

# LV-BERT: Exploiting Layer Variety for BERT

**Weihao Yu**

National University  
of Singapore  
weihaoyu6@gmail.com

**Zihang Jiang**

National University  
of Singapore  
jzihang@u.nus.edu

**Fei Chen**

Huawei Noah’s  
Ark Lab  
chen.f@huawei.com

**Qibin Hou**

National University  
of Singapore  
andrewhoux@gmail.com

**Jiashi Feng**

National University  
of Singapore  
elefjia@nus.edu.sg

## Abstract

Modern pre-trained language models are mostly built upon backbones stacking self-attention and feed-forward layers in an interleaved order. In this paper, beyond this stereotyped layer pattern, we aim to improve pre-trained models by exploiting *layer variety* from two aspects: the layer type set and the layer order. Specifically, besides the original self-attention and feed-forward layers, we introduce convolution into the layer type set, which is experimentally found beneficial to pre-trained models. Furthermore, beyond the original interleaved order, we explore more layer orders to discover more powerful architectures. However, the introduced layer variety leads to a large architecture space of more than billions of candidates, while training a single candidate model from scratch already requires huge computation cost, making it not affordable to search such a space by directly training large amounts of candidate models. To solve this problem, we first pre-train a supernet from which the weights of all candidate models can be inherited, and then adopt an evolutionary algorithm guided by pre-training accuracy to find the optimal architecture. Extensive experiments show that LV-BERT model obtained by our method outperforms BERT and its variants on various downstream tasks. For example, LV-BERT-small achieves 78.8 on the GLUE testing set, 1.8 higher than the strong baseline ELECTRA-small.<sup>1</sup>

## 1 Introduction

In recent years, pre-trained language models, such as the representative BERT (Devlin et al., 2019) and GPT-3 (Brown et al., 2020), have gained great success in natural language processing tasks (Peters et al., 2018a; Radford et al., 2018; Yang et al., 2019; Clark et al., 2020). The backbone architectures of these models mostly adopt a stereotyped

<sup>1</sup><https://github.com/yuweihao/LV-BERT>

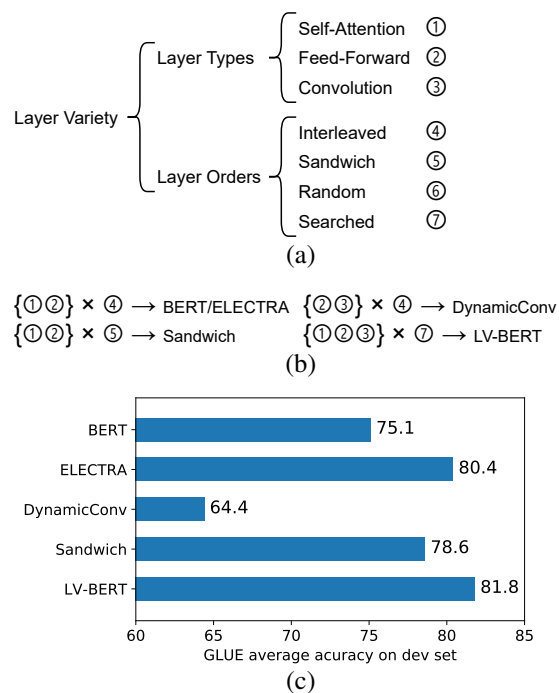


Figure 1: (a) Illustration of layer variety. This concept consists of two aspects: layer type and layer order. (b) Different models represented by layer variety. (c) Performance of different models with hidden size of 256 on GLUE (Wang et al., 2018) development set. Except BERT pre-trained with the Masked Language Modeling objective (Devlin et al., 2019), the other models are pre-trained with Replaced Token Detection objective (Clark et al., 2020) to save computation cost.

layer pattern, in which the *self-attention* and *feed-forward* layers are arrayed in an interleaved order (Vaswani et al., 2017). However, there is no evidence supporting that this layer pattern is optimal (Press et al., 2020). We then consider a straightforward and interesting question: Could we change the layer pattern to improve pre-trained models? We attempt to answer this question by exploiting more *layer variety* from two aspects, as shown in Figure 1(a): the layer type set and the layer order.

We first consider the layer types. In previous pre-trained language models, the most widely-used layer set contains the self-attention layer for capturing global information and the feed-forward layer for non-linear transformation. However, some recent works have unveiled that some self-attention heads in pre-trained models tend to learn local dependencies due to the inherent property of natural language (Kovaleva et al., 2019; Brunner et al., 2020; Jiang et al., 2020), incurring computation redundancy for capturing local information. In contrast, convolution is a local operator (LeCun et al., 1998; Krizhevsky et al., 2012; Simonyan and Zisserman, 2015; He et al., 2016) and has shown effectiveness on extracting local information for language models (Zeng et al., 2014; Kim, 2014; Kalchbrenner et al., 2014; Wu et al., 2018, 2019b; Jiang et al., 2020). Thus, we propose to augment the layer set by including convolution for local information extraction.

For layer orders, most of the existing pre-trained models adopt an interleaved order to arrange the different types of layers. Differently, Press et al. (2020) presented the sandwich order, *i.e.*, stacking consecutive self-attention and feed-forward layers at the bottom and top, respectively, while keeping the interleaved order in the middle. It has been shown that the sandwich order can bring improvement on language modeling task, indicating the layer order contributes to model performance. However, Press et al. (2020) did not show the generalization capability of this order to other tasks. There is still a large room for exploring more effective orders for pre-trained models. We show the different layer variety designs of existing models in Figure 1(b), including BERT (Devlin et al., 2019)/ELECTRA (Clark et al., 2020), Dynamic-Conv (Wu et al., 2018) and Sandwich (Press et al., 2020). Their performance is summarized in Figure 1(c). It can be seen that layer variety significantly influences model performance. We thus claim it is necessary to investigate layer variety for promoting pre-trained models. However, to perform such investigation for a common model backbone, *e.g.*, with 24 layers, we need to evaluate performance of every candidate within an architecture space of  $3^{24} \approx 2.8 \times 10^{11}$  candidates. Pre-training a single language model already needs to consume a large amount of computation, *e.g.*, 2400 P100 GPU days for pre-training BERT (Lin et al., 2020). It is barely affordable to pre-train such a large amount

of model candidates from scratch. To reduce the computation cost, inspired by recent works on Neural Architecture Search (NAS) (Guo et al., 2020; Cai et al., 2019), we construct a supernet according to the layer variety discussed above and pre-train it with Masked Language Modeling (MLM) (Devlin et al., 2019) objective. After obtaining the pre-trained supernet, we develop an evolutionary algorithm guided by MLM evaluation accuracy to search an effective architecture with specific layer variety. We call the resulted model LV-BERT. Extensive experiments show that LV-BERT outperforms BERT and its variants. The contributions of our paper are two-fold. Firstly, to the best of our knowledge, this work is the first to exploit layer variety w.r.t. both layer types and orders for pre-trained language models. We found convolutions and layer orders both benefit pre-trained model performance. We hope our observations would facilitate the development of pre-trained language models. Secondly, our obtained LV-BERT shows superiority over BERT and its variants. For example, LV-BERT-small achieves 79.8 on GLUE testing set, 1.8 higher than the baseline ELECTRA-small (Clark et al., 2020).

