

EfficientTrain: Exploring Generalized Curriculum Learning for Training Visual Backbones

Yulin Wang^{1*} Yang Yue^{1*} Rui Lu¹ Tianjiao Liu² Zhao Zhong²
Shiji Song¹ Gao Huang^{1,3}✉

¹Department of Automation, BNRist, Tsinghua University ²Huawei Technologies Ltd. ³BAAI
{wang-y119, yueyang22}@mails.tsinghua.edu.cn, gaohuang@tsinghua.edu.cn

Abstract

The superior performance of modern deep networks usually comes with a costly training procedure. This paper presents a new curriculum learning approach for the efficient training of visual backbones (e.g., vision Transformers). Our work is inspired by the inherent learning dynamics of deep networks: we experimentally show that at an earlier training stage, the model mainly learns to recognize some ‘easier-to-learn’ discriminative patterns within each example, e.g., the lower-frequency components of images and the original information before data augmentation. Driven by this phenomenon, we propose a curriculum where the model always leverages all the training data at each epoch, while the curriculum starts with only exposing the ‘easier-to-learn’ patterns of each example, and introduces gradually more difficult patterns. To implement this idea, we 1) introduce a cropping operation in the Fourier spectrum of the inputs, which enables the model to learn from only the lower-frequency components efficiently, 2) demonstrate that exposing the features of original images amounts to adopting weaker data augmentation, and 3) integrate 1) and 2) and design a curriculum learning schedule with a greedy-search algorithm. The resulting approach, EfficientTrain, is simple, general, yet surprisingly effective. As an off-the-shelf method, it reduces the wall-time training cost of a wide variety of popular models (e.g., ResNet, ConvNeXt, DeiT, PVT, Swin, and CSWin) by $> 1.5\times$ on ImageNet-1K/22K without sacrificing accuracy. It is also effective for self-supervised learning (e.g., MAE). Code is available at <https://github.com/LeapLabTHU/EfficientTrain>.

1. Introduction

The success of modern visual backbones is largely fueled by the interest in exploring big models on large-scale bench-

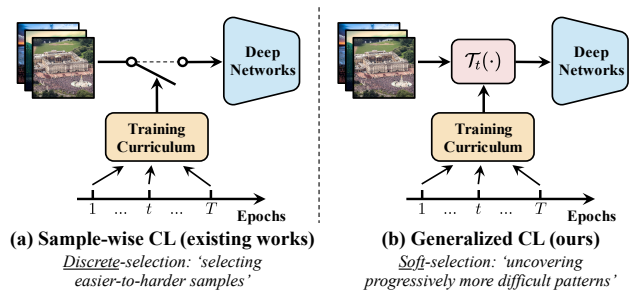


Figure 1: (a) **Sample-wise curriculum learning (CL)**: making a *discrete* decision on whether each example should be leveraged to train the model. (b) **Generalized CL**: we consider a *continuous* function $\mathcal{T}_t(\cdot)$, which only exposes the ‘easier-to-learn’ patterns within each example at the beginning of training (e.g., lower-frequency components; see: Section 4), while gradually introducing relatively *more difficult* patterns as learning progresses.

mark datasets [1, 2, 3, 4]. In particular, the recent introduction of vision Transformers (ViTs) scales up the number of model parameters to more than 1.8 billion, with the training data expanding to 3 billion samples [3, 5]. Although state-of-the-art accuracy has been achieved, this huge-model and high-data regime results in a time-consuming and expensive training process. For example, it takes 2,500 TPUv3-core-days to train ViT-H/14 on JFT-300M [3], which may be unaffordable for practitioners in both academia and industry. Additionally, the power consumption leads to significant carbon emissions [6, 7]. Due to both economic and environmental concerns, there has been a growing demand for reducing the training cost of modern deep networks.

In this paper, we contribute to this issue by revisiting the idea of curriculum learning [8], which reveals that a model can be trained efficiently by starting with the easier aspects of a given task or certain easier subtasks, and increasing the difficulty level gradually. Most existing works implement this idea by introducing easier-to-harder examples progressively during training [9, 10, 11, 12, 13, 14, 15, 16]. However, obtaining a light-weighted and generalizable difficulty measurer is typically non-trivial [9, 10]. In general, these methods have not exhibited the capacity to be a universal

*Equal contribution.

✉Corresponding author.

efficient training technique for modern visual backbones.

In contrast to prior works, this paper seeks a simple yet broadly applicable efficient learning approach with the potential for widespread implementation. To attain this goal, we consider a generalization of curriculum learning. In specific, we extend the notion of training curricula beyond only differentiating between ‘easier’ and ‘harder’ examples, and adopt a more flexible hypothesis, which indicates that the discriminative features of each training sample comprise both ‘easier-to-learn’ and ‘harder-to-learn’ patterns. Instead of making a *discrete* decision on whether each example should appear in the training set, we argue that it would be more proper to establish a *continuous* function that adaptively extracts the simpler and more learnable discriminative patterns within every example. In other words, a curriculum may always leverage all examples at any stage of learning, but it should eliminate the relatively more difficult or complex patterns within inputs at earlier learning stages. An illustration of our idea is shown in Figure 1.

Driven by our hypothesis, a straightforward yet surprisingly effective algorithm is derived. We first demonstrate that the ‘easier-to-learn’ patterns incorporate the lower-frequency components of images. We further show that a lossless extraction of these components can be achieved by introducing a cropping operation in the frequency domain. This operation not only retains exactly all the lower-frequency information, but also yields a smaller input size for the model to be trained. By triggering this operation at earlier training stages, the overall computational/time cost for training can be considerably reduced while the final performance of the model will not be sacrificed. Moreover, we show that the original information before heavy data augmentation is more learnable, and hence starting the training with weaker augmentation techniques is beneficial. Finally, these theoretical and experimental insights are integrated into a unified ‘*EfficientTrain*’ learning curriculum by leveraging a greedy-search algorithm.

One of the most appealing advantages of *EfficientTrain* may be its simplicity and generalizability. Our method can be conveniently applied to most deep networks *without any modification or hyper-parameter tuning*, but significantly improves their training efficiency. Empirically, for the supervised learning on ImageNet-1K/22K [17], *EfficientTrain* reduces the wall-time training cost of a wide variety of popular visual backbones (e.g., ConvNeXt [18], DeiT [19], PVT [20], Swin [4], and CSWin [21]) by *more than 1.5×*, while achieving competitive or better performance compared with the baselines. Importantly, our method is also effective for self-supervised learning (e.g., MAE [22]).

2. Related Work

Curriculum learning is a training paradigm inspired by the organized learning order of examples in human cur-

ricula [23, 24, 8]. This idea has been widely explored in the context of training deep networks from easier data to harder data [25, 11, 16, 12, 26, 27, 9, 10]. Typically, a pre-defined [8, 28, 29, 30] or automatically-learned [25, 31, 32, 11, 33, 16, 14, 34, 35, 12, 36] difficulty measurer is deployed to differentiate between easier and harder samples, while a scheduler [8, 11, 26, 9] is defined to determine when and how to introduce harder training data. Our method is also based on the ‘starting small’ spirit [23], but we always leverage all the training data simultaneously. Our work is also related to curriculum by smoothing [37] and curriculum dropout [38], which do not perform example selection as well. However, our method is orthogonal to them since we reduce the training cost by modifying the model inputs, while they regularize deep features during training (e.g., via anti-aliasing smoothing or feature dropout).

Progressive or modularized training. Deep networks can be trained efficiently by increasing the model size during training, e.g., a growing number of layers [39, 40, 41], a growing width [42], or a dynamically changed network connection topology [43, 44, 45]. These methods are mainly motivated by that smaller models are more efficient to train at earlier epochs. This idea is also explored in language models [46, 47], recommendation systems [48] and graph ConvNets [49]. Locally supervised learning, which trains different model components using tailored objectives, is a promising direction as well [50, 51, 52].

A similar work to us is progressive learning (PL) [53], which down-samples the images to save the training cost. Nevertheless, our work differs from PL in several important aspects: 1) *EfficientTrain* is drawn from a distinctly different motivation of generalized curriculum learning, based on which we present a novel frequency-inspired analysis; 2) we introduce a cropping operation in the frequency domain, which is not only theoretically different from the down-sampling operation in PL (see: Proposition 1), but also outperforms it empirically (see: Table 13 (b)); 3) from the perspective of system-level comparison, we design an *EfficientTrain* curriculum, achieving a significantly higher training efficiency than PL on a variety of state-of-the-art models (see: Tables 9). In addition, FixRes [54] shows that a smaller training resolution may improve the accuracy by fixing the discrepancy between the scale of training and test inputs. However, we do not borrow gains from FixRes as we adopt a standard resolution at the final stages of training. Our method is actually orthogonal to FixRes (see: Table 9).

