of throughput and accuracy. Finally, [54] introduces a benchmark on synthetic data to quantify transformer model performance on long sequences in NLP, evaluating 10 models.

We extend previous work by evaluating more than 45 efficient Transformers for CV, measuring empirical metrics on the same hardware and using real data. The in-depth analysis of our benchmark’s results reveals new insights.

3. Efficient Transformers for Vision

We present our efficiency assessment methodology and briefly describe ViT, along with various efficient transformer variants. These models from the domains of CV and NLP all claim to improve the efficiency of the baseline transformer. We systematically select efficient transformers based on popularity, diversity, and novelty of their approaches. See the supplementary material for more details on model selection.

3.1. Quantifying Efficiency

The term *efficiency* can take on different meanings; it is therefore crucial to consider multiple dimensions when evaluating a model’s efficiency. In this paper, we analyze model speed and memory requirements at training and inference time and contrast them with accuracy.

Due to the complex multidimensional tradeoffs, efficiency cannot be captured by a single number. Hence, we focus on the Pareto front – the boundary of the landscape of efficient models – to identify models that achieve the best trade-offs between two metrics. The Pareto front represents the set of models that take the best compromises between two metrics. A Pareto optimal model outperforms every other model in at least one of the two metrics. See the Pareto front in Figure 1, where for each model that is not on the Pareto front (white dot) there is a model on the Pareto front (black dot) that is both faster and more accurate.

3.2. Core Elements of ViT

We briefly describe the key elements of ViT, that have been studied to make it more efficient, as well as its key bottleneck: the $\mathcal{O}(N^2)$ computational complexity of self-attention. By identifying these components, we can establish where and how the main strategies, that have been proposed to enhance transformer efficiency, change the architecture. ViT is an adaptation of the original transformer model for image processing tasks. Instead of text, it takes an image as input and converts it into a sequence of non-overlapping patches. Each patch is linearly embedded into a token of size d , with a positional encoding being added. A classification token [CLS] is added to the sequence, which is then fed through a transformer encoder. There, the self-attention mechanism computes the attention weights A between tokens, utilizing the query ($Q \in \mathbb{R}^{N \times d}$) and key ($K \in \mathbb{R}^{N \times d}$) matrices, calculated from the input sequence:

$$A = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_{\text{head}}}} \right) \in \mathbb{R}^{N \times N},$$

where the softmax is calculated row-wise. This matrix encodes the global interactions between every possible pair of tokens. It is also the reason why the attention mechanism has an inherent computational (space & time) complexity of $\mathcal{O}(N^2)$. The output sequence of the attention-mechanism

is a weighted sum of the input, using the value matrix (V):

$$X_{\text{out}} = AV = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_{\text{head}}}} \right) V. \quad (1)$$

After self-attention, the sequence is passed through a feed-forward network with MLP layers. In the end, only the [CLS] token is used for the classification decision.

3.3. Efficiency-Improving Changes

After having gone over the backbone that constitutes a ViT, we discuss the most important modifications, which have been proposed to make it more efficient. In order to analyze the overarching strategies that are taken, we classify the efficient models systematically using a two-step approach. First, we identify the specific components — the token mixing mechanism, the token sequence, or the MLP block — of the baseline ViT architecture that are modified (i.e., “where”), as visualized in Figure 2. On the second level, we classify the different strategies to enhance the efficiency (i.e., “how”). While this taxonomy is not meant to be a comprehensive overview of ViT-like models, it is proposed as a tool for identifying the best and most popular strategies to make vision transformers more efficient. Note that our categorization scheme aligns closely with taxonomies that have been recently proposed [19, 47, 55].

(i) Token Mixing

The first and most popular approach is to change the *token mixing* mechanism, which directly tackles the $\mathcal{O}(N^2)$ computational complexity of self-attention. There are multiple strategies through which this can be accomplished. Some methods approximate the attention mechanism with reduced computation, which can be achieved by matrix decomposition, changing the order of operations, or fixing attention values. Other approaches combine attention with CNNs to perform sub-sampling in the attention mechanism or reduce the number of uses of the attention mechanism. Finally, some methods discard the attention mechanism and instead introduce entirely new token mixing strategies.

Low-rank attention leverages the fact, that $QK^\top \in \mathbb{R}^{N \times N}$ in Equation (1) is a matrix of rank $r \leq d \ll N$. The *Linformer* [65] utilizes this to project the sequence direction of K and V down to dimension $k \ll N$, without reducing the informational content of the attention matrix too much. Similarly, the *Nyströmformer* [68] uses the Nyström method of matrix decomposition to approximate the matrix QK^\top . The approximate attention mechanism’s output is then computed with linear complexity by applying the softmax to each part individually. The approach of *XCiT* [17] utilizes a transposed attention mechanism:

$$Y^\top = \text{softmax} \left(\frac{1}{\sqrt{d}} Q^\top K \right) V^\top.$$




Figure 2. List of efficient Transformers (citation key in brackets) categorized at two levels: 1. Where does the approach change ViT? 2. How does the approach change ViT?