## 2 Related Work

**Pre-trained Language Models** Pre-trained language models have achieved great success and promoted the development of NLP techniques. Instead of separate word representation (Mikolov et al., 2013a,b), McCann et al. (2017) and Peters et al. (2018b) propose CoVe and ELMo respectively which both utilize LSTM (Hochreiter and Schmidhuber, 1997) to generate contextualized word representations. Later, Radford et al. (2018) introduce GPT that changes the backbone to transformers where self-attention and feed-forward layers are arrayed interleavedly. They also propose generative pre-training objectives. BERT (Devlin et al., 2019) continues to use the same layer set and order for backbone but employs different pre-training objectives, *i.e.*, Masked Language Modeling and Next Sentence Prediction. Then more works introduce new effective pre-training objectives, like Generalized Autoregressive Pretraining (Yang et al., 2019), Span Boundary Objective (Joshi et al., 2020) and Replaced Token Detection (Clark et al., 2020). Besides designing pre-training objectives, some other works try to extend BERT by incorporating knowledge (Zhang et al., 2019; Peters et al., 2019; Liu

et al., 2020; Xiong et al., 2020) or with multiple languages (Huang et al., 2019; Conneau and Lample, 2019; Chi et al., 2019). These works utilize the stereotyped layer pattern, which is unnecessarily optimal (Press et al., 2020), inspiring us to further investigate more layer variety to improve pre-trained models. To the best of our knowledge, we are the first to exploit layer variety from both the layer type set and the layer order for pre-trained language models.

**Neural Architecture Search** Manually designing neural architecture is a time-consuming and error-prone process (Elsken et al., 2019). To solve this, many neural architecture search algorithms are proposed. Pioneering works utilize reinforcement learning (Zoph and Le, 2017; Baker et al., 2017) or evolutionary algorithm (Real et al., 2017) to sample architecture candidates and train them from scratch, which demand huge computation that ordinary researchers can not afford. To reduce computation cost, recent methods (Pham et al., 2018; Liu et al., 2018; Xie et al., 2018; Brock et al., 2018; Cai et al., 2018; Bender et al., 2018; Wu et al., 2019a; Guo et al., 2020) adopt a *weight sharing* strategy that a supernet subsuming all architectures is trained only once and all architecture candidates can inherit their weights from the supernet. Despite the boom of NAS research, most works focus on computer vision tasks (Chen et al., 2019; Ghiasi et al., 2019; Liu et al., 2019a), while NAS on NLP is not fully investigated. Recently, So et al. (2019) and Wang et al. (2020) search architectures of transformers for translation tasks. Chen et al. (2020) leverage differentiable neural architecture to automatically compress BERT with task-oriented knowledge distillation for specific tasks. Zhu et al. (2020) utilize architecture search to improve models based on pre-trained BERT for the relation classification task. However, these methods only focus on specific tasks or the fine-tuning phase. Besides, Khetan and Karnin (2020) employ pre-training loss to help prune BERT, but their method can not find new architectures. Different from them, our work is the first to use NAS to help explore new architectures in a pre-training scenario for general language understanding.

### 3 Method

An overview of our approach is shown in Figure 2. We first define the layer variety to introduce a large architecture search space, and then pre-train

a supernet subsuming all candidate architectures, followed by an evolutionary algorithm guided by pre-training MLM (Devlin et al., 2019) accuracy to search an effective model. In what follows, we will give detailed descriptions.

#### 3.1 Layer Variety

As shown in Figure 1(a), the proposed layer variety contains two aspects: layer type and layer order, both of which are important for the performance of pre-trained models but not exploited before.

**Layer Type** The layer type set of current BERT-like models consists of self-attention for information communication and feed-forward for non-linear transformation. However, as a global operator, self-attention needs to take as input all tokens to compute attention weights for each token, which is inefficient in capturing local information (Wu et al., 2019b; Jiang et al., 2020). We notice that convolution (LeCun et al., 1998; Krizhevsky et al., 2012), as a local operator, has been successfully applied in language models (Zeng et al., 2014; Kim, 2014; Kalchbrenner et al., 2014; Wu et al., 2018, 2019b; Jiang et al., 2020). A typical example is the dynamic convolution (Wu et al., 2018) for machine translation, language modeling and summarization. Therefore, we augment the layer type set by introducing dynamic convolution as a new layer type. The layer set considered in this work thus contains three types of layers,

$$\mathcal{L}^{\text{type}} = \{L^{\text{SA}}, L^{\text{FF}}, L^{\text{DC}}\}, \quad (1)$$

where the set elements denote self-attention, feed-forward and dynamic convolution layers respectively. See Appendix for more detailed formulation description on them.

**Layer Order** The other variety aspect is layer order. The most widely-used order for pre-trained models is the interleaved order (Vaswani et al., 2017; Devlin et al., 2019). For a model with 24 layers, the interleaved order can be expressed by the following list,

$$[L_1^{\text{SA}}, L_2^{\text{FF}}, L_3^{\text{SA}}, L_4^{\text{FF}}, \dots, L_{23}^{\text{SA}}, L_{24}^{\text{FF}}]. \quad (2)$$

Similarly, the sandwich order (Press et al., 2020) can be expressed as

$$\begin{aligned} & [L_1^{\text{SA}}, L_2^{\text{SA}}, \dots, L_5^{\text{SA}}, \\ & L_6^{\text{SA}}, L_7^{\text{FF}}, L_8^{\text{SA}}, L_9^{\text{FF}}, \dots, L_{18}^{\text{SA}}, L_{19}^{\text{FF}}, \\ & L_{20}^{\text{FF}}, L_{21}^{\text{FF}}, \dots, L_{24}^{\text{FF}}]. \end{aligned} \quad (3)$$




Figure 2: Overview on how to search LV-BERT. ① Construct a supernet with small hidden size by including all types of layers at each layer. ② Pre-train the supernet with Masked Language Modeling (MLM) objective (Devlin et al., 2019) by only uniformly sampling one type of layer into training at each layer. ③ Apply evolutionary algorithm to produce candidate models. ④ The candidate models inherit their weights from the supernet. ⑤ The candidate models with inherited weights are directly evaluated with pre-training MLM accuracy on validation set. ⑥ The accuracy is used to guide the evolutionary algorithm for generating new candidate models. ⑦ After  $T$  iterations, the candidate with best pre-training accuracy is output as LV-BERT-small. ⑧ LV-BERT-small can be scaled up to LV-BERT-medium/base with larger hidden size.

Beyond the above manually designed orders, we take advantage of neural architecture search to identify more effective layer orders for pre-trained models. The order to be discovered can be expressed as

$$[L_1, L_2, \dots, L_i, \dots, L_N], \quad (4)$$

where  $L_i \in \mathcal{L}^{\text{type}}$  and  $N$  is the number of layers. Here,  $N$  is set to 24, following common practice.

