

Preparing Lessons for Progressive Training on Language Models

Yu Pan^{1*}, Ye Yuan^{3,4*}, Yichun Yin⁵, Jiaxin Shi⁶, Zenglin Xu^{1,2†},
Ming Zhang^{3,4†}, Lifeng Shang⁵, Xin Jiang⁵, Qun Liu⁵

¹Harbin Institute of Technology Shenzhen, Shenzhen, Guangdong, China

²Pengcheng Laboratory, Shenzhen, China

³School of Computer Science, Peking University, Beijing, China

⁴Peking University-Anker Embodied AI Lab

⁵Huawei Noah's Ark Lab, Shenzhen, Guangdong, China

⁶Cloud BU, Huawei Technologies

Abstract

The rapid progress of Transformers in artificial intelligence has come at the cost of increased resource consumption and greenhouse gas emissions due to growing model sizes. Prior work suggests using pretrained small models to improve training efficiency, but this approach may not be suitable for new model structures. On the other hand, training from scratch can be slow, and progressively stacking layers often fails to achieve significant acceleration. To address these challenges, we propose a novel method called Apollo, which prepares lessons for expanding operations by learning high-layer functionality during training of low layers. Our approach involves low-value-prioritized sampling (LVPS) to train different depths and weight sharing to facilitate efficient expansion. We also introduce an interpolation method for stable model depth extension. Experiments demonstrate that Apollo achieves state-of-the-art acceleration ratios, even rivaling methods using pretrained models, making it a universal and efficient solution for training deep models while reducing time, financial, and environmental costs.

Introduction

Transformers (Vaswani et al. 2017) have recently achieved a significant impact on the field of artificial intelligence (Wang et al. 2020; Li et al. 2020; Dosovitskiy et al. 2021; Cao et al. 2021; Wang et al. 2023a). Nevertheless, the training cost is increasing in terms of the growing model size, which causes an amount of resourcing consumption and greenhouse gases emission (Schwartz et al. 2019; Pan et al. 2019). Addressing this problem, recent work (Chen et al. 2022; Chen, Goodfellow, and Shlens 2016; Wang et al. 2023b; Pan et al. 2023) suggests to improve training efficiency by reusing a pretrained small model as an initialization method that contains knowledge prior. However, the requirement of a pretrained model can cause fatal obstacles, especially for a

new-designed model structure, which affects the applicability of these studies as general strategies for training. On the other hand, the studies of training from scratch (Gong et al. 2019; Yang et al. 2020) generally stack layers to train Transformer progressively. Nevertheless, these methods usually cannot achieve significant acceleration in training, and are thus slower than training from pretrained models. Against this background, it is an emergency to design a universal method to efficiently train the models with reduced time and financial costs, while benefiting the ecological environment.

To achieve this goal, the progressive expansion of models in depth proves to be a crucial aspect in training from scratch. One noteworthy example is StackBERT (Gong et al. 2019), where a stacking learning strategy contributes to improved training efficiency through two merits: (1) Fewer layers in the initial stages of training require fewer computational resources, leading to faster training; (2) Lower trained weights provide a usable prior that benefits the training of stacked higher weights. While StackBERT undeniably accelerates training from the second merit, there are still two concerns that need to be addressed. Firstly, the suitability of the stacking method is questionable. For instance, directly stacking the 1-st layer onto the 7-th layer of a 12-layer Transformer is not intuitive due to the clear differences in semantic functionality between them (Rogers, Kovaleva, and Rumshisky 2020). Moreover, even though there might be some similarities across the entire model, it has been pointed out that most of the layers are different from each other (Chen et al. 2022). Consequently, it raises doubts about whether the normally trained weights are sufficiently prepared well to be expanded effectively, given the lack of knowledge in higher layers.

Motivated by this consideration, we introduce “Apollo” - a novel approach to preparing lessons for low-layer weights to learn the high-layer functionality in the training process. This strategy proves beneficial in further extending the capabilities of the model. In essence, Apollo involves two key components. Firstly, we employ a low-value-prioritized sampling (LVPS) technique, which randomly selects a depth for training at each step. This helps to ensure a diverse

*These authors contributed equally. This work was done when Yu Pan and Ye Yuan were interns at Huawei Noah's Ark Lab.

†Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

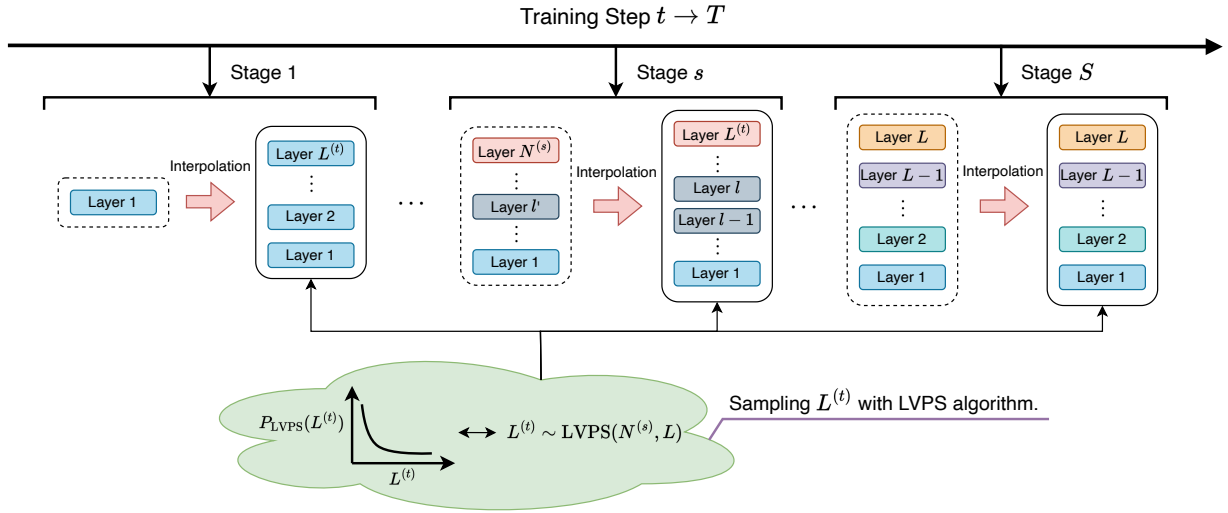


Figure 1: An illustration of the Apollo for training an L -layered model within T steps. We divide this training process into S stages. In the t -th step at the s -th stage, the model weights are $N^{(s)}$ layers (the left layers in each stage in the figure). To let the $N^{(s)}$ layers learn functionality in high layers in advance, we construct $L^{(t)}$ layers (the right layers in each stage in the figure) by sharing the $N^{(s)}$ weights through an interpolation method, where $N^{(s)} \leq L^{(t)}$. As shown in the figure, the same color denotes the same weight. We randomly choose $L^{(t)}$ at t -th step through a probability function Low-Value-Prioritized Sampling (LVPS). Since LVPS tends to select shallower layers, it can greatly save computation costs. Furthermore, we progressively increase the $N^{(s)}$ weights when stepping into the next stage. Since weights in the early stage can learn the properties of higher layers, Apollo can significantly contribute to the training efficiency.

training experience. Subsequently, we share the low-layer weights, enabling them to adapt to the selected layers by LVPS. These shared weights are well-prepared not only for learning high-layer functionality but also for recurrent transformation, as supported by previous work (Lan et al. 2020; Dehghani et al. 2019; Yang et al. 2021). It is worth noting that while weight-sharing strategies have been explored previously to facilitate information exchange across layers, our approach is novel in its dynamic application to sample layers, which is important to learn high-layer property to improve training efficiency. Furthermore, we address the issue of training stability by introducing an interpolation method to extend the depth of the model. This is essential since directly stacking layers can cause large gradients, which can be detrimental to training stability. In all, we summarize our contribution as:

- Through sharing weights in the early stage to learn the functionality of high layers, Apollo effectively expands the depth of networks, resulting in remarkable training acceleration.
- Through LVPS, Apollo achieves a substantial reduction in training FLOPs by predominantly sampling low depth layers, while retaining the benefit of expanding depth.
- Through replacing layer stacking with layer interpolation, Apollo further enhances the stability of the expanded model.
- Experiments show that Apollo attains state-of-the-art training efficiency, surpassing even the methods reliant on pretrained models.

Related Work

Efficient training from scratches. Training from scratch means training without any prior knowledge. Some approaches (Gong et al. 2019; Yang et al. 2020; Li et al. 2022; Gu et al. 2021; Shen et al. 2022) are known as “progressive training,” which involves initially pretraining a smaller scratch model and then gradually increasing its size, resulting in accelerated training. There are various training strategies for such models that are universally applicable and independent of our own method. For instance, employing optimization techniques like Adam (Kingma and Ba 2015) can accelerate the learning process by considering the optimizer’s perspective. Shoenybi et al. (2019) have successfully utilized mixed precision training to enhance training efficiency, while low-rank methods provide a viable option for memory and time-efficient training (Kamalakara et al. 2022). Additionally, Wu et al. (2021) have demonstrated improved data efficiency by taking note of rare words. Other effective training methods include dropping layers (Zhang and He 2020), knowledge inheritance (Qin et al. 2022), and merging tokens (Bolya et al. 2022). Our method embraces the concept of progressive training and attains noteworthy training efficiency by preparing instructive lessons to facilitate the expansion of layer depth.