Here, $Q^T K \in \mathbb{R}^{d \times d}$ is used to replace the low-rank matrix QK^T to define the globally informed filter $\text{softmax}\left(\frac{1}{\sqrt{d}}Q^T K\right) \in \mathbb{R}^{d \times d}$ which is applied to each token individually. Since both Q and K are likely to be of rank d , the former most likely has full rank, which enables more efficient information encoding.

Sparse attention is an alternative approach to these dynamic approximations of global interactions. Normal attention tends to focus on only a few input tokens, giving low scores to most others [27]. Sparse attention takes advantage of this by fixing most of the attention weights at zero and only calculating the most important ones. The parts of the attention matrix, which are determined dynamically, can be based on a fixed pattern, like in the widely used *Swin Transformer* [37], and *SwinV2* [36], where attention is performed only inside local sets of image tokens, or in *HaloNet* [62], where each token can only attend to its neighbors. Alternatively, *Routing Transformer* [45] determines sets of tokens that exchange information by grouping them based on their content. In contrast, *Sinkhorn Transformer* [53] fixes local groups and lets tokens only attend to a different group using a dynamic permutation matrix. *Informer* [80] only calculates attention values for the $c \cdot \ln N$ most important tokens. Another way of keeping the attention matrix sparse, is by subsampling the keys and values. For example, *Wave-ViT* [73] uses a discrete wavelet transform for sub-sampling.

Fixed attention is an extreme example of setting attention values beforehand, which then only depend on the tokens positions. This is explored in the *Synthesizer* [52].

Kernel attention is an approach that changes the order of computations in Equation (1). Instead of using the softmax on the product QK^T , a kernel $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is applied to the queries Q and keys K individually:

$$Y = \varphi(Q)\varphi(K)^T V,$$

which can be calculated with linear complexity $\mathcal{O}(N)$ in the number of tokens. Various kernels have been proposed, including the random Gaussian kernel, introduced in *Performer* [13] and expanded on in the *FourierLearner-Transformer* [12] by adding a relative positional encoding in phase space. *Scatterbrain* [11] combines this attention

mechanism with sparse attention. *Poly-SA* [3] uses the identity function $\varphi \equiv \text{id}$, *Linear Transformer* [25] uses an ELU, and *SLAB* [21] uses a ReLU kernel. *Hydra ViT* [6] uses an L2-normalization kernel and employs $H = d$ heads.

Hybrid attention combines convolutions with the attention mechanism. *EfficientFormerV2* [33] uses convolutions initially to focus on local interactions, and then employs the attention mechanism to capture global interactions, while *EfficientViT* [7] utilizes convolutions first and kernel attention, using the ReLU kernel later on. *Next-ViT* [32] alternates $N_{\text{Conv}} \in \mathbb{N}$ convolution blocks with a single attention block. Another approach is to use convolutions inside the attention mechanism to create locally informed queries, keys, and values, as in *CvT* [67] and *ResT* [77], where it is used to subsample keys and values, and in *CodT* [70].

Fourier attention uses the Fast Fourier Transform (FFT) to reduce attention complexity to $\mathcal{O}(N \log N)$. *FNet* [30,48] directly exploits the FFT for $\mathcal{O}(N \log N)$ complexity token mixing, while *GNet* [43] utilizes it for a global convolution and *AFNO* [20] uses an MLP in Fourier space.

Non-attention shuffling refers to techniques of capturing token interactions without using attention. For example *MLP-Mixer* [56] uses a fully connected layer for global interactions, while *FastViT* [61] mixes tokens using depth-wise convolutions. *EfficientMod* [38] modulates a convolutional context with a value matrix and *FocalNet* [71] modulates a hierarchy of contexts extracted by convolutions. *SwiftFormer* [49] calculates a global query to multiply with the keys and also replaces the MLP block by a CNN.

(ii) Token Sequence

The second category, *token sequence*, is more prevalent in efficient transformers used in CV compared to NLP. The idea is to remove redundant information typically contained in images, and in doing so, reducing computational costs without significantly affecting the model’s performance. Methods aim to reduce the token sequence by removing unnecessary image patches, such as the background, by merging tokens to minimize redundant information, or by summarizing the information into a smaller number of abstract tokens representing higher-level information. These models leverage the $\mathcal{O}(N^2)$ complexity of the self-attention mechanism to

attain a large reduction in computational cost, as removing 30% of the tokens reduces the operations needed by around 50%. While some of these methods can directly be applied to trained models, performance improves when implementing them during training [5], which we do.

Token Removal selectively removes tokens while preserving critical information. *Dynamic ViT* [42] uses Gumbel-softmax for token retention probabilities, and *A-ViT* [74] learns halting probabilities to weight token outputs at different depths. In contrast, *EViT* [34] avoids introducing extra parameters by utilizing the previous layer’s attention matrix.

Token merging is another strategy to remove redundant information. While a version of *EViT* [34] merges the unimportant tokens, *ToMe* [5] merges tokens based on their similarity, using a fast bipartite matching algorithm.

Summary tokens condense the token sequence into a few new tokens, unlike the previous methods that modify existing tokens. *CaiT* [59] uses cross-attention on a single token to gather global information on the classification decision in the last layers. *Token Learner* [46] creates a set number of summary tokens by employing dynamic sums over the image tokens, and *STViT* [10] initializes summary tokens with local information by using strided convolutions, then injecting global information into those using cross-attention.