Frequency-based analysis of deep networks. Our observation that deep networks tend to capture the low-frequency components first is inline with [55], but the discussions in [55] focus on the robustness of ConvNets and are mainly based on some small models and tiny datasets. Towards this direction, several existing works also explore decomposing the inputs of models in the frequency domain

[56, 57, 58] in order to understand or improve the robustness of the networks. In contrast, our aim is to improve the training efficiency of modern deep visual backbones.

3. A Generalization of Curriculum Learning

As uncovered in previous research, machine learning algorithms generally benefit from a ‘starting small’ strategy, *i.e.*, to first learn certain easier aspects of a task, and increase the level of difficulty progressively [23, 24, 8]. The dominant implementation of this idea, curriculum learning, proposes to introduce gradually more difficult examples during training [12, 9, 10]. In specific, a curriculum is defined on top of the training process to determine whether or not each sample should be leveraged at a given epoch (Figure 1 (a)).

On the limitations of sample-wise curriculum learning. Although curriculum learning has been widely explored from the lens of the sample-wise regime, its extensive application is usually limited by two major issues. *First*, differentiating between ‘easier’ and ‘harder’ training data is non-trivial. It typically requires deploying additional deep networks as a ‘teacher’ or exploiting specialized automatic learning approaches [25, 11, 16, 14, 12]. The resulting implementation complexity and the increased overall computational cost are both noteworthy weaknesses in terms of improving the training efficiency. *Second*, it is challenging to attain a principled approach that specifies which examples should be attended to at the earlier stages of learning. As a matter of fact, the ‘easy to hard’ strategy is not always helpful [9]. The hard-to-learn samples can be more informative and may be beneficial to be emphasized in many cases [59, 60, 61, 62, 63, 64], sometimes even leading to a ‘hard to easy’ anti-curriculum [65, 66, 67, 68, 69, 70].

Our work is inspired by the above two issues. In the following, we start by proposing a generalized hypothesis for curriculum learning, aiming to address the second issue. Then we demonstrate that an implementation of our idea naturally addresses the first issue.

Generalized curriculum learning. We argue that simply measuring the easiness of training samples tends to be ambiguous and may be insufficient to reflect the effects of a sample on the learning process. As aforementioned, even the difficult examples may provide beneficial information for guiding the training, and they do not necessarily need to be introduced after easier examples. To this end, we hypothesize that every training sample, either ‘easier’ or ‘harder’, contains both *easier-to-learn* or *more accessible* patterns, as well as certain *difficult* discriminative information which may be challenging for the deep networks to capture. Ideally, a curriculum should be a continuous function on top of the training process, which starts with a focus on the ‘easiest’ patterns of the inputs, while the ‘harder-to-learn’ patterns are gradually introduced as learning progresses.

A formal illustration is shown in Figure 1 (b). Any input

data \mathbf{X} will be processed by a transformation function $\mathcal{T}_t(\cdot)$ conditioned on the training epoch t ($t \leq T$) before being fed into the model, where $\mathcal{T}_t(\cdot)$ is designed to dynamically filter out the excessively difficult and less learnable patterns within the training data. We always let $\mathcal{T}_T(\mathbf{X}) = \mathbf{X}$. Notably, our approach can be seen as a generalized form of the sample-wise curriculum learning. It reduces to example-selection by setting $\mathcal{T}_t(\mathbf{X}) \in \{\emptyset, \mathbf{X}\}$.

Overview. In the rest of this paper, we will demonstrate that an algorithm drawn from our hypothesis dramatically improves the implementation efficiency and generalization ability of curriculum learning. We will show that a zero-cost criterion pre-defined by humans is able to effectively measure the difficulty level of different patterns within images. Based on such simple criteria, even a surprisingly straightforward implementation of introducing ‘easier-to-harder’ patterns yields significant and consistent improvements on the training efficiency of modern visual backbones.

4. The EfficientTrain Approach

To obtain a training curriculum following our aforementioned hypothesis, we need to solve two challenges: 1) identifying the ‘easier-to-learn’ patterns and designing transformation functions to extract them; 2) establishing a curriculum learning schedule to perform these transformations dynamically during training. This section will demonstrate that a proper transformation for 1) can be easily found in both the frequency and the spatial domain, while 2) can be addressed with a greedy search algorithm. Implementation details of the experiments in this section: see Appendix A.

4.1. Easier-to-learn Patterns: Frequency Domain

Image-based data can naturally be decomposed in the frequency domain [71, 72, 73, 74]. In this subsection, we reveal that the patterns in the lower-frequency components of images, which describe the smoothly changing contents, are relatively easier for the networks to learn to recognize.

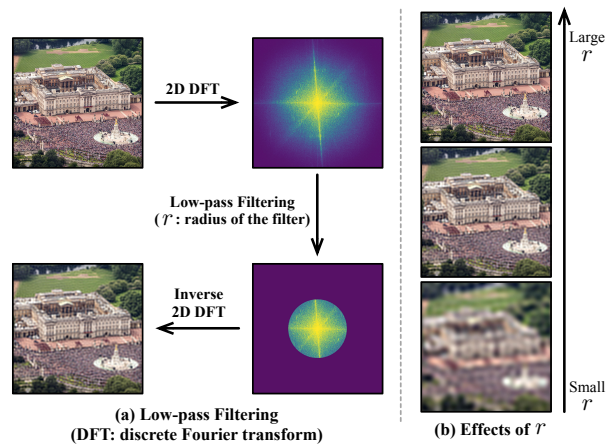
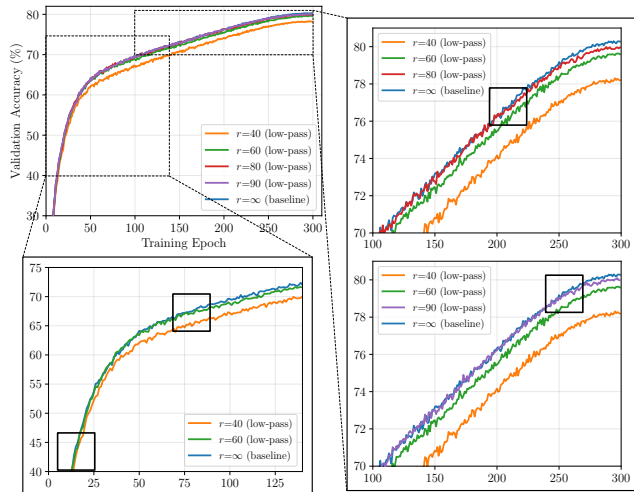


Figure 2: **Low-pass filtering.** Following [55], we adopt a circular filter.



Train & Val. : Low-pass Filtered Images

Figure 3: **Ablation studies with low-pass filtering** (r : bandwidth, see Figure 2). We ablate the higher-frequency components of the inputs of a DeiT-Small [19], and present the curves of validation accuracy v.s. training epochs on ImageNet-1K. Black boxes indicate the separation points of the curves. The two vertical graphs on the right highlight the separation points of $r = 80/r = 90$ and baseline (they will overlap if presented in one graph).

Ablation studies with the low-pass filtered input data.

We first consider an ablation study, where the low-pass filtering is performed on the data we use. As shown in Figure 2 (a), we map the images to the Fourier spectrum with the lowest frequency at the centre, set all the components outside a centred circle (radius: r) to zero, and map the spectrum back to the pixel space. Figure 2 (b) illustrates the effects of r . The curves of accuracy v.s. training epochs on top of the low-pass filtered data are presented in Figure 3. Here both training and validation data is processed by the filter to ensure the compatibility with the i.i.d. assumption.

Lower-frequency components are captured first. The models in Figure 3 are imposed to leverage only the lower-frequency components of the inputs. However, an appealing phenomenon arises: their training process is approximately identical to the original baseline at the beginning of training. Although the baseline finally outperforms, this tendency starts midway in the training process, instead of from the very beginning. In other words, the learning behaviors at earlier epochs remain unchanged even though the higher-frequency components of images are eliminated. Moreover, consider increasing the filter bandwidth r , which preserves progressively more information about the images from the lowest frequency. The separation point between the baseline and the training process on low-pass filtered data moves towards the end of training. To explain these observations, we postulate that, *in a natural learning process where the input images contain both lower- and higher-frequency information, a model tends to first learn to capture the lower-frequency components, while the higher-frequency information is gradually exploited on the basis of them.*

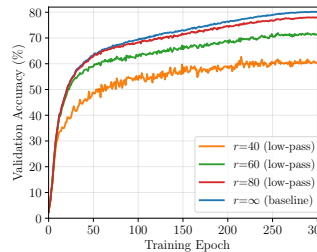


Figure 4: **Performing low-pass filtering only on the validation inputs** (other setups are the same as Figure 3). We train a model using the original images without any filtering (*i.e.*, containing both lower- and higher-frequency components), and evaluate all the intermediate checkpoints on low-pass filtered validation sets with varying bandwidths.