Efficient training by reusing pretrained models. Recent studies have demonstrated the great potentials of developing large pre-trained models via expanding a small model. The pioneering work of Net2Net (Chen, Goodfellow, and Shlens 2016) introduced the concept of function-preserving

transformations which increase the width via neuron splitting and the depth via identity layers. However, the randomly chosen neurons for splitting in Net2Net are not promised to be well performed. Addressing this challenge, a series of studies (Wu, Wang, and Liu 2019; Wu et al. 2020b; Wang et al. 2019b; Wu et al. 2020a) have adopted functional steepest descent to select the optimal subset of neurons for splitting. Following the idea of function preservation, bert2BERT (Chen et al. 2022) expanded small transformers. More recently, LiGO (Wang et al. 2023b) introduced a trainable linear operator to learn an effective expansion formula. Mango (Pan et al. 2023) utilized a tensor ring matrix product operator (TR-MPO) to grow a small pretrained model to a large counterpart for efficient training. In the past, the utilization of pretrained methods consistently yielded greater acceleration than that achieved by training from scratch. Remarkably, our approach demonstrates highly promising results, surpassing even the swiftness of training from a pretrained model.

Method

In this section, we will introduce our method to sample the information of the higher layer to accelerate training.

Notations

To describe the implementation of Apollo, we introduce related notations here. We denote an L -layered network by $f^{(L)}(\cdot)$, and denote the function of the l -th layer of $f^{(L)}$ as $f_l(\cdot)$, where $l \in [L]$. Thus, given an input x , $f^{(L)}(x)$ can be formulated as

$$f^{(L)}(x) = f_L(f_{L-1}(\dots f_l(\dots f_1(x)))). \quad (1)$$

We denote the set of N weights in f as $\{\theta_i\}_{i=1}^N$, and the weights of l -th layer as $\Theta(f_l)$. As our method samples information of higher layers through weight sharing, here we formally consider a mapping function $g(\cdot)$ to arrange the $g(l)$ -th weight to the l -th layer as

$$\Theta(f_l) = \theta_{g(l)}, \quad (2)$$

where $g(\cdot) \in [N]$. Moreover, for the convenience of expressing the process of layer expansion, we denote the layer size of the t -th step as $L^{(t)}$, where $t \in T$ and T is the total steps. Thus, a network at the t -th step can be denoted by $f^{(L^{(t)})}$.

Transformer Architecture. Here, we introduce the structure of the Transformer, each block of which mainly contains two basic types of layers, i.e., the multi-head self-attention (MHSA) layer, and the feed-forward neural network (FFN) layer. Assuming inputs of the l -th layer are query $\mathbf{Q}_l \in \mathbb{R}^d$, key $\mathbf{K}_l \in \mathbb{R}^d$, and value $\mathbf{V}_l \in \mathbb{R}^d$, where d is the dimension, the l -th layer can be formulated as

$$f_l(\mathbf{Q}_l, \mathbf{K}_l, \mathbf{V}_l) = \text{FFN}(\text{MHSA}(\mathbf{Q}_l, \mathbf{K}_l, \mathbf{V}_l)). \quad (3)$$

The parameters of the MHSA layer of the l -th layer are denoted by $\mathbf{W}_l^{\{Q,K,V,O\}} \in \mathbb{R}^{d \times d}$. The weights of the FFN layer are $\mathbf{W}_l^{IN} \in \mathbb{R}^{d \times \alpha d}$ and $\mathbf{W}_l^{OUT} \in \mathbb{R}^{\alpha d \times d}$, where α is an expanding ratio that is often set to 4. Note that, we neglect biases in the formulation for simplification, which does

Algorithm 1: Process of Apollo

Require: the input data x , the ground-truth y , the stage setting $\{s_t, t \in [1, T], s_t \in [1, S]\}$: a non-decreasing list, indicating the stage of each step.

- 1: **for** $t = 1$ **to** T **do**
- 2: **if** $t > 1$ **and** $s_t > s_{t-1}$ **then**
- 3: **for** $n = 1$ **to** $N^{(s)}$ **do**
- 4: $\theta_n := \text{COPY} \left(\theta_{g_{\text{interpolation}}^{N^{(sprev)}:N^{(s)}}(n)} \right)$
- 5: **end for**
- 6: **end if**
- 7: $L^{(t)} = \text{LVPS}(N^{(s)}, L)$
- 8: **for** $l = 1$ **to** $L^{(t)}$ **do**
- 9: $\Theta(f_l) := \text{SHARE} \left(\theta_{g_{\text{interpolation}}^{N^{(s)}:L^{(t)}}(l)} \right)$
- 10: **end for**
- 11: $\mathcal{L} = \text{Loss} \left(f^{(L^{(t)})}(x), y \right)$
- 12: $\{\Delta\theta_i\}^{N^{(t)}} = \mathcal{L}. \text{backward}()$
- 13: Update all the weights $\{\theta_i\}^{N^{(s)}}$ through $\{\Delta\theta_i\}^{N^{(s)}}$
- 14: **end for**

Ensure: The trained model $f^{(L)}$ with $\{\theta_i\}^L$

not affect the generality of the proposed method. Following the above formulation, all weights of the l -th layer can be denoted as $\theta_l = \mathbf{W}_l^{\{Q,K,V,O,IN,OUT\}}$.