(iii) MLP Block

The final way in which proposed methods change the architecture of transformers is by moving compute to the MLP block, which has linear complexity with respect to the sequence length. Despite room for efficiency gains, we have identified only two methods taking this approach, which we include to make this benchmark more diverse.

More MLPs can be used to move compute into these efficient blocks. *Switch Transformer* [18] introduces multiple parameter sets for different sets of tokens, increasing the number of parameters without significantly introducing extra computations. In contrast, *HiViT* [78] exchanges early attention layers for additional MLP blocks.

Out categorization offers a structured framework for understanding approaches of improving ViT efficiency. This enables us to compare broad trends across strategies.

4. Experimental Design

We conduct an extensive series of over 200 experiments on the more than 45 models described in Section 3.3 to evaluate which models are ideal under a given set of constraints. We measure efficiency w.r.t. the full models rather than individual attention modules, since methods that reduce the token sequence (Section 3.3 (ii)), and others, like XCiT, require additional modifications to the architecture, which we include. Note, however, that many models only introduce changes in the attention mechanism, effectively turning our benchmark into a module-level comparison for these.

4.1. Baselines

To find the most efficient models, we need to quantify their efficiency gains by comparing them to the original ViT [16] and its follow-up versions [57,58]. We also include ResNet50 [24] in our evaluation as a representative baseline for CNN architectures and a point of reference across papers. All evaluations are based on the ImageNet-1k dataset for classification [15] as it is one of the most commonly used datasets in this domain. Results on other datasets (Stanford Cars [28], Oxford Flowers 102 [39], MIT Places365 [79]) can be found in the supplementary material.

4.2. Training Pipeline

We compare models on even grounds by training with a standardized pipeline. To reduce bias, our pipeline is relatively simple and only consists of elements commonly used in CV. In particular, we refrain from using knowledge distillation to prevent introducing bias from the choice of teacher model. Any orthogonal techniques, like quantization, sample selection, and others, are not included as they can be applied to every model and would manifest as a systematic offset in the results. To avoid confounding effects from limited training data, we pre-train all models on ImageNet-21k [44]. While data scarcity impacts performance, we focus on evaluating each architecture’s maximum potential.

Models are trained for a total of 140 epochs using the pipeline introduced in [58]. This is an updated version of DeiT [57], which has been adopted throughout the literature for training efficient ViTs (see Table 1). As this pipeline only employs standard CV practices, we consider it a fair point of comparison for evaluating all models.

Pretraining is conducted on ImageNet-21k for 90 epochs at resolutions of 224 and 192 pixels, followed by fine-tuning on ImageNet-1k at 224 and 384 pixels for 50 epochs. In cases where model training is unstable, we default back to the hyperparameters reported in the corresponding publications. All training is conducted using 4 or 8 NVIDIA A100 GPUs. While most of the models work well with our training pipeline, some did not converge. See the supplementary material for more details.

4.3. Efficiency Metrics

As we saw, efficiency has multiple dimensions, and therefore we need to capture a variety of metrics. For comparison with other papers and to evaluate their use as proxy metrics for speed and memory requirements, we capture the theoretical metrics of number of parameters and FLOPS. These are used by the community to provide estimates of the model’s representational capacity and computational requirements [41,43,55]. However, theoretical metrics do not always correlate with real-world performance, especially when comparing different architectures [4,14]. Instead, empirical metrics, obtained by running models on hardware, offer

Table 1. ImageNet-1k accuracy of the original papers and the new training pipeline. ✓ marks original pipelines based on DeiT [57]. Models are trained on 224px images, unless marked with ↑ (384px) or ↓ (112px). Results with ^D use knowledge distillation and ^E uses exponential moving average.

Model	Original		Ours
	DeiT	Accuracy	Accuracy
ViT-S (DeiT)	✓	79.8	82.54
ViT-S (DeiT III)		82.6	82.54
XCiT-S	✓	82.0	83.65
Swin-S	✓	83.0	84.87
SwinV2-Ti		81.7	83.09
Wave-ViT-S		82.7	83.61
Poly-SA-ViT-S		71.48	78.34
SLAB-S	✓	80.0	78.70
EfficientFormer-V2-S0		75.7^D	71.53
CvT-13		83.3[↑]	82.35
CoaT-Ti	✓	78.37	78.42
EfficientViT-B2		82.7[↑]	81.52
NextViT-S		82.5	83.92
ResT-S	✓	79.6	79.92
FocalNet-S		83.4	84.91
SwiftFormer-S		78.5^D	76.41
FastViT-S12	✓	79.8[↑]	78.77
EfficientMod-S	✓	81.0	80.21
GFNet-S		80.0	81.33
EViT	✓	79.4	82.29
DynamicViT-S		83.0^D	81.09
EViT Fuse	✓	79.5	81.96
ToMe-ViT-S	✓	79.42	82.11
TokenLearner-ViT-8		77.87 [↓]	80.66
STViT-Swin-Ti	✓	80.8	82.22
CaiT-S24	✓	82.7	84.91

more informative evaluations, which is why we track GPU time for training, inference throughput at optimal batch sizes, and VRAM requirements. While these empirical metrics can be sensitive to hardware and software configurations, we ensure consistent evaluations by employing the same setup for all models. We evaluate the variance of our metrics in the supplementary material.