Train: Original Images (Low+High Frequency); Val. : Low-pass Filtered Images;

Training Epoch	20 th	50 th	100 th	200 th	300 th (end)
Low-pass Filtered Val. Set	<u>46.9%</u>	<u>58.8%</u>	<u>63.1%</u>	<u>68.6%</u>	<u>71.3%</u>
High-pass Filtered Val. Set	23.9%	43.5%	53.5%	64.3%	71.3%

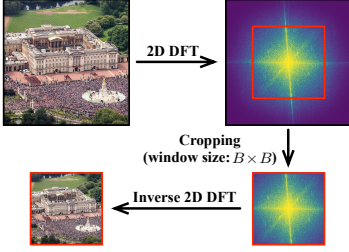
Table 1: **Comparisons: evaluating the model in Figure 4 on low/high-pass filtered validation sets.** Note that the model is trained using the original images without any filtering. The bandwidths of the low/high-pass filters are configured to make the finally trained model (300th epoch) have the same accuracy on the two validation sets (*i.e.*, 71.3%). Here the acc. on the low-pass filtered val. set corresponds to ' $r = 60$ ' in Figure 4.

More evidences. Our assumption can be further confirmed by a well-controlled experiment. Consider training a model using original images, where lower/higher-frequency components are simultaneously provided. In Figure 4, we evaluate all the intermediate checkpoints on low-pass filtered validation sets with varying bandwidths. Obviously, at earlier epochs, only leveraging the low-pass filtered validation data does not degrade the accuracy. This phenomenon suggests that the learning process starts with a focus on the lower-frequency information, even though the model is always accessible to higher-frequency components during training. Furthermore, in Table 1, we compare the accuracies of the intermediate checkpoints on low/high-pass filtered validation sets. We find the accuracy on the low-pass filtered validation set grows much faster at earlier training stages, even though the two final accuracies are the same.

Curricula (ep: epoch)	Low-pass Filtered Training Data	Original Training Data	Final Top-1 Accuracy (r : filter bandwidth)		
			$r=40$	$r=60$	$r=80$
1 st - 300 th ep	-	-	78.3%	79.6%	80.0%
1 st - 225 th ep	226 th - 300 th ep	-	79.4%	<u>80.2%</u>	<u>80.5%</u>
1 st - 150 th ep	151 th - 300 th ep	-	<u>80.1%</u>	80.2%	80.6%
1 st - 75 th ep	76 th - 300 th ep	-	80.3%	80.4%	80.6%
-	1 st - 300 th ep	-	80.3% (baseline)		

Table 2: **Results with the straightforward frequency-based training curricula** (DeiT-Small [19] on ImageNet-1K). Observation: one can eliminate the higher-frequency components of the inputs in 50-75% of the training process without sacrificing the final accuracy (see: underlined data).

Frequency-based curricula. Returning to our hypothesis in Section 3, we have shown that lower-frequency components are naturally captured earlier. Hence, it would be straightforward to consider them as a type of the 'easier-to-learn' patterns. This begs a question: can we design a training curriculum, which starts with providing only the lower-frequency information for the model, while gradually introducing the higher-frequency components? We investigate this idea in Table 2, where we perform low-pass filtering on the training data only in a given number of the beginning



Curricula (ep: epoch)		Final Top-1 Accuracy / Relative Computational Cost for Training (compared to the baseline)							
Low-frequency Cropping (B^2)	Original Training Data ($B=224$)	DeiT-Small [19]				Swin-Tiny [4]			
		$B=96$	$B=128$	$B=160$	$B=192$	$B=96$	$B=128$	$B=160$	$B=192$
1 st – 300 th ep	–	70.5%/0.18	75.3%/0.31	77.9%/0.49	79.1%/0.72	73.3%/0.18	76.8%/0.32	78.9%/0.50	80.5%/0.73
1 st – 225 th ep	226 th – 300 th ep	78.7%/0.38	79.6%/0.48	80.0%/0.62	<u>80.3%/0.79</u>	80.0%/0.38	80.5%/0.49	81.0%/0.63	<u>81.2%/0.80</u>
1 st – 150 th ep	151 th – 300 th ep	79.2%/0.59	79.8%/0.66	<u>80.3%/0.75</u>	80.4%/0.86	80.9%/0.59	80.9%/0.66	<u>81.2%/0.75</u>	81.3%/0.86
1 st – 75 th ep	76 th – 300 th ep	79.6%/0.79	<u>80.2%/0.83</u>	80.4%/0.87	80.3%/0.93	<u>81.2%/0.79</u>	<u>81.2%/0.83</u>	81.3%/0.88	81.3%/0.93
–	1 st – 300 th ep	80.3%/1.00 (baseline)				81.3%/1.00 (baseline)			

Figure 5: **Low-frequency cropping in the frequency domain** (B^2 : bandwidth). Table 3: **Results on ImageNet-1K with the low-pass filtering in Table 2 replaced by the low-frequency cropping**, which yields competitive accuracy with a significantly reduced training cost (see: underlined data).

epochs. The rest of the training process remains unchanged.

Learning from low-frequency information efficiently.

At the first glance, the results in Table 2 may be less dramatic, *i.e.*, by processing the images with a properly-configured low-pass filter at earlier epochs, the accuracy is moderately improved. However, an important observation is noteworthy: the final accuracy of the model can be largely preserved even with aggressive filtering (*e.g.*, $r = 40, 60$) performed in 50-75% of the training process. This phenomenon turns our attention to training efficiency. At earlier learning stages, it is harmless to train the model with only the lower-frequency components. These components incorporate only a selected subset of all the information within the original input images. Hence, *can we enable the model to learn from them efficiently with less computational cost than processing the original inputs?* As a matter of fact, this idea is feasible, and we may have at least two approaches.

- 1) **Down-sampling.** Approximating the low-pass filtering in Table 2 with image down-sampling may be a straightforward solution. Down-sampling preserves much of the lower-frequency information, while it quadratically saves the computational cost for a model to process the inputs [74, 75, 76, 77, 78]. However, it is not an operation tailored for extracting lower-frequency components. Theoretically, it preserves some of the higher-frequency components as well (see: Proposition 1). Empirically, we observe that this issue degrades the performance (see: Table 13 (b)).

- 2) **Low-frequency cropping** (see: Figure 5). We propose a more precise approach that extracts exactly all the lower-frequency information. Consider mapping an $H \times W$ image \mathbf{X} into the frequency domain with the 2D discrete Fourier transform (DFT), obtaining an $H \times W$ Fourier spectrum, where the value in the centre denotes the strength of the component with the lowest frequency. The positions distant from the centre correspond to higher-frequency. We crop a $B \times B$ patch from the centre of the spectrum, where B is the window size ($B < H, W$). Since the patch is still centrosymmetric, we can map it back to the pixel space with the inverse 2D DFT, obtaining a $B \times B$ new image \mathbf{X}_c , namely

$$\mathbf{X}_c = \mathcal{F}^{-1} \circ \mathcal{C}_{B,B} \circ \mathcal{F}(\mathbf{X}) \in \mathbb{R}^{B \times B}, \quad \mathbf{X} \in \mathbb{R}^{H \times W}, \quad (1)$$

where \mathcal{F} , \mathcal{F}^{-1} and $\mathcal{C}_{B,B}$ denote 2D DFT, inverse 2D DFT and B^2 centre-cropping. The computational or the time cost

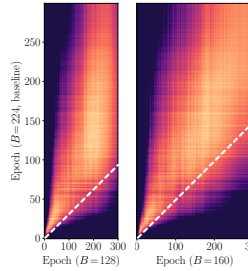


Figure 6: **CKA feature similarity heatmaps** [79, 80] between the DeiT-S [19] trained using the inputs with low-frequency cropping ($B=128, 160$) and the original inputs ($B=224$). The X and Y axes index training epochs (scaled according to the computational cost of training). Here we feed the same original images into all the models (including the ones trained with $B=128, 160$) and take the features from the final layer. The 45° lines are highlighted in white.

for accomplishing Eq. (1) is negligible on GPUs.

Notably, \mathbf{X}_c achieves a lossless extraction of lower-frequency components, while the higher-frequency parts are strictly eliminated. Hence, feeding \mathbf{X}_c into the model at earlier training stages can provide the vast majority of the useful information, such that the final accuracy will be minimally affected or even not affected. In contrast, importantly, due to the reduced input size of \mathbf{X}_c , the computational cost for a model to process \mathbf{X}_c is able to be dramatically saved, yielding a considerably more efficient training process.

Our claims can be empirically supported by Table 3, where we replace the low-pass filtering in Table 2 with the low-frequency cropping. Even such a straightforward implementation yields favorable results: the training cost can be saved by $\sim 20\%$ while a competitive final accuracy is preserved. This phenomenon can be interpreted via Figure 6: at intermediate training stages, the models trained using the inputs with low-frequency cropping can learn similar deep representations to the baseline with a significantly reduced cost, *i.e.*, the bright parts are clearly above the white lines.