Efficient Training by Apollo

Given the potential similarities between high and low layers, the gradual expansion of layers during training can lead to enhanced acceleration, as opposed to the direct approach of training from scratch (Gong et al. 2019). However, as previously discussed, lower layers struggle to effectively capture the intricacies of higher-level features, particularly when utilizing a stacking methodology which can cause training instability. To address this challenge, we introduce Apollo, an innovative approach that facilitates progressive model training by leveraging weight sharing within the training process. This approach enables the accession of high-level functionality prior to layer expansion, thus mitigating the aforementioned issue.

Progressive Training. Specifically in an L -layer network, Apollo divides the whole training process into S stages. We use $N^{(s)}, s \in [S]$ to denote the weight number of the s -th stage, satisfying

$$N^{(s)} < N^{(s+1)}. \quad (4)$$

As restricted in Eq. (4), Apollo increases actual weights when stages go on. This progressive training way is efficient for accelerating language models.

Preparing Lessons. In order to provide lessons for $N^{(s)}$ weights within the s -th stage, thereby facilitating the pre-learning of higher-layer functionalities, Apollo employs a strategic weight-sharing approach. This involves distributing the $N^{(s)}$ weights across $L^{(t)}$ layers, drawn from the




Figure 2: A case of choosing hyper-parameters k of LVPS to sample 1-6 layer number.

range $[N^{(s)}, L]$, thereby establishing a connection to higher-layer elements. The determination of the appropriate $L^{(t)}$ is realized through the utilization of a sampling function denoted as $P(L^{(t)})$ at each training step. The core idea of this sampling function is the prioritization of shallower depths – specifically, $N^{(s)}$ in this context – to strike a balance between computational efficiency and sustained performance. Grounded in this consideration, we introduce a novel approach coined as Low-Value-Prioritized Sampling (LVPS).

In the development of LVPS, we employ an inverse proportional function as the cumulative distribution function for layer selection with a formulation as

$$F_{LVPS}(x) = c - \frac{b}{x+k}, c > 0, b > 0, x \in [N^{(s)}, L], \quad (5)$$

where b, k and c are the hyper-parameters. This function is congruent with our sampling objective, which favors the selection of shallower depths. Then, by setting $F_{LVPS}(N^{(s)}) = 0$ and $F_{LVPS}(L) = 1$. The probability density function P_{LVPS} can be solved as

$$P_{LVPS}(L^{(t)}) = \begin{cases} \frac{b}{(L^{(t)}+k)^2}, & \text{if } L^{(t)} \in [N^{(s)}, L], \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

$$\text{w.r.t } \int P_{LVPS}(L^{(t)}) dL^{(t)} = 1, \quad (7)$$

where b and c can be solved in terms of Eq. (7) as

$$b = \frac{(N^{(s)} + k) * (L + k)}{L - N^{(s)}}, \quad c = \frac{L + k}{L - N^{(s)}}. \quad (8)$$

Therefore, we only need to adjust k to derive different sampling settings as shown in Fig. 2. In this paper, we set $k = 0$ in every experiment to obtain the least computation complexity. We employ the notation $LVPS(\alpha, \beta)$ to signify the process of sampling a value within the interval $[\alpha, \beta]$. $P_{LVPS}(L^{(t)})$ is the probability function of $LVPS(N^{(s)}, L)$.

A comparison with other sampling methods can be found in Table 1 and Fig. 3. In detail, Uniform Sampling (US) samples the layers equally, thereby mitigating the potential layer bias, while Edge Sampling (ES) tends to sample layers




Figure 3: Comparison among US, FS, ES, and LVPS to sample 1-6 layer number.

Method	Probability Density Function
LVPS	$P_{LVPS}(L^{(t)}) = \frac{b}{(L^{(t)}+k)^2}$
ES	$P_{ES}(L^{(t)}) = \frac{1}{k} * \left(\frac{1}{L^{(t)}-N^{(s)}-b} + \frac{1}{L+b-L^{(t)}} \right)$
US	$P_{US}(L^{(t)}) = \frac{1}{L-N^{(s)}}$
FS	$P_{FS}(L) = 1$

Table 1: An overview of sampling methods: (1) Low-Value-Prioritized Sampling (LVPS), (2) Uniform Sampling (US), (3) Edge Sampling (ES), and (4) Full Sampling (FS). $L^{(t)}$ can only derive values in $[N^{(s)}, L]$. Since the integration of the probability density function is 1, b can be solved when k is determined. In this paper, we set $k = 0$ and $k = 10$ for LVPS and ES, respectively. FS always samples $L^{(t)} = L$.

in low and high positions to learn the high-layer information while maintaining efficiency. Alternatively, Full Sampling (FS) always samples the deepest depth, helping early-trained weights adequately acquire the functionality of each layer, which, however, demands significant computational resources due to its selection of the maximum layer number at each step. Among these sampling methods, LVPS can achieve the highest efficiency in progressive training.