### 3.2 Supernet

The layer variety introduced above leads to a huge architecture space of  $3^{24} \approx 2.8 \times 10^{11}$  candidate models to be explored. Thus, it is not affordable to pre-train every candidate model in the space from scratch to evaluate their performance since the pre-training procedure requires huge computations. To reduce the search computations, recent NAS works (Pham et al., 2018; Guo et al., 2020; Cai et al., 2019) exploit a weight sharing strategy. It first trains a supernet subsuming all candidate architectures, and then each candidate architecture can inherit its weights from the trained supernet to avoid training from scratch. Inspired by this strategy, we construct a supernet where each layer contains all types of layers, *i.e.*, self-attention, feed-forward, and dynamic convolution. The supernet

architecture can be expressed as

$$\mathcal{A} = [\{L_1^{\text{SA}}, L_1^{\text{FF}}, L_1^{\text{DC}}\}, \{L_2^{\text{SA}}, L_2^{\text{FF}}, L_2^{\text{DC}}\}, \dots, \{L_N^{\text{SA}}, L_N^{\text{FF}}, L_N^{\text{DC}}\}]. \quad (5)$$

Masked Language Modeling (MLM) (Devlin et al., 2019) is utilized as the pre-training objective to pre-train the supernet since MLM accuracy can reflect the model performance on downstream tasks (Lan et al., 2020). Most weight sharing approaches on NAS (Wu et al., 2019a; Liu et al., 2018) train and optimize the full supernet: the output of each layer is the weighted sum of all types of candidate layers. However, it cannot guarantee the sampled single type of layer also works (Guo et al., 2020).

To handle this issue, we propose to randomly sample a submodel from the supernet to participate in forward and backward propagation per training step (Cai et al., 2018; Guo et al., 2020). The sampled submodel architecture can be expressed as

$$a = [L_1, L_2, \dots, L_i, \dots, L_N], \quad (6)$$

where  $L_i \in \mathcal{L}^{\text{type}} \sim U$  with uniform probability distribution  $Pr = 1/3$ . In this pre-training method, the optimized supernet weights can be expressed

---

**Algorithm 1:** Evolutionary Search Guided by Pre-training MLM Accuracy

---

**Input:**  $W_{\mathcal{A}}$ : supernet weights;  $P$ : population size;  $D_{\text{val}}$ : pre-training validation set;  $T$ : # iteration;  $N^{\text{cro}}$ : # crossover;  $N^{\text{mut}}$ : # mutation;  $p$ : mutation probability;  $k$ : # top candidates for crossover and mutation  
**Output:**  $a^*$ : the architecture with the best pre-training MLM validation accuracy  
 $S_0 := \text{Init}(P)$ ; // Randomly generate  $P$  architecture candidates  
 $S^{\text{topk}} := \emptyset$ ; // The set of top  $k$  candidates  
**for**  $i = 1 : T$  **do**  
     $S_{i-1}^{\text{MLM}} := \emptyset$ ;  
    **for**  $a$  in  $S_{i-1}$  **do**  
         $\text{MLM}_{\text{val}}^a :=$   
        Inference( $\mathcal{N}(a, W_{\mathcal{A}}(a)), D_{\text{val}}$ );  
         $S_{i-1}^{\text{MLM}} := S_{i-1}^{\text{MLM}} \cup \text{MLM}_{\text{val}}^a$ ;  
     $S^{\text{topk}} := \text{Update}(S^{\text{topk}}, S_{i-1}, S_{i-1}^{\text{MLM}})$ ;  
     $S^{\text{cro}} := \text{Crossover}(S^{\text{topk}}, N^{\text{cro}})$ ;  
     $S^{\text{mut}} := \text{Mutation}(S^{\text{topk}}, N^{\text{mut}}, p)$ ;  
     $S_i := S^{\text{cro}} \cup S^{\text{mut}}$ ;  
**return**  $a^* = \text{argmax}_{a \in S^{\text{topk}}} \text{MLM}_{\text{val}}^a$ ;

---

as

$$W_{\mathcal{A}} = \underset{W}{\text{argmin}} \mathbb{E}_{a \sim U(\mathcal{A})} [\mathcal{L}_{\text{pre-train}}(\mathcal{N}(a, W(a)))], \quad (7)$$

where  $W(a)$  denotes the submodel weights inherited from the supernet,  $\mathcal{N}$  means the submodel with specific architecture and weights,  $\mathcal{L}_{\text{pre-train}}$  denotes the pre-training MLM loss and  $a \sim U(\mathcal{A})$  means  $a$  is uniformly sampled from  $\mathcal{A}$ .

### 3.3 Evolutionary Search

Inspired by the recent NAS works (Elsken et al., 2019; Ren et al., 2020; Guo et al., 2020; Wang et al., 2020), we adopt an evolutionary algorithm (EA) to search the model. Previously Real et al. (2017) utilized an evolutionary method in NAS but they trained each candidate model from scratch which is costly and inefficient. Instead, thanks to the supernet mentioned above, we do not need to train the candidate models from scratch since their weights can be inherited from the supernet. Next problem is how to select indicator of the candidate models to guide the EA. Note that our goal is to search a general pre-trained model to benefit a variety of downstream tasks instead of a specific task.

Traditional NAS methods (Chen et al., 2020; Zhu et al., 2020) use downstream task performance as the objective to search for task-specific models. Instead, similar to the work by Khetan and Karnin (2020) that utilize pre-training loss to prune BERT, our method uses pre-training MLM accuracy to search for a unified architecture that can generalize well to different downstream tasks. Besides, using this accuracy, candidate models can be directly evaluated on pre-training validation set without any fine-tuning on specific tasks, which can help save computations.

The detailed algorithm description is shown in Algorithm 1.  $\text{Crossover}(S^{\text{topk}}, N^{\text{cro}})$  means the procedure to generate  $N^{\text{cro}}$  new candidate architectures that two candidate architectures randomly selected from top  $k$  candidate set  $S^{\text{topk}}$  are crossed to produce a new one. Similarly,  $\text{Mutation}(S^{\text{topk}}, N^{\text{mut}}, p)$  denotes the procedure to generate  $N^{\text{mut}}$  new candidates that a random candidate from  $S^{\text{topk}}$  mutates its every layer choice with probability  $p$  to generate a new one. Finally, the candidate architecture with highest pre-training validation accuracy in  $S^{\text{topk}}$  is returned as LV-BERT. The algorithm is set with population size  $P$  of 50, search iteration number  $T$  of 20, crossover number  $N^{\text{cro}}$  of 25, mutation number  $M^{\text{mut}}$  of 25, mutation probability  $p$  of 0.1, top candidate number  $k$  of 10 for crossover and mutation.

## 4 Experiments

### 4.1 Datasets

**Pre-training Datasets** Devlin et al. (2019) propose WikiBooks corpus for training BERT including English Wikipedia and BooksCorpus (Zhu et al., 2015). However, BooksCorpus is no longer publicly available. To ease reproduction, we train models on OpenWebText (Gokaslan and Cohen, 2019) that is open-sourced and of similar size with the corpus used by BERT. When pre-training the supernet, we leave 2% data as our validation set for evolutionary search.

**Fine-tuning Datasets** To compare our model with other pre-trained models, we fine-tune LV-BERT on GLUE (Wang et al., 2018), including various tasks for general language understanding, and SQuAD 1.1/2.0 (Rajpurkar et al., 2016, 2018) for question answering. See Appendix for more details of all tasks.