Proposition 1. *Suppose that $\mathbf{X}_c = \mathcal{F}^{-1} \circ \mathcal{C}_{B,B} \circ \mathcal{F}(\mathbf{X})$, and that $\mathbf{X}_d = \mathcal{D}_{B,B}(\mathbf{X})$, where $B \times B$ down-sampling $\mathcal{D}_{B,B}(\cdot)$ is realized by a common interpolation algorithm (*e.g.*, nearest, bilinear or bicubic). Then we have two properties.*

- a) \mathbf{X}_c is only determined by the lower-frequency spectrum of \mathbf{X} (*i.e.*, $\mathcal{C}_{B,B} \circ \mathcal{F}(\mathbf{X})$). In addition, the mapping to \mathbf{X}_c is reversible. We can always recover $\mathcal{C}_{B,B} \circ \mathcal{F}(\mathbf{X})$ from \mathbf{X}_c .
- b) \mathbf{X}_d has a non-zero dependency on the higher-frequency spectrum of \mathbf{X} (*i.e.*, the regions outside $\mathcal{C}_{B,B} \circ \mathcal{F}(\mathbf{X})$).

Proof. See: Appendix C. □

		Curricula (ep: epoch)		Final Top-1 Accuracy				
Weaker	RandAug	Low-frequency	Original	$(m: \text{magnitude of RandAug})$				
RandAug	$(m=9)$	$(B=128)$	$(B=224)$	$m=1$	$m=3$	$m=5$	$m=7$	$m=9$
1 st - 150 th ep	151 th - 300 th ep	-	1 st -300 th ep	80.4%	80.6%	80.7%	80.5%	80.3%
		1 st -150 th ep	151 th -300 th ep	80.2%	80.2%	80.2%	79.9%	79.8%

Table 4: **Performance of the data-augmentation-based curricula** (DeiT-Small [19] on ImageNet-1K). We test reducing the magnitude of RandAug from 1st-150th training epochs ($m=9$ refers to the baselines).

4.2. Easier-to-learn Patterns: Spatial Domain

Apart from the frequency domain operations, extracting ‘easier-to-learn’ patterns can also be attained through spatial domain transformations. For example, deep networks (e.g., ViTs [19, 4, 21] and ConvNets [18, 81, 82, 83, 84]) are typically trained with strong and delicate data augmentation techniques. We argue that the augmented training data provides a combination of both the information from original samples and the information introduced by augmentation operations. The original patterns may be ‘easier-to-learn’ as they are drawn from real-world distributions. This assumption is supported by the observation that data augmentation is mainly influential at the later stages of training [85].

To this end, following our generalized formulation of curriculum learning in Section 3, a curriculum may adopt a weaker-to-stronger data augmentation strategy during training. We investigate this idea by selecting RandAug [86] as a representative example, which incorporates a family of common spatial-wise data augmentation transformations (rotate, sharpness, shear, solarize, etc.). In Table 4, the magnitude of RandAug is varied in the first half training process. One can observe that this idea improves the accuracy, and the gains are compatible with low-frequency cropping.

4.3. A Unified Training Curriculum

Finally, we integrate the techniques discussed above (i.e., low-frequency cropping at earlier epochs and weaker-to-stronger RandAug) to design a unified efficient training curriculum. We first set the magnitude m of RandAug to be a linear function of the epoch t : $m = (t/T) \times m_0$, with other data augmentation unchanged. Although being simple, this setting yields consistent empirical improvements. We adopt $m_0 = 9$ following the common practice [19, 20, 4, 21, 18].

Then we propose a greedy-search algorithm (Alg. 1) to determine the schedule of the bandwidth B during training for low-frequency cropping. We divide the full training process into several stages and solve for a value of B for each stage (a staircase approximation of the continuous curriculum learning schedule; see Appendix B.2 for more discussions). Alg. 1 starts from the last stage, minimizing B under the constraint of not degrading the performance compared to the baseline accuracy a_0 . Here a_0 is obtained by training a model with a fixed $B = 224$, and does not change throughout Alg. 1. In Alg. 1, ValidationAccuracy(\cdot) refers to training a new model with the given choices of B , and

evaluating the accuracy. The input T is used here. For implementation, we only execute Alg. 1 for a single time. We obtain a schedule on top of Swin-Tiny [4] under the standard 300-epoch training setting [4] with $N = 5$, and directly adopt this schedule for other models or other training settings. Notably, the cost of Alg. 1 is to train a relatively small model (e.g., Swin-Tiny) for a small number of times (e.g., 7 to obtain our curriculum), where we can solve the minimization problems in Alg. 1 via simple linear search.

Algorithm 1 Algorithm to Solve for the Curriculum.

- 1: **Input:** Number of training epochs T and training stages N (i.e., T/N epochs for each stage).
- 2: **Input:** Baseline accuracy a_0 (with 224^2 images).
- 3: **To solve:** The value of B for i^{th} training stage: \hat{B}_i .
- 4: **Initialize:** $\hat{B}_1 = \dots = \hat{B}_N = 224$
- 5: **for** $i = N - 1$ **to** 1 **do**
- 6: $\hat{B}_i = \underset{B_1 = \dots = B_i = B, B_j = \hat{B}_j, i < j \leq N}{\text{minimize}} B,$
 s.t. ValidationAccuracy(B_1, \dots, B_N) $\geq a_0$
- 7: **end for**
- 8: **Output:** $\hat{B}_1, \dots, \hat{B}_N$

Derived from the aforementioned procedure, our finally proposed curriculum is presented in Table 5, which is named as *EfficientTrain*. Notably, despite its simplicity, it is general and surprisingly effective. It can be directly applied to most visual backbones under various training settings (e.g., different training budgets, varying amounts of training data, and supervised/self-supervised learning algorithms) *without tuning additional hyper-parameters*, and contributes to a significantly improved training efficiency.

Epochs	Low-frequency Cropping	RandAug
1 st - 180 th	$B = 160$	$m = 0 \rightarrow 9$
181 th - 240 th	$B = 192$	Increase linearly.
241 th - 300 th	$B = 224$	

Table 5: **EfficientTrain curriculum** (with the standard 300-epoch training and 224^2 final input size [4]). Notably, it can straightforwardly adapt to varying epochs or final input sizes by simple linear scaling (see: Table 8).

5. Experiments

Datasets. Our main experiments are based on the large-scale ImageNet-1K/22K [17] datasets, which consist of $\sim 1.28\text{M}/\sim 14.2\text{M}$ images in 1K/ $\sim 22\text{K}$ classes. We verify the transferability of pre-trained models on MS COCO [87], CIFAR [88], Flowers-102 [89], and Stanford Dogs [90].

Models. A wide variety of visual backbones are considered in our experiments, including ResNet [1], ConvNeXt [18], DeiT [19], PVT [20], Swin [4] and CSWin [21] Transformers. We adopt the training pipeline in [4, 18], where EfficientTrain only modifies the terms mentioned in Table 5. Unless otherwise specified, we report the results of our implementation for both our method and the baselines. More implementation details can be found in Appendix A.

5.1. Supervised Learning

Training various visual backbones on ImageNet-1K. Table 6 presents the results of applying our method to

Model	Input Size (inference)	#Param.	#FLOPs	Top-1 Accuracy (300 epochs)			Training Speedup		
				Original Paper	Baseline (ours)	EfficientTrain	Computation	Wall-time	
ConvNets	ResNet-50 [1]	224 ²	26M	4.1G	–	78.8%	79.4%	1.53×	1.44×
	ConvNeXt-Tiny [18]	224 ²	29M	4.5G	82.1%	82.1%	82.2%	1.53×	1.49×
	ConvNeXt-Small [18]	224 ²	50M	8.7G	83.1%	83.1%	83.2%	1.53×	1.50×
	ConvNeXt-Base [18]	224 ²	89M	15.4G	83.8%	83.8%	83.7%	1.53×	1.48×
Isotropic ViTs	DeiT-Tiny [19]	224 ²	5M	1.3G	72.2%	72.5%	73.3%	1.59×	1.55×
	DeiT-Small [19]	224 ²	22M	4.6G	79.9%	80.3%	80.4%	1.56×	1.51×
Multi-stage ViTs	PVT-Tiny [20]	224 ²	13M	1.9G	75.1%	75.5%	75.5%	1.55×	1.48×
	PVT-Small [20]	224 ²	25M	3.8G	79.8%	79.9%	80.4%	1.55×	1.56×
	PVT-Medium [20]	224 ²	44M	6.7G	81.2%	81.8%	81.8%	1.54×	1.49×
	PVT-Large [20]	224 ²	61M	9.8G	81.7%	82.3%	82.3%	1.54×	1.53×
	Swin-Tiny [4]	224 ²	28M	4.5G	81.3%	81.3%	81.4%	1.55×	1.49×
	Swin-Small [4]	224 ²	50M	8.7G	83.0%	83.1%	83.2%	1.54×	1.50×
	Swin-Base [4]	224 ²	88M	15.4G	83.5%	83.4%	83.6%	1.54×	1.50×
	CSWin-Tiny [21]	224 ²	23M	4.3G	82.7%	82.7%	82.8%	1.56×	1.55×
	CSWin-Small [21]	224 ²	35M	6.9G	83.6%	83.4%	83.6%	1.56×	1.51×
CSWin-Base [21]	224 ²	78M	15.0G	84.2%	84.3%	84.3%	1.55×	1.56×	

Table 6: **Results on ImageNet-1K (IN-1K)**. We train the models w/ or w/o EfficientTrain on the IN-1K training set, and report the accuracy on the IN-1K validation set. For fair comparisons, we also report the baselines in the original papers. The training wall-time is benchmarked on NVIDIA 3090 GPUs.