Stack V.S. Interpolation

In the training process of Apollo, it is important to select a feasible method for expanding the layer for sharing in each step and initializing weights of the next stage. As shown in Fig. 4, the common-used methods are stacking and interpolating layers. Given a target to expand L_1 to L_2 , the stacking method can be formed as

$$g_{\text{stack}}^{L_1:L_2}(l_2) = l_2 \bmod L_1, \quad (9)$$

where $l_2 \in [L_2]$ is the index of L_2 . The stacking method is usually adopted in language models, e.g., BERT (Gong et al. 2019). By contrast, the interpolation method is often used in the computer vision field, e.g., ResNet (Chang et al. 2018),




Figure 4: A case of expanding 3 layers to 6 layers. The same color denotes the same weight. The stacking method recurrently arranges the layers, e.g., the 1-st layer \rightarrow the 4-th layer. By contrast, the interpolation method arranges the layers in a neighbor, e.g., the 1-st layer \rightarrow the 2-nd layer.

and can be formulated as

$$g_{\text{interpolation}}^{L_1:L_2}(l_2) = \lfloor \frac{l_2 * L_1}{L_2} \rfloor, \quad (10)$$

where $*$ means dot production, and $\lfloor \cdot \rfloor$ denotes the rounding operation. Although the two expanding methods have both shown good performance in their fields, there is still a lack of analysis on the comparison between them, especially in the applicability of language models. In this paper, we investigate the influence of the stability and performance of these methods. We defer an analysis experiment to the experiment section. Here, we would like to give the conclusion: (1) There is only a small performance gap between them; (2) interpolation can reach better stability than the stacking method. As a result, we adopt interpolation instead of the stacking method as the expanding method of the proposed Apollo training for language models.

Experiment

In this section, we conduct experiments to validate the performance of our proposed method.

Common setting: We implement a series of experiments on BERT and GPT for validation. We use AdamW as the optimizer with a learning rate of 10^{-4} and weight decay of 10^{-2} in all the experiments. We chose the training batch sizes of 768 and 512 for BERT (Devlin et al. 2019) and GPT (Radford et al. 2019) models, respectively. We use the Scratch model, StackBERT, bert2BERT, and LiGO as baselines. Layer numbers of Apollo are [1, 3, 6, 12] and change at epoch [2, 4, 10]. LiGO is warmly trained for 100 steps as claimed in the original paper (Wang et al. 2023b). The training dataset is a concatenation of English Wikipedia and Toronto Book Corpus (Zhu et al. 2015).

Experiment on Expanding Method

We implement the experiment to show the influence of stacking and interpolating layers. We use BERT as the backbone. We train a 6-layer BERT model called BERT-Base/2 which is the half of BERT-Base for 10 epochs with a 768 batch size. Then, we expand the trained BERT-Base/2 into the 12-layer BERT-Base through the stacking (Eq. (9)) and




Figure 5: Distribution of output activations. BERT-Base/2 is half of BERT-Base. BERT-Base/2-S and BERT-Base/2-I denote to stack and interpolate BERT-Base/2 to BERT-Base, respectively. After stacking BERT-Base/2, the distribution of output activations changes a lot, while the interpolation method keeps the distribution well.

Model	Trainable	Layer	Loss	Gradient (1e-3)
BERT-Base Ra.	-	12	10.56	11.45 \pm 26.46
BERT-Base/2	-	6	1.82	1.66 \pm 4.71
BERT-Base/2-S	\times	12	7.32	55.28 \pm 241.52
BERT-Base/2-I	\checkmark	12	3.97	44.01 \pm 225.76

Table 2: Results of the analysis of expanding methods. BERT-Base Ra. means a randomly initialized BERT-Base. When expanding BERT-Base/2 from 6 to 12, the loss and gradients of the stacking method rise sharply, which causes failure in further training. By contrast, the interpolation method achieves smaller gradients and loss and is thus trainable later. However, both of the two methods cause higher gradients than a random model, indicating an unstable state.

interpolating (Eq. (10)) methods. We use 500 data samples for validation. Moreover, we apply Apollo to train BERT-Base with stacking and interpolation methods.

Illustrated in Fig.5, directly stacking BERT-Base/2 exhibits a notable alteration in the distribution of output activations, whereas the interpolation method preserves this distribution. This distinction results in interpolation yielding smaller losses and gradients in comparison to the stacking technique. This advantage contributes to the enhanced trainability of BERT-Base-I, as evidenced in Table2. Additionally, in Table 3, where expanding BERT-Base/2 increases both loss and gradients, Apollo showcases a reverse trend. Here, layer expansion within Apollo leads to a reduction in the loss and gradient for the lessons on the functionality of 12 layers, ensuring robust stability in progressive learning. Notably, interpolation further diminishes gradient values and improves the acceleration ratio, measured in terms of FLOPs. Therefore, we adopt interpolation as the preferred expansion method for Apollo considering these findings.