Model	Layer Variety				Params		GLUE
	DC	SA	FF	Order	Word Emb	Backbone	
BERT-small (Devlin et al., 2019)		✓	✓	Interleaved		9.5M	75.1
ELECTRA-small (Clark et al., 2020)		✓	✓	Interleaved	3.9M	9.5M	80.4
DynamicConv-small* (Wu et al., 2018)	✓		✓	Interleaved		9.6M	64.4
Sandwich-small* (Press et al., 2020)	✓		✓	Sandwich		9.5M	78.6
LV-BERT-small variants		✓	✓	Random		9.5M	80.8
			✓	Randomly searched		9.8M	81.1
			✓	EA searched		10.3M	81.2
	✓		✓	Random		9.6M	64.9
	✓		✓	Randomly searched		9.6M	65.4
	✓		✓	EA searched	3.9M	9.6M	65.7
	✓	✓		Random		6.4M	79.7
	✓	✓		Randomly searched		6.4M	79.9
	✓	✓		EA searched		6.4M	79.8
	✓	✓	✓	Random		7.7M	80.6
LV-BERT-small	✓	✓	✓	EA searched	3.9M	8.5M	<b>81.8</b>

Table 1: Performance of the models with different layer types and orders on the GLUE development set. DC, SA and FF denote dynamic convolution, self-attention and feed-forward layers respectively. For each design of layer type set, “Random” means the best order among five randomly generated ones that are estimated by training model from scratch. “Randomly searched” or “EA searched” are both based on the supernet. “Randomly searched” denotes the orders searched at random while “EA searched” denotes ones searched by evolutionary algorithm. \* denotes the methods implemented by us for language pre-training. All models are pre-trained on OpenWebText by 1M steps with sequence length 128 using ELECTRA (Clark et al., 2020) pre-training objective except BERT-small using MLM objective.

Model	Size	Params		CoLA	MPRC	MNLI	SST	RTE	QNLI	QQP	STS	Avg.
		Word Emb	Backbone									
ELECTRA (Clark et al., 2020)	Small	3.9M	9.5M	56.8	87.4	78.9	88.3	68.5	87.9	88.3	86.8	80.4
	Medium*	3.9M	21.3M	61.2	89.5	82.1	89.1	65.7	88.9	90.5	89.3	82.0
	Base*	23.4M	85.0M	64.8	88.5	85.7	92.6	76.5	91.7	<b>91.1</b>	89.9	85.1
DynamicConv <sup>†</sup> (Wu et al., 2018)	Small	3.9M	9.6M	60.2	69.2	56.6	85.6	49.5	68.0	82.1	44.1	64.4
	Medium	3.9M	21.4M	61.5	67.9	55.7	85.9	49.1	68.3	83.3	51.6	65.4
	Base	23.4M	85.2M	62.1	70.6	61.0	88.5	51.3	72.0	85.6	64.7	69.5
Sandwich <sup>†</sup> (Press et al., 2020)	Small	3.9M	9.5M	53.2	87.1	77.5	88.1	63.9	86.4	88.3	84.6	78.6
	Medium	3.9M	21.3M	55.6	86.2	81.5	90.3	63.0	88.9	89.6	86.6	80.2
	Base	23.4M	85.0M	58.8	89.7	83.8	91.9	72.6	90.2	90.1	88.5	83.2
LV-BERT	Small	3.9M	8.5M	62.3	86.9	81.1	89.9	69.0	88.9	89.3	87.4	81.8
	Medium	3.9M	19.0M	64.4	88.0	82.4	90.5	68.6	89.4	90.1	89.7	82.9
	Base	23.4M	75.7M	<b>66.8</b>	<b>90.3</b>	<b>86.3</b>	<b>93.2</b>	<b>76.9</b>	<b>92.3</b>	90.9	<b>90.8</b>	<b>85.9</b>

Table 2: Performance of different models in different sizes on GLUE development set. \* denotes results obtained by running official code. <sup>†</sup> denotes the methods implemented by us for language pre-training. All models are pre-trained on OpenWebText by 1M steps with sequence length 128 using ELECTRA (Clark et al., 2020) pre-training objective.

## 4.2 Implementation Details

**Model Size** Similar to Devlin et al. (2019), Clark et al. (2020) and Jiang et al. (2020), we define different model sizes, *i.e.*, “small”, “medium” and “base”, with the same layer number of 24 but different hidden sizes of 256, 384, and 768, respectively. The detailed hyperparameters are shown in Appendix.

**Pre-training Supernet** To reduce training cost, we construct the supernet only in small size. Since the layer number of models in medium and base sizes are the same as that of the small-sized one,

the obtained architecture of LV-BERT-small can be easily scaled up to the ones of medium and base sizes. We use Adam (Kingma and Ba, 2015) to pre-train the supernet with MLM loss (Devlin et al., 2019), learning rate of  $2e-4$ , batch size of 128, max sequence length of 128 and pre-training step number of 2 million. See Appendix for more details.

**Evaluation Setup** To compare with other pre-trained models, we pre-train the searched LV-BERT architecture for 1M steps from scratch on the OpenWebText (Gokaslan and Cohen, 2019) using Re-

placed Token Detection (Clark et al., 2020) since it can save computation cost. We fine-tune LV-BERT on GLUE (Wang et al., 2018) and SQuAD (Rajpurkar et al., 2016, 2018) downstream tasks with most hyperparameters the same as those of ELECTRA (Clark et al., 2020) for fair comparison. For GLUE tasks, the evaluation metrics are Matthews correlation for CoLA, Spearman correlation for STS, and accuracy for other tasks, which are averaged to get GLUE score. We utilize evaluation metrics of Exact-Match (EM) and F1 for SQuAD 1.1/2.0. Some of the fine-tuning datasets are small, and consequently, the results may vary substantially for different random seeds. Similar to ELECTRA (Clark et al., 2020), we report the median of 10 fine-tuning runs from the same pre-trained model for each result. See Appendix for more evaluation details.

### 4.3 Ablation Study

**Layer Variety** Various models are constructed with different layer variety designs, and their results on GLUE development set are shown in Table 1. For the layer types, if only two layer types are provided, selecting self-attention and feed-forward yields the best result, which can always achieve performance higher than 80 under different search methods. With only dynamic convolution and feed-forward, the performance drops dramatically to around 65. Surprisingly, without feed-forward, the layer set of dynamic convolution and self-attention can still achieve relatively good score, near 80. When using all the three layer types, we can obtain the best 81.8 score, 1.4 higher than the strong baseline ELECTRA (80.4) and 0.6 higher than the model searched with only self-attention and feed-forward (81.2). This indicates that it is effective to augment the layer type set by including convolution to extract local information for pre-trained models.

For layer orders, with the same layer types, the models with either EA or randomly searched orders perform better than those with randomly sampled orders, reflecting the importance of investigating layer orders. For example, with the same layer types of self-attention and feed-forward, the EA searched model obtains 81.2 score, improving BERT/ELECTRA by 6.1/0.8 as well as Sandwich by 2.6.

**Search Method** Table 1 shows the results with different search methods. “Random” means for




Figure 3: The pre-training MLM validation accuracy comparison between random search and evolutionary search with the layer set of all three types of layers. Blue and yellow dots denote the accuracy of top 10 candidates for each method respectively, while the plots mean their averages.

each design of layer type set, the order is the best one among 5 randomly generated orders that are estimated by training models from scratch. “Randomly searched” and “EA searched” are both supernet-based methods, in which the weights of candidate models are inherited from the supernet. “Randomly searched” produces candidate models at random for estimation while “EA searched” generates candidate models with evolutionary algorithm guided by the pre-training MLM accuracy. With the same layer types, EA searched orders are generally better than randomly searched ones while the randomly searched ones are generally better than random ones. Figure 3 plots the pre-training MLM evaluation accuracy over search iterations with both random and evolutionary search methods. It shows that the accuracy of evolutionary search is obviously higher than that of random search, demonstrating the effectiveness of evolutionary search.