Model	Input Size (inference)	#Param.	#FLOPs	Top-1 Accuracy (fine-tuned to ImageNet-1K)		Wall-time Pre-training Cost (in GPU-days)		Time Saved (for an 8-GPU node)
				Baseline	EfficientTrain	Baseline	EfficientTrain	
ConvNeXt-Base [18]	224 ²	89M	15.4G	85.6%	85.7%	170.6	114.8 (↓1.49×)	6.98 Days
	384 ²	89M	45.1G	86.7%	86.8%			
ConvNeXt-Large [18]	224 ²	198M	34.4G	86.4%	86.3%	347.6	225.5 (↓1.54×)	15.26 Days
	384 ²	198M	101.0G	87.3%	87.3%			
CSWin-Base [21]	224 ²	78M	15.0G	85.5%	85.6%	238.9	157.7 (↓1.52×)	10.16 Days
	384 ²	78M	47.0G	86.7%	87.0%			
CSWin-Large [21]	224 ²	173M	31.5G	86.5%	86.6%	469.5	307.7 (↓1.53×)	20.23 Days
	384 ²	173M	96.8G	87.6%	87.8%			

Table 7: **Results with ImageNet-22K (IN-22K) pre-training**. The models are pre-trained on IN-22K w/ or w/o EfficientTrain, fine-tuned on the ImageNet-1K (IN-1K) training set, and evaluated on the IN-1K validation set. The wall-time pre-training cost is benchmarked on NVIDIA 3090 GPUs.

train representative deep networks on ImageNet-1K. EfficientTrain achieves consistent improvements across different models, *i.e.*, it reaches a competitive or better validation accuracy compared to the baselines (*e.g.*, 83.6% v.s. 83.4% on the CSWin-Small network), while saving the training cost by 1.5 – 1.6×. Importantly, the practical speedup is consistent with the theoretical results. The detailed training runtime (GPU-hours) is deferred to Appendix B.1.

ImageNet-22K pre-training. Our method exhibits excellent scalability with a growing training data scale or an increasing model size. In Table 7, a number of large backbones are pre-trained on ImageNet-22K w/ or w/o EfficientTrain, and evaluated by being fine-tuned to ImageNet-1K. Note that pre-training accounts for the vast majority of the total computation/time cost in this procedure. Our method performs at least on par with the baselines, while achieving a significant training speedup. A highlight from the results is that EfficientTrain saves the real training time considerably, *e.g.*, 162 GPU-days (307.7 v.s. 469.5) for CSWin-Large, corresponding to ~20 days for a 8-GPU node.

Adapted to varying epochs. EfficientTrain can conveniently adapt to a varying length of training schedule, *i.e.*, by simply scaling the indices of epochs in Table 5. As shown in Table 8 (a), the advantage of our method is even

Models	Input Size (inference)	Top-1 Accuracy (baseline / EfficientTrain)			Wall-time Training Speedup
		100 epochs	200 epochs	300 epochs	
DeiT-Tiny [19]	224 ²	65.8% / 68.1%	70.5% / 71.8%	72.5% / 73.3%	1.55×
DeiT-Small [19]	224 ²	75.5% / 76.4%	79.0% / 79.1%	80.3% / 80.4%	1.51×
Swin-Tiny [4]	224 ²	78.4% / 78.5%	80.6% / 80.6%	81.3% / 81.4%	1.49×
Swin-Small [4]	224 ²	80.6% / 80.7%	82.7% / 82.6%	83.1% / 83.2%	1.50×
Swin-Base [4]	224 ²	80.7% / 81.1%	83.2% / 83.2%	83.4% / 83.6%	1.50×

(a) **Reducing training cost with the same number of epochs.**

Models	Input Size (inference)	Top-1 Accuracy (ep: epochs)		Wall-time Training Speedup
		Baseline _{300ep}	EfficientTrain _{450ep}	
DeiT-Tiny [19]	224 ²	72.5%	74.3% (+1.8)	1.04×
DeiT-Small [19]	224 ²	80.3%	80.9% (+0.6)	1.01×

(b) **Higher accuracy with the same training cost.**

Models	Method	Top-1 Accuracy / Wall-time Training Speedup		
		224 ²	384 ²	512 ²
Swin-Base [4]	Baseline	83.4%/1.00×	84.5%/1.00×	84.7%/1.00×
	EfficientTrain	83.6%/1.50×	84.7%/2.91×	85.1%/3.37×

(c) **Adapted to different final input sizes.** Swin-B is selected as a representative example since larger models benefit more from larger input sizes.

Table 8: **ImageNet-1K results with varying epochs and final input sizes.**

more significant with fewer training epochs, *e.g.*, it outperforms the baseline by 0.9% (76.4% v.s. 75.5%) for the 100-epoch trained DeiT-Small (the speedup is the same as 300-epoch). We attribute this to the greater importance of efficient training algorithms in the scenarios of limited training

Models	Training Approach	Training Epochs	Aug-Regs	Top-1 Accuracy	Wall-time Training Speedup
ResNet-18 [1]	Curriculum by Smoothing [37] (NearIPS20)	90	✗	71.0%	1.00×
	EfficientTrain	90	✗	71.0%	1.48×
	Self-paced Learning [25] (NearIPS10)	200	✗	73.2%	1.15×
	Minimax Curriculum Learning [67] (ICLR18)	200	✗	75.1%	1.97×
	DIH Curriculum Learning [64] (NearIPS20)	200	✗	76.3%	2.45×
ResNet-50 [1]	EfficientTrain	200	✗	77.5%	1.44×
	CurriculumNet [91] (ECCV18)	90	✗	76.1%	< 2.22×
	Label-sim. Curriculum Learning [92] (ECCV20)	90	✗	76.9%	2.22×
	EfficientTrain	90	✗	77.0%	3.21×
	Progressive Learning [53] (ICML21)	350	✓	78.4%	1.21×
DeiT-Tiny [19]	EfficientTrain	300	✓	79.4%	1.44×
	Auto Progressive Learning [7] (CVPR22)	300	✓	72.4%	1.51×
DeiT-Small [19]	Progressive Learning [53] (ICML21)	100	✓	72.6%	1.54×
	Auto Progressive Learning [7] (CVPR22)	100	✓	74.4%	1.41×
	Budgeted ViT [93] (ICLR23)	128	✓	74.5%	1.34×
	EfficientTrain	100	✓	76.4%	1.51×
	Budgeted Training† [94] (ICLR20)	225	✓	79.6%	1.33×
	Progressive Learning† [53] (ICML21)	300	✓	79.5%	1.49×
	Auto Progressive Learning [7] (CVPR22)	300	✓	79.8%	1.42×
	DeiT III [95] (ECCV22)	300	✓	79.9%	1.00×
	Budgeted ViT [93] (ICLR23)	303	✓	80.1%	1.34×
	EfficientTrain	300	✓	80.4%	1.51×
CSWin-Tiny [21]	Progressive Learning† [53] (ICML21)	300	✓	82.3%	1.51×
	EfficientTrain	300	✓	82.8%	1.55×
	FixRes† [54] (NearIPS19)	300	✓	82.9%	1.00×
CSWin-Small [21]	EfficientTrain+FixRes	300	✓	83.1%	1.55×
	Progressive Learning† [53] (ICML21)	300	✓	83.3%	1.48×
	EfficientTrain	300	✓	83.6%	1.51×
CSWin-Small [21]	FixRes† [54] (NearIPS19)	300	✓	83.7%	1.00×
	EfficientTrain+FixRes	300	✓	83.8%	1.51×

Table 9: **EfficientTrain v.s. state-of-the-art efficient training algorithms on ImageNet-1K.** Here ‘AugRegs’ denotes the holistic combination of various model regularization and data augmentation techniques, which is widely applied to train deep networks effectively [19, 20, 4, 21, 18]. In particular, some baselines are not developed on top of this state-of-the-art training pipeline. To ensure a fair comparison with them, we also implement our method without AugRegs (notably, the ‘RandAug’ in Table 5 is removed in this scenario as well). †: our reproduced baselines.

resources. Table 8 (b) further shows that our method can significantly improve the accuracy with the same training wall-time as the baseline (e.g., by 1.8% for DeiT-Tiny).