Figure 6: Results of BERT-Base and GPT-Base. Apollo achieves the highest acceleration on BERT-Base and GPT-Base in terms of FLOPs at 41.6% and 47.9%, respectively. In addition, Apollo can keep the best training efficiency on wall time for BERT-Base at 41.1%. Apollo surpasses methods (i.e., bert2BERT and LiGO) relying on pretrained models in all cases.

Model	Acc. Ratio	Layer	Loss	Gradient (1e-3)
Apollo-S	39.7%	6	1.82	1.65 ± 3.72
		12	1.76	1.42 ± 3.50
Apollo-I	41.6%	6	1.82	1.59 ± 3.33
		12	1.76	1.39 ± 3.01

Table 3: Results of expanding methods for Apollo. “S” and “I” mean the stacking and the interpolation methods, respectively. Apollo can decrease the gradient values loss after expanding layers since learning the functionality of a 12-layer network. Moreover, the interpolation method can derive a comparably better acceleration ratio in terms of FLOPs and smaller gradient values.

Experiment on Sampling Method

We construct the experiment on BERT-Base to investigate the influence of the sampling methods including ES, US, FS, and LVPS, with a comparison to the w/o sampling case.

Results are shown in Fig. 7. The LVPS sampling method leads with the highest acceleration ratio, attaining an impressive 41.6%. Both ES and LVPS attach the top two positions for acceleration, underscoring the utility of sampling lower layers. Significantly, LVPS outperforms Apollo w/o sampling by a substantial margin of +9.3%, affirming the pronounced benefits of extracting high-layer functionalities through sampling. US performing normally may indicate that sampling middle layers is not really helpful. FS demonstrates the least acceleration due to its consistent sampling of the maximum layers, incurring a substantial FLOP cost. In summation, LVPS strategically harnesses the ability to capture high-layer functionalities while simultaneously minimizing FLOP usage by primarily sampling lower layers.

Experiment on BERT

We construct this experiment to show the training efficiency of BERT. We train Apollo and StackBERT from scratch,




Figure 7: Results of the analysis on sampling methods. LVPS achieves the highest acceleration ratio at 41.6%, surpassing ES and FS at 4.1% and +17.1%, respectively. Most sampling methods are faster than Apollo w/o Sampling. FS performs the lowest acceleration by always sampling the deepest depth, which is resource-consuming.

while training bert2BERT and LiGO from BERT-Small to BERT-Base for 40 epochs. Additionally, we also implement short training on BERT-Large for 5 epochs.

As shown in Fig. 6, Apollo achieves the highest acceleration rate of 41.6% in FLOPs saving, which outperforms bert2BERT by an additional +6.0%. Through training from a pretrained BERT-Small, bert2BERT and LiGO exhibit superior performance compared to the progressive training approach StackBERT, resulting in acceleration ratios of 29.5% with corresponding improvements of +6.1% and +4.0%, respectively. However, the convergence speed of both bert2BERT and LiGO is found to be slower in comparison to Apollo, which shows the accelerated ability of the Apollo method through the “lessons” of higher layers. To substantiate the effectiveness of Apollo, comprehensive evaluations were conducted using the SQuAD and GLUE

Model	Saving (FLOPs)	Saving (Wall Time)	SQuADv1.1 (F1)	SQuADv2.0 (F1)	SST-2 (Acc)	MNLI (Acc)	MRPC (Acc)	COLA (Mcc)	QNLI (Acc)	STS-B (Acc)	QQP (Acc)	GLUE Avg.	SQuAD Avg.
Scratch	-	-	89.05	77.49	92.04	84.05	87.65	56.95	91.39	89.16	91.17	84.63	83.27
Training from the Pretrained Model: BERT-Small→BERT-Base													
bert2BERT	35.6%	35.2%	90.02	78.99	92.89	84.92	86.91	60.32	91.81	88.11	90.72	85.10	84.50
LiGO	33.5%	33.2%	90.09	78.34	92.75	84.99	87.44	61.10	91.33	87.94	90.42	85.14	84.22
Progressive Training form Scratch													
StackBERT	29.5%	28.9%	89.82	78.21	92.94	84.63	87.65	61.61	90.95	87.13	90.20	85.01	84.01
Apollo	41.6%	41.1%	89.87	78.42	92.28	84.81	87.06	60.57	91.43	88.27	90.69	85.02	84.15

Table 4: Experiments on downstream tasks of BERT-Base on GLUE (Wang et al. 2019a), SQuADv1.1 (Rajpurkar et al. 2016), and SQuADv2.0 (Rajpurkar, Jia, and Liang 2018) dataset. The terms “Training from the Pretrained Model” and “Progressive Training from Scratch” denote the pretraining type of the methods in the table. Compared with baselines, Apollo can achieve the highest FLOPs saving under similar downstream performance, even better than training from the pretrained model.