### 4.4 LV-BERT Architecture

As shown in Table 1, LV-BERT achieves the best performance. Its architecture is

$$\begin{aligned}
 & [L_1^{DC}, L_2^{DC}, L_3^{SA}, L_4^{FF}, L_5^{FF}, L_6^{SA}, \\
 & L_7^{DC}, L_8^{FF}, L_9^{FF}, L_{10}^{SA}, L_{11}^{DC}, L_{12}^{DC}, \\
 & L_{13}^{SA}, L_{14}^{FF}, L_{15}^{DC}, L_{16}^{FF}, L_{17}^{SA}, L_{18}^{DC}, \\
 & L_{19}^{FF}, L_{20}^{DC}, L_{21}^{SA}, L_{22}^{SA}, L_{23}^{FF}, L_{24}^{SA} ].
 \end{aligned} \tag{8}$$

Pre-trained with MLM from scratch by 1M steps (sequence length 128) on OpenWebText, LV-BERT-small can achieve 61.2% MLM accuracy while BERT-small is 60.4%. More specific architectures of the models in Table 1 are listed in Appendix.

Model	Train FLOPs	Params	CoLA	MPRC	MNLI	SST	RTE	QNLI	QQP	STS	Avg.
TinyBERT* (Jiao et al., 2020)	6.4e19+ (54x+)	15M	51.1	82.6	84.6	93.1	70.0	90.4	89.1	83.7	80.6
MobileBERT* (Sun et al., 2020)	6.4e19+ (54x+)	25M	51.1	84.5	84.3	92.6	70.4	91.6	88.3	84.8	81.0
ELECTRA-small (Clark et al., 2020)	1.4e18 (1.2x)	14M	54.6	83.7	79.7	89.1	60.8	87.7	88.0	80.2	78.0
GPT (Radford et al., 2018)	4.0e19 (33x)	117M	45.4	75.7	82.1	91.3	56.0	88.1	88.5	80.0	75.9
BERT-base (Devlin et al., 2019)	6.4e19 (54x)	110M	52.1	84.8	84.6	93.5	66.4	90.5	89.2	85.8	80.9
ELECTRA-base (Clark et al., 2020)	6.4e19 (54x)	110M	59.7	86.7	85.8	93.4	73.1	<b>92.7</b>	89.1	87.7	83.5
LV-BERT-small	1.2e18 (1x) <sup>†</sup>	13M	57.2	84.1	81.0	90.4	64.6	88.9	88.2	83.8	79.8
LV-BERT-medium	3.1e18 (2.6x) <sup>†</sup>	23M	60.1	85.0	82.0	91.4	67.6	89.7	88.9	85.9	81.3
LV-BERT-base	1.8e19 (15x) <sup>†</sup>	100M	<b>64.0</b>	<b>87.9</b>	<b>86.4</b>	<b>94.7</b>	<b>77.0</b>	92.6	<b>89.5</b>	<b>88.8</b>	<b>85.1</b>

Table 3: Performance of models with similar size on GLUE testing set. \* denotes knowledge distillation methods that rely on large pre-trained teacher models and are orthogonal to other methods. <sup>†</sup> We set the sequence length as 128 for pre-training to save computation although it hurts the performance.

Model	Train FLOPs	Params	SQuAD 1.1		SQuAD 2.0	
			EM	F1	EM	F1
DistillBERT* (Sanh et al., 2019)	6.4e19+ (54x+)	52M	71.8	81.2	60.6	64.1
TinyBERT* (Jiao et al., 2020)	6.4e19+ (54x+)	15M	72.7	82.1	65.3	68.8
MobileBERT* (Sun et al., 2020)	6.4e19+ (54x+)	25M	83.4	90.3	77.6	80.2
ELECTRA-small <sup>†</sup> (Clark et al., 2020)	1.4e18 (1.2x)	14M	74.3	81.8	66.8	69.4
BERT-base (Devlin et al., 2019)	6.4e19 (54x)	110M	80.7	88.4	74.2	77.1
ELECTRA-base (Clark et al., 2020)	6.4e19 (54x)	110M	84.5	<b>90.8</b>	80.5	83.3
LV-BERT-small	1.2e18 (1x) <sup>‡</sup>	13M	77.1	84.1	71.0	73.7
LV-BERT-medium	3.1e18 (2.6x) <sup>‡</sup>	23M	79.6	86.4	74.9	77.5
LV-BERT-base	1.8e19 (15x) <sup>‡</sup>	100M	<b>84.8</b>	<b>90.8</b>	<b>80.9</b>	<b>83.7</b>

Table 4: Performance of models with similar model size on SQuAD 1.1/2.0 development set. \* denotes knowledge distillation methods that rely on large pre-trained teacher models and are orthogonal to other methods. <sup>†</sup> denotes results obtained by running official code. <sup>‡</sup> We set the sequence length as 128 for pre-training to save computation although it hurts the performance.

When running the evolutionary method with different seeds, we see that the resulting models prefer stacking dynamic convolutions at the bottom two layers for extracting local information and self-attention at the top layer to fuse the global information. According to these observation, for ELECTRA-small, if we replace the bottom two layers with dynamic convolutions or the top layer with self-attention, the performance can be improved by 0.3 or 0.5 respectively on GLUE development set. If we replace the bottom 8 layers with manually designed ‘ccsfccsf’ (‘c’, ‘s’ and ‘f’ denote dynamic convolution, self-attention and feed-forward layers, respectively) and replace the top 8 layers with manually designed ‘ssfsfsfs’ together, we observe 0.7 performance improvement. These results show that it is helpful to stack dynamic convolution at the bottom and self-attention at the top.

#### 4.5 Generalization to Larger Models

We only investigate layer variety and search models in a small-sized setting to save computation cost. It is interesting to know whether the searched models can be generalized to larger models with large hidden size. The results are shown in Table

2. For larger model size “medium” and “base”, LV-BERTs still outperform other baseline models, demonstrating the good generalization in terms of model size.

#### 4.6 Comparison with State-of-the-arts

We compare LV-BERT with state-of-the-art pre-trained models (Radford et al., 2018; Devlin et al., 2019; Clark et al., 2020; Sanh et al., 2019; Jiao et al., 2020; Sun et al., 2020) on GLUE testing set and SQuAD 1.1/2.0 to show its advantages. Although more pre-training data/steps and larger model size can significantly help improve performance (Yang et al., 2019; Liu et al., 2019b; Lan et al., 2020), due to the computation resource limit, we only pre-train our models in small/medium/base sizes for 1M steps with OpenWebText (Gokaslan and Cohen, 2019). We leave evaluating models with more pre-training data/steps and larger model size for future work. We also list some knowledge distillation methods for comparison. However, note that these methods rely on a pre-trained large teacher network and thus are orthogonal to LV-BERT and other methods.

Table 3 presents the performance of LV-BERT



and other pre-trained models on GLUE testing set. It shows that LV-BERT outperforms other pre-trained models with similar model size. Remarkably, LV-BERT-small/base achieve 79.8/85.1, 1.8/1.6 higher than strong baselines ELECTRA-small/base. Even compared with knowledge distillation based model MobileBERT (Sun et al., 2020), LV-BERT-medium still outperforms it by 0.3.

