

AdaBERT: Task-Adaptive BERT Compression with Differentiable Neural Architecture Search

Daoyuan Chen, Yaliang Li, Minghui Qiu, Zhen Wang, Bofang Li,
Bolin Ding, Hongbo Deng, Jun Huang, Wei Lin, Jingren Zhou
Alibaba Group

{daoyuanchen.cdy, yaliang.li, minghui.qmh, jones.wz, bofang.lbf}@alibaba-inc.com
{bolin.ding, dhb167148, huangjun.hj, weilin.lw, jingren.zhou}@alibaba-inc.com

Abstract

Large pre-trained language models such as BERT have shown their effectiveness in various natural language processing tasks. However, the huge parameter size makes them difficult to be deployed in real-time applications that require quick inference with limited resources. Existing methods compress BERT into small models while such compression is task-independent, i.e., the same compressed BERT for all different downstream tasks. Motivated by the necessity and benefits of task-oriented BERT compression, we propose a novel compression method, AdaBERT, that leverages differentiable Neural Architecture Search to automatically compress BERT into task-adaptive small models for specific tasks. We incorporate a task-oriented knowledge distillation loss to provide search hints and an efficiency-aware loss as search constraints, which enables a good trade-off between efficiency and effectiveness for task-adaptive BERT compression. We evaluate AdaBERT on several NLP tasks, and the results demonstrate that those task-adaptive compressed models are 12.7x to 29.3x faster than BERT in inference time and 11.5x to 17.0x smaller in terms of parameter size, while comparable performance is maintained.

1 Introduction

Nowadays, pre-trained contextual representation encoders, such as ELMo (Peters et al., 2018), GPT (Radford et al., 2018), BERT (Devlin et al., 2019), and RoBERTa (Liu et al., 2019c), have been widely adopted in a variety of Natural Language Processing (NLP) tasks (Wang et al., 2019a). Despite their effectiveness, these models are built upon large-scale datasets and they usually have parameters in the billion scale. For example, the BERT-base and BERT-large models are with 109M and 340M parameters respectively. This makes it difficult to deploy such large-scale models in real-time appli-

cations that have tight constraints on computation resource and inference time.

To fulfill the deployment in real-time applications, recent studies compress BERT into a relatively small model to reduce the computational workload and accelerate the inference time. BERT-PKD (Sun et al., 2019) distills BERT into a small Transformer-based model that mimics intermediate layers from original BERT. TinyBERT (Jiao et al., 2019) uses two-stage knowledge distillation and mimics attention matrices and embedding matrices of BERT. And in Michel et al. (2019), the authors propose a method to iteratively prune the redundant attention heads of BERT.

However, these existing studies compress BERT into a task-independent model, i.e., the same compressed BERT model for all different tasks. Recall that BERT learns various knowledge from the large-scale corpus, while only certain parts of the learned knowledge are needed for a specific downstream task (Tenney et al., 2019). Further, Jawahar et al. (2019); Liu et al. (2019b) show that different hidden layers of BERT learned different levels of linguistic knowledge, and Voita et al. (2019) demonstrate that the importance degrees of attention heads in BERT varies for different tasks. All these findings shed light on the task-adaptive BERT compression: different NLP tasks use BERT in different ways, and it is necessary to compress large-scale models such as BERT for specific downstream tasks respectively. By doing so, the task-adaptive compressed BERT can remove task-specific redundant parts in original large-scale BERT, which leads to better compression and faster inference.

Motivated by this, we propose a novel **Adaptive BERT** compression method, *AdaBERT*, that leverages differentiable Neural Architecture Search (NAS) to automatically compress BERT into task-adaptive small models for specific tasks. We incorporate a task-oriented knowledge distillation loss that depends on the original BERT model

to provide search hints, as well as an efficiency-aware loss based on network structure as search constraints. These two loss terms work together and enable the proposed compression method to achieve a good trade-off between efficiency and effectiveness for different downstream tasks. To be more specific, we adopt a lightweight CNN-based search space and explicitly model the efficiency metrics with respect to searched architectures, which has not been considered in previous BERT compression studies. Further, we hierarchically decompose the learned general knowledge of BERT into task-oriented useful knowledge with a set of probe models for knowledge distillation loss, such that the architecture search space can be reduced into a small task-oriented sub-space. Finally, by relaxing the discrete architecture parameters into continuous distribution, the proposed method can efficiently find task-adaptive compression structures through the gradient-based optimization.