Adapted to any final input size γ . To adapt to a final input size γ , the value of B for the three stages of EfficientTrain can be simply adjusted to $160, (160 + \gamma)/2, \gamma$. As shown in Table 8 (c), our method outperforms the baseline by large margins for $\gamma > 224$ in terms of training efficiency.

Orthogonal to 224^2 pre-training + γ^2 fine-tuning. In particular, in some cases, existing works find it efficient to fine-tune 224^2 pre-trained models to a target test input size γ^2 [4, 18, 21]. Here our method can be directly leveraged for more efficient pre-training (e.g., $\gamma = 384$ in Table 7).

Comparisons with state-of-the-art efficient training methods are summarized in Table 9. EfficientTrain outperforms all the recently proposed baselines in terms of both accuracy and training efficiency. Moreover, the simplicity of our method enables it to be conveniently applied to different models and training settings, which may be non-trivial for other methods (e.g., the network-growing method [7]).

Method	#Param.	Pre-training Epochs	Top-1 Accuracy (fine-tuning)		Pre-training Speedup	
			Baseline	EfficientTrain	Computation	Wall-time
MAE (ViT-B) [22]	86M	1600	83.5%†	83.6%	1.54×	1.52×
MAE (ViT-L) [22]	307M	1600	85.9%†	85.8%	1.53×	1.55×

Table 10: **Self-supervised learning results with MAE [22].** Following [22], the models are pre-trained on ImageNet-1K w/ or w/o EfficientTrain, and evaluated by end-to-end fine-tuning. †: our reproduced baselines.

Backbone	Pre-training Method	Pre-training Speedup	Top-1 Accuracy (fine-tuned to downstream datasets)			
			CIFAR-10	CIFAR-100	Flowers-102	Stanford Dogs
DeiT-S [19]	Baseline	1.0×	98.39%	88.65%	96.57%	90.72%
	EfficientTrain	1.5×	98.47%	88.93%	96.62%	91.12%

Table 11: **Transferability to downstream image recognition tasks.** The backbone model is pre-trained on ImageNet-1K w/ or w/o EfficientTrain, and fine-tuned to the downstream datasets to report the accuracy.

Method	Backbone	Pre-training Method	Pre-training Speedup	Apbox						
				AP ₅₀	AP ₇₅	AP _{mask}	AP ₅₀ ^{mask}	AP ₇₅ ^{mask}		
RetinaNet [96] (1× schedule)	Swin-T [4]	Baseline	1.0×	41.7	62.9	44.5	–	–	–	
		EfficientTrain	1.5×	41.8	63.3	44.6	–	–	–	
Cascade Mask-RCNN [97] (1× schedule)	Swin-T [4]	Baseline	1.0×	48.1	67.0	52.1	41.5	64.2	44.9	
		EfficientTrain	1.5×	48.2	67.5	52.3	41.8	64.6	45.0	
Mask-RCNN [97] (1× schedule)	Swin-S [4]	Baseline	1.0×	50.0	69.1	54.4	43.1	66.3	46.3	
		EfficientTrain	1.5×	50.6	69.7	55.2	43.6	66.9	47.3	
	Swin-B [4]	Baseline	1.0×	50.9	70.2	55.5	44.0	67.4	47.4	
		EfficientTrain	1.5×	51.3	70.6	56.0	44.2	67.7	47.7	

Table 12: **Object detection and instance segmentation on COCO.** We implement representative detection/segmentation algorithms on top of the backbones pre-trained w/ or w/o EfficientTrain on ImageNet-1K.

Orthogonal to FixRes. FixRes [54] reveals that there exists a discrepancy in the scale of images between the training and test inputs, and thus the inference with a larger resolution will yield a better test accuracy. However, EfficientTrain does not leverage the gains of FixRes. We adopt the original inputs (e.g., 224^2) at the final training stage, and hence the finally-trained model resembles the 224^2 -trained networks, while FixRes is orthogonal to it. This fact can be confirmed by the evidence in both Table 9 (see: FixRes v.s. EfficientTrain+FixRes on top of the state-of-the-art CSWin Transformers [21]) and Table 7 (see: Input Size= 384^2).

5.2. Self-supervised Learning

Results with Masked Autoencoders (MAE). Since our method only modifies the training inputs, it can also be applied to self-supervised learning algorithms. As a representative example, we deploy EfficientTrain on top of MAE [22] in Table 10. Our method reduces the training cost significantly while preserving a competitive accuracy.

5.3. Transfer Learning

Downstream image recognition tasks. In Table 11, we fine-tune the models trained with EfficientTrain to downstream classification datasets. Following [19], the 32^2 CIFAR images [88] are resized to 224^2 , and hence their discriminative patterns are concentrated in the lower-frequency components. On the contrary, Flowers-102 [89] and Stanford Dogs [90] are fine-grained visual recognition datasets where the high-frequency clues contain important discriminative information. As shown in Table 11, EfficientTrain

Low-frequency Cropping	Linear RandAug	Training		Top-1 Accuracy (100ep: 100-epoch; Others: 300-epoch)				
		Speedup	DeiT-T	DeiT-S _{100ep}	DeiT-S	Swin-T	Swin-S	CSWin-T
		1.0×	72.5%	75.5%	80.3%	81.3%	83.1%	82.7%
✓		1.5×	72.4%	75.5%	80.0%	81.1%	83.0%	82.6%
✓	✓	1.5×	73.3%	76.4%	80.4%	81.4%	83.2%	82.8%

(a) Ablating low-frequency cropping or linearly increased RandAug.

Low-frequency Information Extraction in EfficientTrain	Training		Top-1 Accuracy (100ep: 100-epoch; Others: 300-epoch)				
	Speedup	DeiT-S _{100ep}	DeiT-S	Swin-T	Swin-S	CSWin-L [†]	
Image Down-sampling	1.5×	75.9%	80.3%	81.0%	83.0%	86.4%	
Low-frequency Cropping	1.5×	76.4%	80.4%	81.4%	83.2%	86.6%	

(b) Design choices of the operation for extracting low-frequency information. †: pre-trained on ImageNet-22K.

Schedule of B in EfficientTrain	Training		Top-1 Accuracy		
	Speedup	DeiT-T	DeiT-S	Swin-T	
Linear Increasing [53]	1.5×	72.8%	79.9%	81.0%	
Obtained from Algorithm 1	1.5×	73.3%	80.4%	81.4%	

(c) Schedule of B . For fair comparisons, the linear increasing schedule is configured to have the same training cost as our proposed schedule.

Table 13: Ablation studies of EfficientTrain on ImageNet-1K.

yields competitive transfer learning accuracy on both types of datasets. In other words, although our method learns to exploit the lower/higher-frequency information via an ordered curriculum, the finally obtained models can leverage both types of information effectively.

Object detection & instance segmentation. Table 12 investigates fine-tuning our pre-trained models to other computer vision tasks. Our method reduces the pre-training cost by 1.5×, but performs at least on par with the baselines in terms of the detection/segmentation performance.

5.4. Discussions

Ablation study. In Table 13 (a), we show that the major gain of training efficiency comes from low-frequency cropping, which effectively reduces the training cost at the price of a slight drop of accuracy. On top of this, linear RandAug further improves the accuracy. Moreover, replacing low-frequency cropping with image down-sampling consistently degrades the accuracy (see: Table 13 (b)), since down-sampling cannot strictly filter out all the higher-frequency information (see: Proposition 1), yielding a sub-optimal implementation of our idea. In addition, as shown in Table 13 (c), the schedule of B found by Alg. 1 outperforms the heuristic design choices (e.g., the linear schedule in [53]).

Curves of val. accuracy during training are shown in Figure 7. The horizontal axis denotes the wall-time training cost. The low-frequency cropping in EfficientTrain is performed on both the training and test inputs.

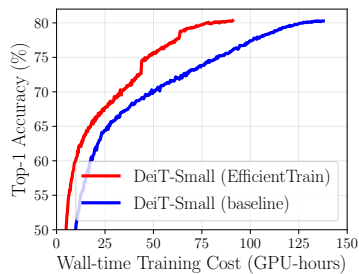


Figure 7: Curves of acc. during training. Our method is able to learn discriminative representations more efficiently at earlier epochs.

6. Conclusion

This paper investigated a novel generalized curriculum learning approach. The proposed algorithm, *EfficientTrain*, always leverages all the data at any training stage, but only exposes the ‘easier-to-learn’ patterns of each example at the beginning of training, and gradually introduces more difficult patterns as learning progresses. Our method significantly improves the training efficiency of state-of-the-art deep networks on the large-scale ImageNet-1K/22K datasets, for both supervised and self-supervised learning.

Acknowledgements

This work is supported in part by the National Key R&D Program of China under Grant 2021ZD0140407, the National Natural Science Foundation of China under Grants 62022048 and 62276150, the National Defense Basic Science and Technology Strengthening Program of China, Beijing Academy of Artificial Intelligence (BAAI), and Huawei Technologies Ltd.