5. Results

5.1. Improved Training Pipeline

To validate the fairness of our training pipeline, in Table 1 we compare ImageNet-1k accuracy reported in the original papers (whenever reported) with the one we obtained to see if there are any obvious outliers. Note, that a substantial fraction (13 out of 26) of papers explicitly base their training pipelines on DeiT [57], making them a good fit for training with the updated version in DeiT III [58]. We observe that almost all models trained with this pipeline achieve a higher accuracy; 0.85% on average, with the largest gain being +6.86% for Poly-SA. The models reporting higher performance than what we obtain use knowledge distillation, which we avoid to not induce bias, or a higher resolution, which we find is inefficient.

Most categories contain models that achieve peak accu-

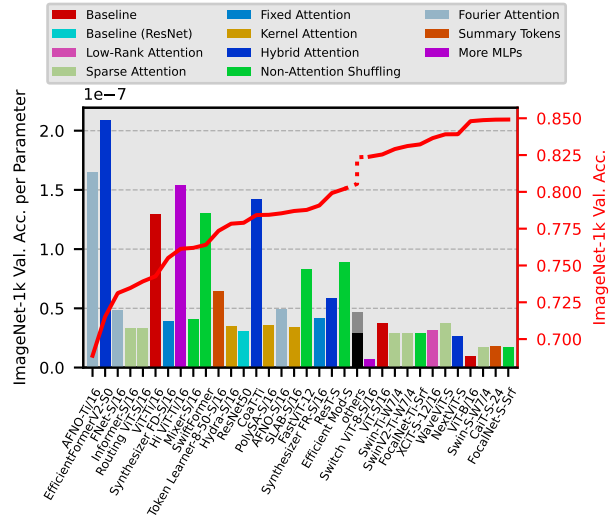


Figure 3. Accuracy (red line, right y-axis) and accuracy per parameter (bars, left y-axis) of models ordered by accuracy at a resolution of 224px. 14 of 47 models of intermediate size are grouped into *others*. See the supplementary material for the full plot.

rary of approximately 85% (see the supplementary material). However, two notable exceptions are observed: the *Kernel Attention* class, which might pose challenges in optimizing models, and the *Fixed Attention* class, where the use of a constant attention matrix inherently results in a lower accuracy. Overall, our selection of the training pipeline as a reference for a fair comparison is supported by the improvements in results compared to the original papers, regardless of architectural differences.

5.2. Number of Parameters

The number of parameters is used in the literature as a proxy metric for tracking the model complexity and the overall computational cost of using a model. When analyzing the parameter efficiency of different models (Figure 3), it is evident that for most smaller models, the accuracy per parameter remains relatively constant at a baseline level of about $4 \times 10^{-8} \frac{\%}{\text{param}}$. This value is approximately halved for larger models, indicating diminishing returns on scaling the model size. For ViT, there is a noticeable drop in accuracy per parameter across the model sizes. While ViT-Ti outperforms models of comparable accuracy, and ViT-S performs on par with similar models, ViT-B slightly underperforms when compared to other, larger models. Significant outliers are the smaller *Hybrid Attention* models EfficientFormerV2-S0 and CoaT-Ti, which exhibit the highest accuracy per parameter, as well as those *Non-attention Shuffling* models which incorporate convolutions, like SwiftFormer and FastViT. In the *MLP* category, we find HiViT to have significantly higher and Switch ViT with significantly lower accuracy per parameter than surrounding models. This suggests that the




Figure 4. Pareto front of finetuning time and accuracy for models which need less than 50 hours for finetuning. We include the full plot in the supplementary material.

combination of attention and convolutions allows for the development of highly parameter-efficient models.

5.3. Speed

Whether driven by strict requirements for real-time processing or the desire to obtain model outputs within reasonable timeframes, inference speed directly impacts the usability of deployed models. The models we evaluate often claim a superior throughput vs. accuracy trade-off compared to ViT. However, our evaluation in Figure 1 reveals that ViT remains Pareto optimal at all model sizes. Only few models, namely Synthesizer-FR, NextViT, and some *Sequence Reduction* models, show improvements in the Pareto front when compared to a ViT of comparable size.

We also analyze the Pareto optimality with models finetuned on Stanford Cars [28], Oxford Flowers 102 [39], and MIT Places365 [79] in the supplementary material, observing a high Spearman correlation between the accuracy of different datasets. In particular, the Pareto optimal models of throughput and accuracy are largely the same, which is a strong indicator for the generalizability of our results. Even when assessing throughput on a CPU instead of a GPU, we observe only a high Spearman correlation with only few outliers. Again, we find the Pareto front to be robust to the choice of hardware. See the supplementary material for more information about these results.

In the Pareto front of fine-tuning time in Figure 4, we see many similarities to the one of inference time. Here, some of the *token sequence* based models are highly efficient. In particular, TokenLearner emerges as a standout performer

with the fastest fine-tuning speed while achieving a competitive accuracy of 77.35%. For larger models, we find ViT-B and FocalNet (-Ti and -S) to be Pareto optimal.