Since there is nearly no single model submission on SQuAD leaderboard<sup>2</sup>, we only compare LV-BERT with other pre-trained models on the development sets. The results are shown in Table 4. We find that LV-BERT-small outperforms ELECTRA-small significantly, like F1 score 73.7 versus 69.4 on SQuAD 2.0. However, when we generalize LV-BERT-small to base size, the gap between LV-BERT and ELECTRA with base size is narrower than that with small size. One reason may be LV-BERT-small is searched by our method while LV-BERT-base is only generalized from LV-BERT-small with larger hidden size.

## 5 Conclusion

We are the first to exploit layer variety for improving pre-trained language models, from two aspects, *i.e.*, layer types and layer orders. For layer types, we augment the layer type set by including convolution for local information extraction. For layer orders, beyond the stereotyped interleaved one, we explore more effective orders by using an evolutionary based search algorithm. Experiment results show our obtained model LV-BERT outperforms BERT and its variants on various downstream tasks.

## Acknowledgments

We would like to thank the anonymous reviewers for their insightful comments and suggestions. This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-100E/2019-035). Jiashi Feng was partially supported by MOE2017-T2-2-151, NUS\_ECRA\_FY17\_P08 and CRP20-2017-0006. The authors also thank Qunhong Fu and Jian Liang for the help to improve the technical writing aspect of this paper. The computational work for this article was partially performed on resources of the National Supercomputing Centre, Singapore (<https://www.nsc.sg>). Weihao Yu would like to thank TPU Research Cloud (TRC)

<sup>2</sup>[rajpurkar.github.io/SQuAD-explorer/](https://rajpurkar.github.io/SQuAD-explorer/)

program and Google Cloud Research Credits Program for the support of computational resources.

## References

- Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. 2017. [Designing neural network architectures using reinforcement learning](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. 2018. Understanding and simplifying one-shot architecture search. In *International Conference on Machine Learning*, pages 550–559.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2009. The fifth pascal recognizing textual entailment challenge. In *TAC*.
- Andrew Brock, Theodore Lim, James Millar Ritchie, and Nicholas J Weston. 2018. Smash: One-shot model architecture search through hypernetworks. In *6th International Conference on Learning Representations*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Gino Brunner, Yang Liu, Damian Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. 2020. [On identifiability in transformers](#). In *International Conference on Learning Representations*.
- Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. 2019. Once-for-all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*.
- Han Cai, Ligeng Zhu, and Song Han. 2018. Proxylessnas: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and](#)

- crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Daoyuan Chen, Yaliang Li, Minghui Qiu, Zhen Wang, Bofang Li, Bolin Ding, Hongbo Deng, Jun Huang, Wei Lin, and Jingren Zhou. 2020. *Adabert: Task-adaptive BERT compression with differentiable neural architecture search*. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 2463–2469. ijcai.org.
- Yukang Chen, Tong Yang, Xiangyu Zhang, Gaofeng Meng, Chunhong Pan, and Jian Sun. 2019. Detnas: Neural architecture search on object detection. *arXiv preprint arXiv:1903.10979*, 1(2):4–1.
- Zihan Chen, Hongbo Zhang, Xiaoji Zhang, and Leqi Zhao. 2018. Quora question pairs.
- Zewen Chi, Li Dong, Furu Wei, Wenhui Wang, Xian-Ling Mao, and Heyan Huang. 2019. Cross-lingual natural language generation via pre-training. *arXiv preprint arXiv:1909.10481*.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*, pages 7057–7067.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pages 177–190. Springer.
- Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language modeling with gated convolutional networks. In *International conference on machine learning*, pages 933–941. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- William B. Dolan and Chris Brockett. 2005. *Automatically constructing a corpus of sentential paraphrases*. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Thomas Elsken, Jan Hendrik Metzen, Frank Hutter, et al. 2019. Neural architecture search: A survey. *J. Mach. Learn. Res.*, 20(55):1–21.
- Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. 2019. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7036–7045.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pages 1–9. Association for Computational Linguistics.
- Aaron Gokaslan and Vanya Cohen. 2019. Openwebtext corpus. <http://SkyLion007.github.io/OpenWebTextCorpus>.
- Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. 2020. Single path one-shot neural architecture search with uniform sampling. In *European Conference on Computer Vision*, pages 544–560. Springer.
- R Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. 2006. The second pascal recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2485–2494.
- Zi-Hang Jiang, Weihao Yu, Daquan Zhou, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. 2020. Convbart: Improving bert with span-based dynamic convolution. *Advances in Neural Information Processing Systems*, 33.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. *Tinybert: Distilling BERT for natural language understanding*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online*

- Event, 16-20 November 2020, pages 4163–4174. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [Spanbert: Improving pre-training by representing and predicting spans](#). *Trans. Assoc. Comput. Linguistics*, 8:64–77.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. [A convolutional neural network for modelling sentences](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 655–665, Baltimore, Maryland. Association for Computational Linguistics.
- Ashish Khetan and Zohar Karnin. 2020. [schubert: Optimizing elements of bert](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2807–2818.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of bert](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4365–4374.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. [Imagenet classification with deep convolutional neural networks](#). In *NeurIPS*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. [Gradient-based learning applied to document recognition](#). *Proceedings of the IEEE*, 86(11):2278–2324.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The winograd schema challenge](#). In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*.
- Jiahuang Lin, Xin Li, and Gennady Pekhimenko. 2020. [Multi-node bert-pretraining: Cost-efficient approach](#). *arXiv preprint arXiv:2008.00177*.
- Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L Yuille, and Li Fei-Fei. 2019a. [Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 82–92.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2018. [Darts: Differentiable architecture search](#). In *International Conference on Learning Representations*.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. [K-bert: Enabling language representation with knowledge graph](#). In *AAAI*, pages 2901–2908.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. [Learned in translation: Contextualized word vectors](#). In *Advances in Neural Information Processing Systems*, pages 6294–6305.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. [Distributed representations of words and phrases and their compositionality](#). In *Advances in neural information processing systems*, pages 3111–3119.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018a. [Deep contextualized word representations](#). In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018b. [Deep contextualized word representations](#). In *Proc. of NAACL*.
- Matthew E Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. 2019. [Knowledge enhanced contextual word representations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 43–54.
- Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. 2018. [Efficient neural architecture search via parameters sharing](#). In *International Conference on Machine Learning*, pages 4095–4104.

- Ofir Press, Noah A. Smith, and Omer Levy. 2020. [Improving transformer models by reordering their sub-layers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2996–3005. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100, 000+ questions for machine comprehension of text. In *EMNLP*.
- Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. 2017. Large-scale evolution of image classifiers. In *International Conference on Machine Learning*, pages 2902–2911.
- Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. 2020. A comprehensive survey of neural architecture search: Challenges and solutions. *arXiv preprint arXiv:2006.02903*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Karen Simonyan and Andrew Zisserman. 2015. [Very deep convolutional networks for large-scale image recognition](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- David So, Quoc Le, and Chen Liang. 2019. The evolved transformer. In *International Conference on Machine Learning*, pages 5877–5886.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*, pages 1631–1642.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. [Mobilebert: a compact task-agnostic BERT for resource-limited devices](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2158–2170. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30:5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. 2020. [HAT: Hardware-aware transformers for efficient natural language processing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7675–7688, Online. Association for Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *NAACL-HLT*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. 2019a. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10734–10742.
- Felix Wu, Angela Fan, Alexei Baevski, Yann Dauphin, and Michael Auli. 2018. Pay less attention with lightweight and dynamic convolutions. In *International Conference on Learning Representations*.
- Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han. 2019b. Lite transformer with long-short range attention. In *International Conference on Learning Representations*.
- Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. 2018. Snas: stochastic neural architecture search. In *International Conference on Learning Representations*.
- Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. 2020. [Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model](#). In *International Conference on Learning Representations*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5753–5763.

Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. [Relation classification via convolutional deep neural network](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [Ernie: Enhanced language representation with informative entities](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451.

Wei Zhu, Xiaoling Wang, Xipeng Qiu, Yuan Ni, and Guotong Xie. 2020. [Auroc: Improving bert based relation classification models via architecture search](#). *arXiv preprint arXiv:2009.10680*.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.

Barret Zoph and Quoc V. Le. 2017. [Neural architecture search with reinforcement learning](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

## A Details about Layer Types

For a layer, assume its input is  $I \in \mathbb{R}^{s \times c}$  and output is  $O \in \mathbb{R}^{s \times c}$ , where  $s$  is the sequence length and  $c$  is the hidden size (channel dimension). For simplicity,  $c$  takes the same value for the input and output.

**Self-Attention** The self-Attention layer, also known as multi-head self-attention (Vaswani et al., 2017), transforms the input by three linear transformations into the key  $K$ , query  $Q$  and value  $V$  vectors respectively,

$$\begin{aligned} K &= \text{Reshape}(IW^K + b^K) \\ Q &= \text{Reshape}(IW^Q + b^Q) \\ V &= \text{Reshape}(IW^V + b^V), \end{aligned} \quad (9)$$

where  $K, Q, V \in \mathbb{R}^{h \times s \times d}$ ,  $W^K, W^Q, W^V \in \mathbb{R}^{c \times c}$ , and  $b^K, b^Q, b^V \in \mathbb{R}^c$ . Notice that  $h \times d = c$  where  $h$  is the number of heads and  $d$  is the head dimension.

The above  $K$  and  $Q$  are used to compute their similarity matrix  $M$  which is then used to generate

new value  $V'$ :

$$\begin{aligned} M &= \text{Softmax}(KQ^\top / \sqrt{d}) \\ V' &= \text{Reshape}(MV), \end{aligned} \quad (10)$$

where  $M \in \mathbb{R}^{h \times s \times s}$  and  $V' \in \mathbb{R}^{s \times c}$ . Finally, a linear transformation is used to exchange information between different heads, followed by shortcut connection and layer normalization,

$$O = \text{Norm}(V'W_O + b_O + I), \quad (11)$$

where  $W_O \in \mathbb{R}^{c \times c}$  and  $b_O \in \mathbb{R}^c$ .

**Feed-Forward** The feed-forward layer (Vaswani et al., 2017) includes two linear transformations with a non-linear activation, followed by a shortcut connection and layer normalization,

$$\begin{aligned} N &= \text{GELU}(IW_1 + b_1) \\ O &= \text{Norm}(NW_2 + b_2 + I), \end{aligned} \quad (12)$$

where  $W_1 \in \mathbb{R}^{c \times rc}$  and  $W_2 \in \mathbb{R}^{rc \times c}$  with a ratio  $r$ .  $\text{GELU}(\cdot)$  denotes the Gaussian Error Linear Unit (Hendrycks and Gimpel, 2016).

**Dynamic Convolution** Dynamic convolution is introduced by Wu et al. (2018) to replace self-attention, which shows strong competitiveness in the tasks of machine translation, language modeling and summarization. The dynamic convolution first uses gated linear unit (GLU) (Dauphin et al., 2017) to generate new representation,

$$V = \text{GLU}(I). \quad (13)$$

Different from the vanilla dynamic convolution that directly generates dynamic kernel from  $V \in \mathbb{R}^{s \times c}$ , in this work, we supplement a separate convolution (Howard et al., 2017) with depthwise weights  $W^{\text{Dep}} \in \mathbb{R}^{k \times c}$  ( $k$  is the convolution kernel size, set as 9 in this paper) and pointwise weights  $W^{\text{Poi}} \in \mathbb{R}^{c \times c}$  to extract local information to help the following kernel generation. Denoting the output as  $S \in \mathbb{R}^{s \times c}$ , the separate convolution can be formulated as

$$S_{i,:} = \left( \sum_{j=1}^k W_{j,:}^{\text{Dep}} \cdot V_{i+j-\frac{k+1}{2},:} \right) W^{\text{Poi}}. \quad (14)$$

Then the output of separate convolution is used to generate dynamic kernels,

$$D = \text{Softmax}(\text{Reshape}(SW^{\text{Dyn}})), \quad (15)$$

where  $W^{Dyn} \in \mathbb{R}^{c \times hk}$  and  $D \in \mathbb{R}^{h \times s \times k}$ . Then lightweight convolution is applied to the reshaped  $V' = \text{Reshape}(V) \in \mathbb{R}^{h \times s \times d}$ . The output  $C \in \mathbb{R}^{h \times s \times d}$  can be expressed as

$$C_{p,i,:} = \sum_{j=1}^k D_{p,i,j} \cdot V'_{p,i+j-\frac{k+1}{2},:} \quad (16)$$

Finally,  $C$  is reshaped to  $C' = \text{Reshape}(C) \in \mathbb{R}^{s \times c}$  and a linear transformer is applied to fuse the information among multiple heads, followed by a short connection and layer normalization,

$$O = \text{Norm}(C'W^{\text{Out}} + b^{\text{Out}} + I), \quad (17)$$

where  $W^{\text{Out}} \in \mathbb{R}^{c \times c}$  and  $b^{\text{Out}} \in \mathbb{R}^c$ .

## B Details about Datasets

### B.1 GLUE Dataset

Introduced by Wang et al. (2018), General Language Understanding Evaluation (GLUE) benchmark is a collection of nine tasks for natural language understanding, where testing set labels are hidden and predictions need to be submitted to the evaluation server<sup>3</sup>. We provide details about the GLUE tasks below.

**CoLA** The Corpus of Linguistic Acceptability (Warstadt et al., 2019) is a binary single-sentence classification dataset for predicting whether a sentence is grammatical or not. The samples are from books and journal articles on linguistic theory.

**MRPC** The Microsoft Research Paraphrase Corpus (Dolan and Brockett, 2005) is a dataset for the task to predict whether two sentences are semantically equivalent or not. It is extracted from online news sources with human annotations.

**MNLI** The Multi-Genre Natural Language Inference Corpus (Williams et al., 2018) is a dataset of sentence pairs. Each pair has a premise sentence and a hypothesis sentence, requiring models to predict its relationships containing *entailment*, *contradiction* or *neutral*. It is from ten distinct genres of spoken and written English.

**SST** The Stanford Sentiment Treebank (Socher et al., 2013) is a dataset for the task to predict whether a sentence is positive or negative in sentiment. The dataset is from movie reviews with human annotations.