References

- [1] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1, 6, 7, 8, 14, 16
- [2] Gao Huang, Zhuang Liu, Geoff Pleiss, Laurens Van Der Maaten, and Kilian Weinberger. Convolutional networks with dense connectivity. *TPAMI*, pages 1–1, 2019. 1
- [3] 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 *ICLR*, 2021. 1
- [4] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 1, 2, 5, 6, 7, 8, 14, 16
- [5] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *CVPR*, 2022. 1
- [6] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp. In *ACL*, pages 3645–3650, 2019. 1
- [7] Changlin Li, Bohan Zhuang, Guangrun Wang, Xiaodan Liang, Xiaojun Chang, and Yi Yang. Automated progressive learning for efficient training of vision transformers. In *CVPR*, 2022. 1, 8
- [8] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, pages 41–48, 2009. 1, 2, 3

- [9] Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *TPAMI*, 2021. 1, 2, 3
- [10] Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey. *International Journal of Computer Vision*, pages 1–40, 2022. 1, 2, 3
- [11] Alex Graves, Marc G Bellemare, Jacob Menick, Remi Munos, and Koray Kavukcuoglu. Automated curriculum learning for neural networks. In *ICML*, pages 1311–1320, 2017. 1, 2, 3
- [12] Guy Hacohen and Daphna Weinshall. On the power of curriculum learning in training deep networks. In *ICML*, pages 2535–2544, 2019. 1, 2, 3
- [13] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G Hauptmann. Self-paced curriculum learning. In *AAAI*, 2015. 1
- [14] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313, 2018. 1, 2, 3
- [15] Chen Gong, Dacheng Tao, Stephen J Maybank, Wei Liu, Guoliang Kang, and Jie Yang. Multi-modal curriculum learning for semi-supervised image classification. *TIP*, 25(7):3249–3260, 2016. 1
- [16] Yang Fan, Fei Tian, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. Learning to teach. In *ICLR*, 2018. 1, 2, 3
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 2, 6, 14, 15
- [18] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 2, 6, 7, 8, 14, 15, 16
- [19] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357, 2021. 2, 4, 5, 6, 7, 8, 14, 16
- [20] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 568–578, 2021. 2, 6, 7, 8, 14, 16
- [21] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *CVPR*, 2022. 2, 6, 7, 8, 14, 15, 16
- [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 2, 8
- [23] Jeffrey L Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99, 1993. 2, 3
- [24] Kai A Krueger and Peter Dayan. Flexible shaping: How learning in small steps helps. *Cognition*, 110(3):380–394, 2009. 2, 3
- [25] M Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *NeurIPS*, 2010. 2, 3, 8
- [26] Emmanouil Antonios Platanios, Otilia Stretcu, Graham Neubig, Barnabás Póczos, and Tom M Mitchell. Competence-based curriculum learning for neural machine translation. In *NAACL-HLT (1)*, 2019. 2
- [27] Yizeng Han, Yifan Pu, Zihang Lai, Chaofei Wang, Shiji Song, Junfeng Cao, Wenhui Huang, Chao Deng, and Gao Huang. Learning to weight samples for dynamic early-exiting networks. In *ECCV*, pages 362–378, 2022. 2
- [28] Xinlei Chen and Abhinav Gupta. Webly supervised learning of convolutional networks. In *ICCV*, pages 1431–1439, 2015. 2
- [29] Yunchao Wei, Xiaodan Liang, Yunpeng Chen, Xiaohui Shen, Ming-Ming Cheng, Jiashi Feng, Yao Zhao, and Shuicheng Yan. Stc: A simple to complex framework for weakly-supervised semantic segmentation. *TPAMI*, 39(11):2314–2320, 2016. 2
- [30] Radu Tudor Ionescu, Bogdan Alexe, Marius Leordeanu, Marius Popescu, Dim P Papadopoulos, and Vittorio Ferrari. How hard can it be? estimating the difficulty of visual search in an image. In *CVPR*, pages 2157–2166, 2016. 2
- [31] Jonathan G Tullis and Aaron S Benjamin. On the effectiveness of self-paced learning. *Journal of memory and language*, 64(2):109–118, 2011. 2
- [32] Lu Jiang, Deyu Meng, Teruko Mitamura, and Alexander G Hauptmann. Easy samples first: Self-paced reranking for zero-example multimedia search. In *ACM MM*, pages 547–556, 2014. 2
- [33] Daphna Weinshall, Gad Cohen, and Dan Amir. Curriculum learning by transfer learning: Theory and experiments with deep networks. In *ICML*, pages 5238–5246, 2018. 2
- [34] Mengye Ren, Wenyan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pages 4334–4343, 2018. 2
- [35] Dingwen Zhang, Junwei Han, Long Zhao, and Deyu Meng. Leveraging prior-knowledge for weakly supervised object detection under a collaborative self-paced curriculum learning framework. *IJCV*, 127(4):363–380, 2019. 2
- [36] Tabet Matisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher–student curriculum learning. *IEEE transactions on neural networks and learning systems*, 31(9):3732–3740, 2019. 2
- [37] Samarth Sinha, Animesh Garg, and Hugo Larochelle. Curriculum by smoothing. In *NeurIPS*, pages 21653–21664, 2020. 2, 8
- [38] Pietro Morerio, Jacopo Cavazza, Riccardo Volpi, René Vidal, and Vittorio Murino. Curriculum dropout. In *ICCV*, pages 3544–3552, 2017. 2

- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2014. 2
- [40] Guangcong Wang, Xiaohua Xie, Jianhuang Lai, and Jiaxuan Zhuo. Deep growing learning. In *ICCV*, pages 2812–2820, 2017. 2
- [41] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018. 2
- [42] Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. In *ICLR*, 2015. 2
- [43] Tao Wei, Changhu Wang, Yong Rui, and Chang Wen Chen. Network morphism. In *ICML*, pages 564–572, 2016. 2
- [44] Tao Wei, Changhu Wang, and Chang Wen Chen. Modularized morphing of deep convolutional neural networks: A graph approach. *IEEE Transactions on Computers*, 70(2):305–315, 2020. 2
- [45] Le Yang, Haojun Jiang, Ruojin Cai, Yulin Wang, Shiji Song, Gao Huang, and Qi Tian. Condensenet v2: Sparse feature reactivation for deep networks. In *CVPR*, pages 3569–3578, 2021. 2
- [46] Linyuan Gong, Di He, Zhuohan Li, Tao Qin, Liwei Wang, and Tieyan Liu. Efficient training of bert by progressively stacking. In *ICML*, pages 2337–2346, 2019. 2
- [47] Minjia Zhang and Yuxiong He. Accelerating training of transformer-based language models with progressive layer dropping. In *NeurIPS*, pages 14011–14023, 2020. 2
- [48] Jiachun Wang, Fajie Yuan, Jian Chen, Qingyao Wu, Min Yang, Yang Sun, and Guoxiao Zhang. Stackrec: Efficient training of very deep sequential recommender models by iterative stacking. In *ACM SIGIR*, pages 357–366, 2021. 2
- [49] Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. L2-gcn: Layer-wise and learned efficient training of graph convolutional networks. In *CVPR*, pages 2127–2135, 2020. 2
- [50] Yulin Wang, Zanlin Ni, Shiji Song, Le Yang, and Gao Huang. Revisiting locally supervised learning: an alternative to end-to-end training. In *ICLR*, 2021. 2
- [51] Eugene Belilovsky, Michael Eickenberg, and Edouard Oyallon. Decoupled greedy learning of cnns. In *ICML*, pages 736–745, 2020. 2
- [52] Zanlin Ni, Yulin Wang, Jiangwei Yu, Haojun Jiang, Yue Cao, and Gao Huang. Deep incubation: Training large models by divide-and-conquering. *arXiv preprint arXiv:2212.04129*, 2022. 2
- [53] Mingxing Tan and Quoc Le. Efficientnetv2: Smaller models and faster training. In *ICML*, pages 10096–10106, 2021. 2, 8, 9
- [54] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. In *NeurIPS*, 2019. 2, 8
- [55] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *CVPR*, pages 8684–8694, 2020. 2, 3
- [56] Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier perspective on model robustness in computer vision. In *NeurIPS*, 2019. 3
- [57] Raphael Gontijo Lopes, Dong Yin, Ben Poole, Justin Gilmer, and Ekin D Cubuk. Improving robustness without sacrificing accuracy with patch gaussian augmentation. *arXiv preprint arXiv:1906.02611*, 2019. 3
- [58] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. *arXiv preprint arXiv:2105.07581*, 2(3), 2021. 3
- [59] Yoav Freund, Robert E Schapire, et al. Experiments with a new boosting algorithm. In *ICML*, pages 148–156, 1996. 3
- [60] Guillaume Alain, Alex Lamb, Chinnadhurai Sankar, Aaron Courville, and Yoshua Bengio. Variance reduction in sgd by distributed importance sampling. *arXiv preprint arXiv:1511.06481*, 2015. 3
- [61] Ilya Loshchilov and Frank Hutter. Online batch selection for faster training of neural networks. *arXiv preprint arXiv:1511.06343*, 2015. 3
- [62] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, pages 761–769, 2016. 3
- [63] Siddharth Gopal. Adaptive sampling for sgd by exploiting side information. In *ICML*, pages 364–372, 2016. 3
- [64] Tianyi Zhou, Shengjie Wang, and Jeffrey Bilmes. Curriculum learning by dynamic instance hardness. In *NeurIPS*, pages 8602–8613, 2020. 3, 8
- [65] Te Pi, Xi Li, Zhongfei Zhang, Deyu Meng, Fei Wu, Jun Xiao, and Yueting Zhuang. Self-paced boost learning for classification. In *IJCAI*, pages 1932–1938, 2016. 3
- [66] Stefan Braun, Daniel Neil, and Shih-Chii Liu. A curriculum learning method for improved noise robustness in automatic speech recognition. In *EUSIPCO*, pages 548–552, 2017. 3
- [67] Tianyi Zhou and Jeff Bilmes. Minimax curriculum learning: Machine teaching with desirable difficulties and scheduled diversity. In *ICLR*, 2018. 3, 8
- [68] Xuan Zhang, Gaurav Kumar, Huda Khayrallah, Kenton Murray, Jeremy Gwinnup, Marianna J Martindale, Paul McNamee, Kevin Duh, and Marine Carpuat. An empirical exploration of curriculum learning for neural machine translation. *arXiv preprint arXiv:1811.00739*, 2018. 3
- [69] Wei Wang, Isaac Caswell, and Ciprian Chelba. Dynamically composing domain-data selection with clean-data selection by “co-curricular learning” for neural machine translation. In *ACL*, pages 1282–1292, 2019. 3
- [70] Fengbei Liu, Yu Tian, Yuanhong Chen, Yuyuan Liu, Vasileios Belagiannis, and Gustavo Carneiro. Aclpl: Anti-curriculum pseudo-labelling for semi-supervised medical image classification. In *CVPR*, pages 20697–20706, 2022. 3