Generally, ViT is still a solid choice for speed, with some *Token Sequence* models trading off accuracy for speed.

5.4. Memory

Figure 5 (right) reveals that the tradeoffs when optimizing for VRAM usage during inference are very different from the ones for speed or training VRAM (left side), making it the only metric where ViT is *not* Pareto optimal. At inference time, the *Hybrid Attention* models CoaT, CvT, ResT, and NextViT stand out from the rest. Additionally, CaiT, and EViT@384px for the larger sizes and EfficientMod for the smaller ones are Pareto optimal. This observation suggests that, similar to the number of parameters, *Hybrid Attention* models excel in low memory environments.

In contrast, the aspect of training memory (Figure 5; left) exhibits a similar pattern as observed in throughput. Here, CoaT, CvT, and ResT need more memory than other models of comparable accuracy and again ViT is Pareto optimal, along with some *Sequence Reduction* models and Synthesizer-FR. NextViT is the only model that is Pareto optimal for both training and inference memory. When training at a resolution of 384px most models need more than 225GB of VRAM and therefore are not plotted.

Overall, inference memory is the most different from the other Pareto fronts with *Hybrid Attention* models excelling.

5.5. Scaling Behaviors

Our observations reveal that fine-tuning at a higher resolution is inefficient. While it may result in improved accuracy, it entails a significant increase in computational cost, leading to a substantial reduction in throughput. In turn, scaling up the model ends up being more efficient. This is shown in Figure 1, where the best models using high resolution images are two to four times slower than a low resolution model at the same accuracy (compare *Pareto front* and *HR Pareto front*). Consequently, no model fine-tuned at the higher resolution is Pareto optimal. For example, when scaling from tiny to small (A to C in Figure 1), we gain a lot more accuracy for the lost speed than when scaling up the image size (A to B). For training memory, this trend is even more pronounced (see *HR Pareto front* in Figure 5), with most of the high resolution models requiring more than 225GB of VRAM for training and thus being off the chart. Overall, many papers focus on reducing the $\mathcal{O}(N^2)$ complexity of self-attention to deal with long sequences, however sequences in image classification are typically not long enough to show the theoretical gains. On the contrary, it is more efficient to use shorter sequences instead. We believe that efficient mechanisms focusing on these shorter sequences are a fruitful area for future research.

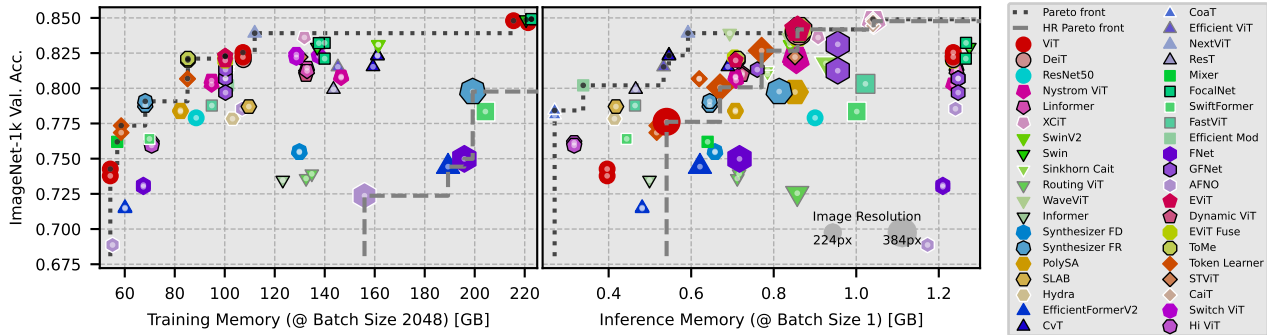


Figure 5. Pareto front (dotted line) of training memory at our default batch size of 2048 (left) and inference memory at the minimum batch size of 1 (right) and accuracy for models with less than 225GB of VRAM for training and 1.25GB for inference.

Table 2. Correlation between the number of floating point operations, number of parameters (θ), fine-tuning time, training memory, inference time, and inference memory.

$\text{corr}(x, y)$	θ	Training		Inference	
		Time	Mem	Time	Mem
FLOPS	0.30	0.72	0.85	0.48	0.42
θ		0.05	0.18	0.02	0.40
Training Time			0.89	0.81	0.17
Training Mem				0.71	0.48
Inference Time					0.13

5.6. Correlation of Metrics

We run a correlation analysis between the efficiency metrics in Table 2. The highest correlation of 0.89 is between fine-tuning time and training memory. This suggests a common underlying factor or bottleneck, possibly related to the necessity of memory reads during training. Understanding this relationship can provide valuable insights into the factors influencing training efficiency. Intriguingly, the highest correlation coefficient involving a theoretical metric is 0.85 between FLOPS and training memory, suggesting the possibility for a rough estimation. For other metrics, we reaffirm the limited reliability of estimating computational costs solely based on theoretical metrics [4, 14]. Consequently, assessing model efficiency in practice requires the empirical measurement of throughput and memory requirements.