Figure 8: Results of BERT-Large.

benchmark datasets, as illustrated in Table 4. For a larger BERT-Large in Fig. 8, Apollo also shows a significant acceleration ratio at 55.1%. The results clearly demonstrate the ability of Apollo to exhibit proficient transfer learning capabilities, coupled with faster convergence rates, thus making it promising for practical applications.

Experiment on GPT

We also implement the experiment to validate the performance of GPT. We train Apollo, StackBERT and Scratch from a random initialization, and bert2BERT and LiGO from GPT-Small. The training epoch is 35.

As shown in Fig. 6(c), we construct a comparison among the Scratch model, StackBERT, bert2BERT, LiGO, and Apollo. Notably, our proposed Apollo method exhibits a remarkable 47.9% acceleration ratio. Despite the structural disparities between GPT and BERT, encompassing distinct features such as the mask method, and the position of layer normalization, Apollo consistently maintains the highest performance across the evaluated models. Specifically, Apollo shows a significant advancement over Stack-

BERT, achieving an improvement of +11.9% for the prepared lessons before expanding. Furthermore, compared to bert2BERT and LiGO, Apollo attains a considerable acceleration advantage with +6.8% and +9.9% higher performance, respectively. These findings highlight the substantial acceleration capabilities of the Apollo method, affirming its effectiveness for broader applications.

Conclusion

Training language models imposes a substantial demand on computational resources. Compared to training from pre-trained models, progressive training from scratch offers a universal and flexible solution to accelerate training, as it does not require a pretrained model. However, previous progressive methods suffered from inefficient layer expansion, leading to suboptimal training efficiency. In light of these challenges, we propose Apollo, a novel approach that imparts lessons for weights in the early stage, achieving a more natural and efficient expansion. Additionally, we conduct a thorough analysis of the influence between stacking and interpolating methods for expanding model depth, advocating the use of interpolation for improved stability in progressive training. Experimental results consistently demonstrate that Apollo achieves remarkable acceleration across various language models, even surpassing methods that rely on pre-trained models. This implies the significant effectiveness of Apollo in enhancing training efficiency. In the future, we expect that Apollo will contribute to the realization of green AI by significantly reducing the cost of training Transformers.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (No. 62276002), a key program of fundamental research from Shenzhen Science and Technology Innovation Commission (No. JCYJ20200109113403826), the Major Key Project of PCL (No. 2022ZD0115301), and an Open Research Project of Zhejiang Lab (NO.2022RC0AB04).