- [71] Fergus W Campbell and John G Robson. Application of fourier analysis to the visibility of gratings. *The Journal of physiology*, 197(3):551, 1968. 3
- [72] Wim Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM journal on mathematical analysis*, 29(2):511–546, 1998. 3
- [73] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999. 3
- [74] Yunpeng Chen, Haoqi Fan, Bing Xu, Zhicheng Yan, Yanis Kalantidis, Marcus Rohrbach, Shuicheng Yan, and Jiashi Feng. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In *ICCV*, pages 3435–3444, 2019. 3, 5
- [75] Le Yang, Yizeng Han, Xi Chen, Shiji Song, Jifeng Dai, and Gao Huang. Resolution adaptive networks for efficient inference. In *CVPR*, pages 2369–2378, 2020. 5
- [76] Gao Huang, Yulin Wang, Kangchen Lv, Haojun Jiang, Wenhui Huang, Pengfei Qi, and Shiji Song. Glance and focus networks for dynamic visual recognition. *TPAMI*, 45(4):4605–4621, 2023. 5
- [77] Yulin Wang, Kangchen Lv, Rui Huang, Shiji Song, Le Yang, and Gao Huang. Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. In *NeurIPS*, pages 2432–2444, 2020. 5
- [78] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. In *NeurIPS*, volume 34, pages 11960–11973, 2021. 5
- [79] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *JMLR*, 13:795–828, 2012. 5
- [80] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *ICML*, pages 3519–3529, 2019. 5
- [81] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, pages 113–123, 2019. 6
- [82] Xinyu Zhang, Qiang Wang, Jian Zhang, and Zhao Zhong. Adversarial autoaugment. In *ICLR*, 2019. 6
- [83] Yulin Wang, Xuran Pan, Shiji Song, Hong Zhang, Gao Huang, and Cheng Wu. Implicit semantic data augmentation for deep networks. In *NeurIPS*, volume 32, 2019. 6
- [84] Yulin Wang, Gao Huang, Shiji Song, Xuran Pan, Yitong Xia, and Cheng Wu. Regularizing deep networks with semantic data augmentation. *IEEE TPAMI*, 44(7):3733–3748, 2021. 6
- [85] Keyu Tian, Chen Lin, Ming Sun, Luping Zhou, Junjie Yan, and Wanli Ouyang. Improving auto-augment via augmentation-wise weight sharing. In *NeurIPS*, pages 19088–19098, 2020. 6
- [86] Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. Randaugment: Practical automated data augmentation with a reduced search space. In *NeurIPS*, pages 18613–18624, 2020. 6, 14
- [87] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014. 6
- [88] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6, 8
- [89] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729, 2008. 6, 8
- [90] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization: Stanford dogs. In *CVPRW*, 2011. 6, 8
- [91] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. 8
- [92] Ürün Dogan, Aniket Anand Deshmukh, Marcin Bronislaw Machura, and Christian Igel. Label-similarity curriculum learning. In *ECCV*, pages 174–190, 2020. 8
- [93] zhuofan xia, Xuran Pan, Xuan Jin, Yuan He, Hui Xue’, Shiji Song, and Gao Huang. Budgeted training for vision transformer. In *ICLR*, 2023. 8
- [94] Mengtian Li, Ersin Yumer, and Deva Ramanan. Budgeted training: Rethinking deep neural network training under resource constraints. In *ICLR*, 2020. 8
- [95] Hugo Touvron, Matthieu Cord, and Herve Jegou. Deit iii: Revenge of the vit. In *ECCV*, 2022. 8
- [96] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 8, 15
- [97] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: high quality object detection and instance segmentation. *TPAMI*, 43(5):1483–1498, 2019. 8, 15
- [98] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 14
- [99] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032, 2019. 14
- [100] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, pages 13001–13008, 2020. 14
- [101] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 14

- [102] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *ECCV*, pages 646–661, 2016. [14](#), [15](#)
- [103] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *ICCV*, pages 32–42, 2021. [14](#)
- [104] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992. [14](#)
- [105] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. In *ICLR*, 2018. [14](#)
- [106] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021. [15](#)
- [107] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *CVPR*, 2022. [15](#)
- [108] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *NeurIPS*, 2021. [15](#)
- [109] Yulin Wang, Zhaoxi Chen, Haojun Jiang, Shiji Song, Yizeng Han, and Gao Huang. Adaptive focus for efficient video recognition. In *ICCV*, pages 16249–16258, 2021. [20](#)
- [110] Yulin Wang, Yang Yue, Yuanze Lin, Haojun Jiang, Zihang Lai, Victor Kulikov, Nikita Orlov, Humphrey Shi, and Gao Huang. Adafocus v2: End-to-end training of spatial dynamic networks for video recognition. In *CVPR*, pages 20030–20040. IEEE, 2022. [20](#)
- [111] Yulin Wang, Yang Yue, Xinhong Xu, Ali Hassani, Victor Kulikov, Nikita Orlov, Shiji Song, Humphrey Shi, and Gao Huang. Adafocusv3: On unified spatial-temporal dynamic video recognition. In *ECCV 2022*, pages 226–243. Springer, 2022. [20](#)
- [112] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, volume 34, pages 13937–13949, 2021. [20](#)
- [113] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Yitian Zhang, and Haojun Jiang. Spatially adaptive feature refinement for efficient inference. *IEEE Transactions on Image Processing*, 30:9345–9358, 2021. [20](#)
- [114] Yizeng Han, Zhihang Yuan, Yifan Pu, Chenhao Xue, Shiji Song, Guangyu Sun, and Gao Huang. Latency-aware spatial-wise dynamic networks. *arXiv preprint arXiv:2210.06223*, 2022. [20](#)
- [115] Ziwei Zheng, Le Yang, Yulin Wang, Miao Zhang, Lijun He, Gao Huang, and Fan Li. Dynamic spatial focus for efficient compressed video action recognition. *IEEE TCSVT*, 2023. [20](#)
- [116] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. Dynamic neural networks: A survey. *IEEE TPAMI*, 44(11):7436–7456, 2021. [20](#)
- [117] Yifan Pu, Yiru Wang, Zhuofan Xia, Yizeng Han, Yulin Wang, Weihao Gan, Zidong Wang, Shiji Song, and Gao Huang. Adaptive rotated convolution for rotated object detection. *arXiv preprint arXiv:2303.07820*, 2023. [20](#)
- [118] Yizeng Han, Dongchen Han, Zeyu Liu, Yulin Wang, Xuran Pan, Yifan Pu, Chao Deng, Junlan Feng, Shiji Song, and Gao Huang. Dynamic perceiver for efficient visual recognition. *arXiv preprint arXiv:2306.11248*, 2023. [20](#)