6. Discussion & Conclusion

In this paper, we provide a thorough benchmark of efficiency oriented transformers for vision, facilitating a fair, competitive comparison of contributions from NLP and CV and revealing some unexpected trends. Our analysis shows ViT’s continued Pareto optimality across multiple metrics, reaffirming its effectiveness as a baseline, despite other models claiming to be more efficient. Results also prove that scaling the model size is, in most cases, more efficient than using higher resolution images. This goes against the trend




Figure 6. Suggestions on *Which Transformer to Favor* in image classification. First, choose the main constraint (1.), speed, or memory, and the phase you want to optimize this constraint at (2.): Inference or training. Based on that, we provide recommendations on which transformer to favor (3.) from the taxonomy classes of **Baseline**, **Sequence Reduction**, **Hybrid Attention**, **Non-Attention Shuffling** and **Sparse Attention**. See the supplementary material for a full list of Pareto optimal models for each metric.

of efficient models being evaluated using high resolution images [17, 37, 59, 61]. Comparing different efficiency metrics, we reaffirm the limitations of estimating computational costs solely based on theoretical metrics.

When addressing the question of *which transformer to favor*, our benchmark offers actionable insights in the form of models and strategies to use (see Figure 6). ViT remains the preferred choice overall. However, *Token Sequence* methods can become viable alternatives when speed and training efficiency are of importance. For scenarios with significant inference memory constraints, considering *Hybrid CNN*-attention models can prove advantageous.

Looking ahead, we posit that the establishment of this comparable, fair benchmark and analysis will drive substantial advancements in efficiency-oriented transformers for CV, benefitting researchers and practitioners making architectural decisions for resource-constrained environments.

Acknowledgements

This work was funded by the Carl-Zeiss Foundation under the Sustainable Embedded AI project (P2021-02-009), by the EU project SustainML (Horizon Europe grant agreement No 101070408), and by the Albatross project (01IW24002).

References

- [1] Abdulaziz Amer Aleissae, Amandeep Kumar, Rao Muhammad Anwer, Salman Khan, Hisham Cholakkal, Gui-Song Xia, and Fahad Shahbaz Khan. Transformers in remote sensing: A survey. *Remote Sensing*, 15(7):1860, 2023. 2
- [2] Anas M. Ali, Bilel Benjdira, Anis Koubaa, Walid El-Shafai, Zahid Khan, and Wadii Boulila. Vision transformers in image restoration: A survey. *Sensors*, 23(5), 2023. 2
- [3] Francesca Babiloni, Ioannis Marras, Jiankang Deng, Filippos Kokkinos, Matteo Maggioni, Grigorios Chrysos, Philip Torr, and Stefanos Zafeiriou. Linear complexity self-attention with 3rd order polynomials. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–12, 2023. 4
- [4] Brian R Bartoldson, Bhavya Kailkhura, and Davis Blalock. Compute-efficient deep learning: Algorithmic trends and opportunities. *Journal of Machine Learning Research*, 24(122):1–77, 2023. 2, 5, 8
- [5] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. In *International Conference on Learning Representations*, 2023. 5
- [6] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, and Judy Hoffman. Hydra attention: Efficient attention with many heads. In Leonid Karlinsky, Tomer Michaeli, and Ko Nishino, editors, *Computer Vision – ECCV 2022 Workshops*, pages 35–49, Cham, 2022. Springer Nature Switzerland. 4
- [7] Han Cai, Junyan Li, Muyan Hu, Chuang Gan, and Song Han. Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17302–17313, October 2023. 4
- [8] Alfredo Canziani, Adam Paszke, and Eugenio Culurciello. An analysis of deep neural network models for practical applications, 2016. 2
- [9] Silvia Casola, Ivano Lauriola, and Alberto Lavelli. Pre-trained transformers: an empirical comparison. *Machine Learning with Applications*, 9:100334, 2022. 2
- [10] Shuning Chang, Pichao Wang, Ming Lin, Fan Wang, David Junhao Zhang, Rong Jin, and Mike Zheng Shou. Making vision transformers efficient from a token sparsification view. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6195–6205, June 2023. 2, 5
- [11] Beidi Chen, Tri Dao, Eric Winsor, Zhao Song, Atri Rudra, and Christopher Ré. Scatterbrain: Unifying sparse and low-rank attention approximation. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 4
- [12] Krzysztof Marcin Choromanski, Shanda Li, Valerii Likhoshervstov, Kumar Avinava Dubey, Shengjie Luo, Di He, Yiming Yang, Tamas Sarlos, Thomas Weingarten, and Adrian Weller. Learning a fourier transform for linear relative positional encodings in transformers. In *AISTATS 2024*, 2024. 4
- [13] Krzysztof Marcin Choromanski, Valerii Likhoshervstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukas Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021. 4
- [14] Mostafa Dehghani, Yi Tay, Anurag Arnab, Lucas Beyer, and Ashish Vaswani. The efficiency misnomer. In *International Conference on Learning Representations*, 2022. 2, 5, 8
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009. 5
- [16] 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 *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 1, 5
- [17] Alaaeldin El-Nouby, Hugo Touvron, Mathilde Caron, Piotr Bojanowski, Matthijs Douze, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Jakob Verbeek, and Hervé Jegou. Xcit: Cross-covariance image transformers. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 3, 8
- [18] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022. 5
- [19] Quentin Fournier, Gaétan Marceau Caron, and Daniel Aloise. A practical survey on faster and lighter transformers. *ACM Comput. Surv.*, 3 2023. 2, 3
- [20] John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro. Efficient token mixing for transformers via adaptive fourier neural operators. In *International Conference on Learning Representations*, 2022. 4
- [21] Jialong Guo, Xinghao Chen, Yehui Tang, and Yunhe Wang. Slab: Efficient transformers with simplified linear attention and progressive re-parameterized batch normalization. In *International Conference on Machine Learning*, 2024. 4
- [22] K. Han, Y. Wang, H. Chen, X. Chen, J. Guo, Z. Liu, Y. Tang, A. Xiao, C. Xu, Y. Xu, Z. Yang, Y. Zhang, and D. Tao. A survey on vision transformer. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 45(01):87–110, 1 2023. 2
- [23] Kelei He, Chen Gan, Zhuoyuan Li, Islem Rekik, Zihao Yin, Wen Ji, Yang Gao, Qian Wang, Junfeng Zhang, and Dinggang Shen. Transformers in medical image analysis. *Intelligent Medicine*, 3(1):59–78, 2023. 2
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [25] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. *International Conference on Machine Learning*, pages 5156–5165, 2020. 4