<sup>3</sup><https://gluebenchmark.com>

Hyperparameter	Supernet	Small	Medium	Base
Layer number	24	24	24	24
Word emb. size	128	128	128	768
Hidden size	256	256	384	768
FF inner hidden size	1024	1024	1536	3072
Generator size	N/A	1/4	1/3	1/3
Head number	4	4	6	12
Head size	64	64	64	64
Learning rate	2e-4	5e-4	5e-4	2e-4
Learning rate decay	Linear	Linear	Linear	Linear
Warmup steps	10000	10000	10000	10000
Adam $\epsilon$	1e-6	1e-6	1e-6	1e-6
Adam $\beta_1$	0.9	0.9	0.9	0.9
Adam $\beta_2$	0.999	0.999	0.999	0.999
Dropout	0.1	0.1	0.1	0.1
Batch size	128	128	128	256
Input sequence length	128	128	128	128

Table 5: Pre-training hyperparameters. Generator size means the multiplier for hidden size, feed-forward inner hidden size and head number to construct generator for Replaced Token Detection pre-training objective (Clark et al., 2020).

Hyperparameter	Value
Learning rate	3e-4 for small/medium size 1e-4 (except 2e-4 for SQuAD) for base size
Adam $\epsilon$	1e-6
Adam $\beta_1$	0.9
Adam $\beta_2$	0.999
Layerwise LR decay	0.8 for every two layers
Learning rate decay	Linear
Warmup fraction	0.1
Attention Dropout	0.1
Dropout	0.1
Weight decay	0.01
Batch size	32
Train epochs	10 for RTE and STS, 2 for SQuAD, and 3 for other tasks

Table 6: Fine-tuning hyperparameters.

**RTE** The Recognizing Textual Entailment (RTE) dataset is for the task to determine whether the relationship of a pair of premise and hypothesis sentences is entailment. The dataset is from several annual textual entailment challenges including RTE1 (Dagan et al., 2005), RTE2 (Haim et al., 2006), RTE3 (Giampiccolo et al., 2007), and RTE5 (Bentivogli et al., 2009).

**QNLI** Question Natural Language Inference is a dataset converted from The Stanford Question Answering Dataset (Rajpurkar et al., 2016). An example is a pair of a context sentence and a question, requiring to predict whether the context sentence contains the answer to the given question.

**QQP** The Quora Question Pairs dataset (Chen et al., 2018) is the dataset from Quora, requiring to

Model	Layer Variety			Architecture	GLUE
	DC	SA	FF Order		
BERT-small	✓	✓	Interleaved	[1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2]	75.1
ELECTRA-small		✓	Interleaved	[1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2]	80.4
DynamicConv-small*	✓		Interleaved	[0, 2, 0, 2, 0, 2, 0, 2, 0, 2, 0, 2, 0, 2, 0, 2, 0, 2, 0, 2]	64.4
Sandwich-small*	✓		Sandwich	[1, 1, 1, 1, 1, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 1, 2, 2, 2, 2, 2, 2]	78.6
LV-BERT-small variants		✓	Random	[1, 1, 1, 2, 2, 1, 2, 2, 2, 1, 1, 2, 2, 2, 1, 1, 1, 1, 2, 2, 2, 1, 1]	80.8
		✓	Randomly searched	[1, 2, 1, 1, 2, 2, 1, 2, 1, 2, 2, 1, 2, 2, 2, 1, 1, 2, 2, 1, 1, 2, 2, 1]	81.1
		✓	EA searched	[1, 2, 1, 2, 2, 1, 2, 2, 1, 2, 2, 1, 2, 2, 1, 2, 2, 1, 2, 1, 2, 2, 2]	81.2
		✓	Random	[2, 0, 0, 2, 2, 0, 0, 2, 2, 2, 0, 2, 2, 0, 2, 0, 0, 0, 2, 2, 0, 2, 0, 0]	64.9
		✓	Randomly searched	[2, 2, 0, 2, 2, 0, 2, 2, 2, 2, 0, 0, 2, 2, 0, 0, 0, 0, 2, 0, 0, 2, 0, 0, 2]	65.4
		✓	EA searched	[0, 0, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 0, 0, 0, 0, 0, 0, 2, 0, 0, 0, 2]	65.7
		✓	Random	[0, 1, 1, 0, 1, 1, 0, 0, 1, 0, 0, 0, 0, 0, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1]	79.7
		✓	Randomly searched	[0, 1, 0, 1, 0, 1, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 1, 0, 1, 1, 0, 1, 0]	79.9
		✓	EA searched	[0, 1, 0, 0, 0, 1, 1, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1]	79.8
		✓	Random	[1, 1, 0, 0, 0, 0, 1, 1, 0, 1, 1, 0, 2, 2, 2, 1, 0, 1, 0, 1, 0, 2, 2, 1]	80.6
LV-BERT-small	✓	✓	EA searched	[0, 0, 1, 2, 2, 1, 0, 2, 2, 1, 0, 0, 1, 2, 0, 2, 1, 0, 2, 0, 1, 1, 2, 1]	<b>81.8</b>

Table 7: Architectures of different models and their performance on GLUE development set. In Architecture column, 0, 1, and 2 denote dynamic convolution, self-attention, and feed-forward layers respectively \* denotes methods implemented by us for language pre-training.

determine whether a pair of questions are semantically equivalent or not.

**STS** The Semantic Textual Similarity Benchmark (Cer et al., 2017) is a collection of sentence pairs with human-annotated similarity score on a 1-5 scale.

**WNLI** Winograd NLI (Levesque et al., 2012) is a small dataset for natural language inference. However, there are issues with the construction of this dataset<sup>4</sup>. Therefore, this dataset is excluded in this paper for comparison as BERT (Devlin et al., 2019) etc.

### B.1.1 SQuAD dataset

The Stanford Question Answering Dataset (SQuAD 1.1) (Rajpurkar et al., 2016) is a dataset of more than 100K questions which all can be answered by locating a span of text from the corresponding context passage. Besides this data, the upgraded version SQuAD 2.0 (Rajpurkar et al., 2018) supplements it with over 50K unanswerable questions.

## C Pre-training Details

For supernet, We pre-train it for 2M steps with hyperparameters listed in Table 5, using Masked Language Modeling (MLM) pre-training objective (Devlin et al., 2019). This objective masks 15% input tokens that require the model to predict. The reason to use this objective is that the MLM valida-

tion accuracy can reflect the performance of models on downstream tasks (Lan et al., 2020).

For pre-training LV-BERTs and other compared baselines like DynamicConv (Wu et al., 2018) and Sandwich (Press et al., 2020) from scratch, we utilize Replaced Token Detection (RTE) pre-training objective (Clark et al., 2020). This objective employs a small generator to predict masked tokens and utilize a larger discriminator to determine predicted tokens from the generator are the same as original ones or not. RTE can help save computation cost but achieve good performance (Clark et al., 2020). We pre-train the models for 1M steps, mostly using the same hyperparameters as ELECTRA (Clark et al., 2020). We set the pre-training sequence length 128 that can help us save computation cost. For downstream task SQuAD 1.1/2.0 that needs longer input sequence length, we pre-train more 10% steps with the sequence length of 512 to learn the position embedding before fine-tuning. The hyperparameters are listed in Table 5.

## D Fine-tuning Details

For fine-tuning on downstream tasks, most of the hyperparameters are the same as ELECTRA (Clark et al., 2020). See Table 6.

## E Searched Architectures

The different searched architectures are listed in Table 7.

<sup>4</sup><https://gluebenchmark.com/faq>