To evaluate the proposed method, we compress BERT for several NLP tasks including sentiment classification, entailment recognition, and semantic equivalence classification. Empirical results on six datasets show that the proposed compression method can find task-adaptive compression models that are 12.7x to 29.3x faster than BERT in inference time and 11.5x to 17.0x smaller than BERT in terms of parameter size while maintaining comparable performance.

2 Methodology

2.1 Overview

As shown in Figure 1, we aim to compress a given large BERT model into an effective and efficient task-adaptive model for a specific task. The structures of the compressed models are searched in a differentiable manner with the help of task-oriented knowledge from the large BERT model while taking the model efficiency into consideration.

Formally, let’s denote a BERT model fine-tuned on the target data \mathcal{D}_t as $BERT_t$, an architecture searching space as \mathcal{A} . Our task is to find an optimal architecture $\alpha \in \mathcal{A}$ by minimizing the following loss function:

$$\mathcal{L} = (1 - \gamma)\mathcal{L}_{CE}(\alpha, w_\alpha, \mathcal{D}_t) + \gamma\mathcal{L}_{KD}(\alpha, w_\alpha, BERT_t) + \beta\mathcal{L}_E(\alpha), \quad (1)$$

where w_α is the trainable network weights of the architecture α (e.g., weights of a feed forward layer), \mathcal{L}_{CE} , \mathcal{L}_{KD} , and \mathcal{L}_E are losses for the target task,

task-oriented knowledge distillation and efficiency respectively. Specifically, \mathcal{L}_{CE} is the cross-entropy loss w.r.t. labels from the target data \mathcal{D}_t , \mathcal{L}_{KD} is the task-oriented knowledge distillation (KD) loss that provides hints to find suitable structures for the task, and \mathcal{L}_E is the efficiency-aware term to provide constraints to help search lightweight and efficient structures. γ and β are hyper-parameters to balance these loss terms.

In the following, we first introduce the architecture search space \mathcal{A} , then present the task-oriented KD loss \mathcal{L}_{KD} and efficiency-aware loss \mathcal{L}_E , and finally describe the differentiable search method.

2.2 Search Space \mathcal{A}

Most neural architecture search methods focus on cell-based micro search space (Pham et al., 2018). That is, the searching target is a cell and the network architecture is stacked by the searched cell over pre-defined K_{max} layers, where the cell structure parameter α_c is shared for all layers. In this work, we consider a *macro search space* over the entire network to enhance the structure exploration. Specifically, besides for searching shared cell parameter α_c for stacking, we also search the number of stacking layers $K \in [1, 2, \dots, K_{max}]$. The K is crucial for finding a trade-off between model expressiveness and efficiency, as a larger K leads to higher model capacity but slower inference.

As depicted in Figure 2, the searched cell is represented as a directed acyclic graph. Each node within the cell indicates a latent state h and the edge from node i to node j indicates operation $o_{i,j}$ that transforms h_i to h_j . For the cell at k -th layer ($k > 1$), we define two input nodes c_{k-2} and c_{k-1} as layer-wise residual connections, and an output node c_k that is obtained by attentively summarized over all the intermediate nodes. For the cell at first layer, node 0 and node 1 are task-dependent input embeddings. Formally, let’s denote \mathcal{O} as the set of candidate operations. We assume a topological order among N intermediate nodes, i.e., $o_{i,j} \in \mathcal{O}$ exists when $i < j$ and $j > 1$, and the search space α can be formalized as:

$$\alpha = \{K, \alpha_c\}, \quad K \leq K_{max}, \quad (2)$$

$$\alpha_c = [o_{0,2}, o_{1,2}, \dots, o_{i,j}, \dots, o_{N+1,N+2}].$$

2.3 Task-oriented Knowledge Distillation

To encourage the learned structure to be suitable for the target task, we introduce the task-oriented knowledge distillation loss, denoted as \mathcal{L}_{KD} in Equation (1), to guide the structure search process.