- [26] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Computing Surveys*, 54(10s):1–41, 2022. [2](#)
- [27] Kyungmin Kim, Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Zhicheng Yan, Peter Vajda, and Seon Kim. Rethinking the self-attention in vision transformers. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 2021. [4](#)
- [28] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. [5](#), [7](#)
- [29] Siddique Latif, Aun Zaidi, Heriberto Cuayahuitl, Fahad Shamshad, Moazzam Shoukat, and Junaid Qadir. Transformers in speech processing: A survey, 2023. [2](#)
- [30] James P Lee-Thorp, Joshua Ainslie, Ilya Eckstein, and Santiago Ontañón. Fnet: Mixing tokens with fourier transforms. In *NAACL*, 2022. [4](#)
- [31] Jun Li, Junyu Chen, Yucheng Tang, Ce Wang, Bennett A. Landman, and S. Kevin Zhou. Transforming medical imaging with transformers? a comparative review of key properties, current progresses, and future perspectives. *Medical Image Analysis*, 85:102762, 2023. [2](#)
- [32] Jiashi Li, Xin Xia, Wei Li, Huixia Li, Xing Wang, Xuefeng Xiao, Rui Wang, Min Zheng, and Xin Pan. Next-vit: Next generation vision transformer for efficient deployment in realistic industrial scenarios. 2022. [4](#)
- [33] Yanyu Li, Ju Hu, Yang Wen, Georgios Evangelidis, Kamyar Salahi, Yanzhi Wang, Sergey Tulyakov, and Jian Ren. Rethinking vision transformers for mobilenet size and speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16889–16900, October 2023. [4](#)
- [34] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *International Conference on Learning Representations*, 2022. [2](#), [5](#)
- [35] Yang Liu, Yao Zhang, Yixin Wang, Feng Hou, Jin Yuan, Jiang Tian, Yang Zhang, Zhongchao Shi, Jianping Fan, and Zhiqiang He. A survey of visual transformers. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21, 2023. [2](#)
- [36] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer v2: Scaling up capacity and resolution. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11999–12009. arXiv, 2022. [4](#)
- [37] 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 *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, Los Alamitos, CA, USA, 10 2021. IEEE Computer Society. [2](#), [4](#), [8](#)
- [38] Xu Ma, Xiyang Dai, Jianwei Yang, Bin Xiao, Yinpeng Chen, Yun Fu, and Lu Yuan. Efficient modulation for vision networks. In *The Twelfth International Conference on Learning Representations*, 2024. [2](#), [4](#)
- [39] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008. [5](#), [7](#)
- [40] Arshi Parvaiz, Muhammad Anwaar Khalid, Rukhsana Zafar, Huma Ameer, Muhammad Ali, and Muhammad Moazzam Fraz. Vision transformers in medical computer vision—a contemplative retrospection. *Engineering Applications of Artificial Intelligence*, 122:106126, 2023. [2](#)
- [41] Badri N. Patro and Vijay Srinivas Agneeswaran. Efficiency 360: Efficient vision transformers, 2023. [2](#), [5](#)
- [42] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 13937–13949. Curran Associates, Inc., 2021. [5](#)
- [43] Yongming Rao, Wenliang Zhao, Zheng Zhu, Jiwen Lu, and Jie Zhou. Global filter networks for image classification. *Advances in Neural Information Processing Systems*, 34, 2021. [2](#), [4](#), [5](#)
- [44] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. [5](#)
- [45] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 9:53–68, 2021. [4](#)
- [46] Michael S. Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. Tokenlearner: What can 8 learned tokens do for images and videos? *NeurIPS 2021*, 2021. [5](#)
- [47] Javier Selva, Anders S. Johansen, Sergio Escalera, Kamal Nasrollahi, Thomas B. Moeslund, and Albert Clapés. Video transformers: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–20, 2023. [3](#)
- [48] Nurullah Sevim, Ege Ozan Özyedek, Furkan Şahinuç, and Aykut Koç. Fast-fnet: Accelerating transformer encoder models via efficient fourier layers, 2022. [4](#)
- [49] Abdelrahman Shaker, Muhammad Maaz, Hanoona Rasheed, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Swiftformer: Efficient additive attention for transformer-based real-time mobile vision applications. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [4](#)
- [50] Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu. Transformers in medical imaging: A survey. *Medical Image Analysis*, page 102802, 2023. [2](#)
- [51] Hamid Tabani, Ajay Balasubramaniam, Shabbir Marzban, Elahe Arani, and Bahram Zonooz. Improving the efficiency of transformers for resource-constrained devices. In *2021 24th Euromicro Conference on Digital System Design (DSD)*, pages 449–456, 2021. [2](#)