References

- Bolya, D.; Fu, C.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2022. Token Merging: Your ViT But Faster. *CoRR*, abs/2210.09461.
- Cao, H.; Wang, Y.; Chen, J.; Jiang, D.; Zhang, X.; Tian, Q.; and Wang, M. 2021. Swin-Unet: Unet-like Pure Transformer for Medical Image Segmentation. *CoRR*, abs/2105.05537.
- Chang, B.; Meng, L.; Haber, E.; Tung, F.; and Begert, D. 2018. Multi-level Residual Networks from Dynamical Systems View. In *ICLR (Poster)*. OpenReview.net.
- Chen, C.; Yin, Y.; Shang, L.; Jiang, X.; Qin, Y.; Wang, F.; Wang, Z.; Chen, X.; Liu, Z.; and Liu, Q. 2022. bert2BERT: Towards Reusable Pretrained Language Models. In *ACL (1)*, 2134–2148. Association for Computational Linguistics.
- Chen, T.; Goodfellow, I. J.; and Shlens, J. 2016. Net2Net: Accelerating Learning via Knowledge Transfer. In *ICLR*.
- Dehghani, M.; Gouws, S.; Vinyals, O.; Uszkoreit, J.; and Kaiser, L. 2019. Universal Transformers. In *ICLR (Poster)*. OpenReview.net.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT (1)*, 4171–4186. Association for Computational Linguistics.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*. OpenReview.net.
- Gong, L.; He, D.; Li, Z.; Qin, T.; Wang, L.; and Liu, T. 2019. Efficient Training of BERT by Progressively Stacking. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, 2337–2346. PMLR.
- Gu, X.; Liu, L.; Yu, H.; Li, J.; Chen, C.; and Han, J. 2021. On the Transformer Growth for Progressive BERT Training. In *NAACL-HLT*, 5174–5180. Association for Computational Linguistics.
- Kamalakara, S. R.; Locatelli, A.; Venkitesh, B.; Ba, J.; Gal, Y.; and Gomez, A. N. 2022. Exploring Low Rank Training of Deep Neural Networks. *CoRR*, abs/2209.13569.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR (Poster)*.
- Lan, Z.; Chen, M.; Goodman, S.; Gimpel, K.; Sharma, P.; and Soricut, R. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *ICLR*. OpenReview.net.
- Li, C.; Zhuang, B.; Wang, G.; Liang, X.; Chang, X.; and Yang, Y. 2022. Automated Progressive Learning for Efficient Training of Vision Transformers. In *CVPR*, 12476–12486. IEEE.
- Li, N.; Pan, Y.; Chen, Y.; Ding, Z.; Zhao, D.; and Xu, Z. 2020. Heuristic Rank Selection with Progressively Searching Tensor Ring Network. *CoRR*, abs/2009.10580.
- Pan, Y.; Xu, J.; Wang, M.; Ye, J.; Wang, F.; Bai, K.; and Xu, Z. 2019. Compressing Recurrent Neural Networks with Tensor Ring for Action Recognition. In *AAAI*, 4683–4690. AAAI Press.
- Pan, Y.; Yuan, Y.; Yin, Y.; Xu, Z.; Shang, L.; Jiang, X.; and Liu, Q. 2023. Reusing Pretrained Models by Multi-linear Operators for Efficient Training. In *NeurIPS*.
- Qin, Y.; Lin, Y.; Yi, J.; Zhang, J.; Han, X.; Zhang, Z.; Su, Y.; Liu, Z.; Li, P.; Sun, M.; and Zhou, J. 2022. Knowledge Inheritance for Pre-trained Language Models. In *NAACL-HLT*, 3921–3937. Association for Computational Linguistics.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rajpurkar, P.; Jia, R.; and Liang, P. 2018. Know What You Don’t Know: Unanswerable Questions for SQuAD. In *ACL (2)*, 784–789. Association for Computational Linguistics.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100, 000+ Questions for Machine Comprehension of Text. In *EMNLP*, 2383–2392. The Association for Computational Linguistics.
- Rogers, A.; Kovaleva, O.; and Rumshisky, A. 2020. A Primer in BERTology: What We Know About How BERT Works. *Trans. Assoc. Comput. Linguistics*, 8: 842–866.
- Schwartz, R.; Dodge, J.; Smith, N. A.; and Etzioni, O. 2019. Green AI. *CoRR*, abs/1907.10597.
- Shen, S.; Walsh, P.; Keutzer, K.; Dodge, J.; Peters, M. E.; and Beltagy, I. 2022. Staged Training for Transformer Language Models. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, 19893–19908. PMLR.
- Shoeybi, M.; Patwary, M.; Puri, R.; LeGresley, P.; Casper, J.; and Catanzaro, B. 2019. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism. *CoRR*, abs/1909.08053.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NIPS*, 5998–6008.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019a. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *ICLR (Poster)*. OpenReview.net.
- Wang, D.; Li, M.; Wu, L.; Chandra, V.; and Liu, Q. 2019b. Energy-Aware Neural Architecture Optimization with Fast Splitting Steepest Descent. *CoRR*, abs/1910.03103.
- Wang, M.; Pan, Y.; Yang, X.; Li, G.; Xu, Z.; and Cichocki, A. 2023a. Tensor Networks Meet Neural Networks: A Survey and Future Perspectives. *CoRR*, abs/2302.09019.
- Wang, M.; Su, Z.; Luo, X.; Pan, Y.; Zheng, S.; and Xu, Z. 2020. Concatenated Tensor Networks for Deep Multi-Task Learning. In *ICONIP (5)*, volume 1333 of *Communications in Computer and Information Science*, 517–525. Springer.
- Wang, P.; Panda, R.; Hennigen, L. T.; Greengard, P.; Karlinsky, L.; Feris, R.; Cox, D. D.; Wang, Z.; and Kim, Y. 2023b. Learning to Grow Pretrained Models for Efficient Transformer Training. In *ICLR*. OpenReview.net.
- Wu, L.; Liu, B.; Stone, P.; and Liu, Q. 2020a. Firefly Neural Architecture Descent: a General Approach for Growing Neural Networks. In *NeurIPS*.

- Wu, L.; Wang, D.; and Liu, Q. 2019. Splitting Steepest Descent for Growing Neural Architectures. In *NeurIPS*, 10655–10665.
- Wu, L.; Ye, M.; Lei, Q.; Lee, J. D.; and Liu, Q. 2020b. Steepest Descent Neural Architecture Optimization: Escaping Local Optimum with Signed Neural Splitting. *CoRR*, abs/2003.10392.
- Wu, Q.; Xing, C.; Li, Y.; Ke, G.; He, D.; and Liu, T. 2021. Taking Notes on the Fly Helps Language Pre-Training. In *ICLR*. OpenReview.net.
- Yang, C.; Wang, S.; Yang, C.; Li, Y.; He, R.; and Zhang, J. 2020. Progressively Stacking 2.0: A Multi-stage Layer-wise Training Method for BERT Training Speedup. *CoRR*, abs/2011.13635.
- Yang, S.; Hou, L.; Song, X.; Liu, Q.; and Zhou, D. 2021. Speeding up Deep Model Training by Sharing Weights and Then Unsharing. *CoRR*, abs/2110.03848.
- Zhang, M.; and He, Y. 2020. Accelerating Training of Transformer-Based Language Models with Progressive Layer Dropping. In *NeurIPS*.
- Zhu, Y.; Kiros, R.; Zemel, R. S.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books. In *ICCV*, 19–27. IEEE Computer Society.