- [52] Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, Che Zheng, Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. Synthesizer: Rethinking self-attention in transformer models. In *International conference on machine learning*, pages 10183–10192. PMLR, 2021. 4
- [53] Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. Sparse sinkhorn attention. In *International Conference on Machine Learning*, pages 9438–9447. PMLR, arXiv, 2020. 4
- [54] Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021. 2
- [55] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *ACM Comput. Surv.*, 4 2022. 2, 3, 5
- [56] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Peter Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. MLP-mixer: An all-MLP architecture for vision. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, pages 24261–24272, 2021. 4
- [57] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10347–10357. PMLR, 7 2021. 5, 6
- [58] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 516–533, Cham, 2022. Springer Nature Switzerland. 5, 6
- [59] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Herve Jegou. Going deeper with image transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, 2021. 5, 8
- [60] Anwaar Ulhaq, Naveed Akhtar, Ganna Pogrebna, and Ajmal Mian. Vision transformers for action recognition: A survey, 2022. 2
- [61] Pavan Kumar Anasosalu Vasu, James Gabriel, Jeff Zhu, Oncel Tuzel, and Anurag Ranjan. Fastvit: A fast hybrid vision transformer using structural reparameterization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5785–5795, October 2023. 2, 4, 8
- [62] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12889–12899, Los Alamitos, CA, USA, 6 2021. IEEE Computer Society. 4
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 1
- [64] Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. HAT: Hardware-aware transformers for efficient natural language processing. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetraault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7675–7688, Online, July 2020. Association for Computational Linguistics. 2
- [65] Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity, June 2020. 3
- [66] Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in time series: A survey. In *the 32nd International Joint Conference on Artificial Intelligence (IJCAI 2023)*, 2022. 2
- [67] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22–31, 2021. 4
- [68] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021. 2, 3
- [69] Peng Xu, Xiayan Zhu, and David A. Clifton. Multimodal learning with transformers: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10):12113–12132, 2023. 2
- [70] Weijian Xu, Yifan Xu, Tyler Chang, and Zhuowen Tu. Co-scale conv-attentional image transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9981–9990, 2021. 4
- [71] Jianwei Yang, Chunyuan Li, Xiyang Dai, Lu Yuan, and Jianfeng Gao. Focal modulation networks. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 4
- [72] Yuting Yang, Licheng Jiao, Xu Liu, Fang Liu, Shuyuan Yang, Zhixi Feng, and Xu Tang. Transformers meet visual learning understanding: A comprehensive review, 2022. 2
- [73] Ting Yao, Yingwei Pan, Yehao Li, Chong-Wah Ngo, and Tao Mei. Wave-vit: Unifying wavelet and transformers for visual representation learning. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 328–345, Cham, 2022. Springer Nature Switzerland. 4
- [74] Hongxu Yin, Arash Vahdat, Jose M. Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. A-vit: Adaptive tokens for efficient vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10809–10818, 6 2022. 5
- [75] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive

- captioners are image-text foundation models. *Transactions on Machine Learning Research*, 2022. [1](#)
- [76] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12104–12113, 2022. [1](#)
- [77] Qinglong Zhang and Yu-Bin Yang. Rest: An efficient transformer for visual recognition. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*. arXiv, 2021. [4](#)
- [78] Xiaosong Zhang, Yunjie Tian, Lingxi Xie, Wei Huang, Qi Dai, Qixiang Ye, and Qi Tian. Hivit: A simpler and more efficient design of hierarchical vision transformer. In *The Eleventh International Conference on Learning Representations*, 2023. [5](#)
- [79] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. [5](#), [7](#)
- [80] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 11106–11115. AAAI Press, 2021. [4](#)
- [81] Bohan Zhuang, Jing Liu, Zizheng Pan, Haoyu He, Yuetian Weng, and Chunhua Shen. A survey on efficient training of transformers. In Edith Elkind, editor, *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 6823–6831. International Joint Conferences on Artificial Intelligence Organization, 8 2023. Survey Track. [2](#)
- [82] Shuangquan Zuo, Yun Xiao, Xiaojun Chang, and Xuanhong Wang. Vision transformers for dense prediction: A survey. *Knowledge-Based Systems*, 253:109552, 2022. [2](#)