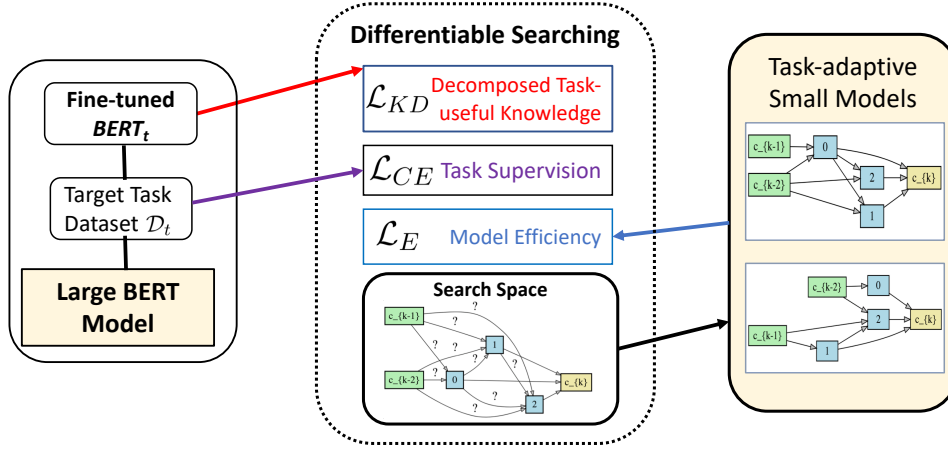


Figure 1: The overview of AdaBERT. By considering task-useful knowledge from original BERT, as well as model efficiency, our method searches suitable small models for target tasks in a differentiable way.

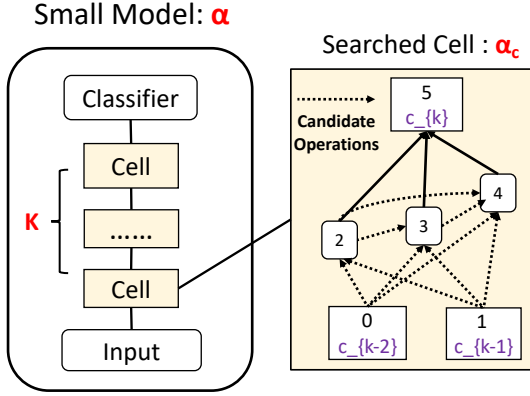


Figure 2: Search space including stacked layers and stacked cells.

Task-useful Knowledge Probe We leverage a set of probe classifiers to hierarchically decompose the task-useful knowledge from the teacher model $BERT_t$, and then distill the knowledge into the compressed model. Specifically, we freeze the parameters of $BERT_t$, and train a Softmax probe classifier for each hidden layer w.r.t. the ground-truth task labels. In total we have J classifiers, ($J = 12$ in BERT-base), and the classification logits of j -th classifier can be regarded as the learned knowledge from j -th layer. Given an input instance m , denote $C_{j,m}^T$ as the hidden representation from the j -th layer of $BERT_t$, $C_{i,m}^S$ as the attentively summed hidden state on i -th layer of the compressed student model, we distill the task-useful knowledge (classification logits) as:

$$\mathcal{L}_{KD}^{i,m} = -P_j^T(C_{j,m}^T) \cdot \log(P_i^S(C_{i,m}^S)/T), \quad (3)$$

where T is the temperature value, P_j^T is the j -th teacher probe classifier, P_i^S is the trainable student probe on i -th layer of the compressed model. Clearly $\mathcal{L}_{KD}^{i,m}$ represents the mastery degree of teacher knowledge from the i -th layer of the compressed model, i.e., how similar are the logits of the two models predict for the same input m . Here we set j to be proportional to the layers of the two models, i.e., $j = \lceil i \times J/K \rceil$, such that the compressed model can learn the decomposed knowledge smoothly and hierarchically.

Attentive Hierarchical Transfer The usefulness of each layer of BERT is varied for different tasks as shown in (Liu et al., 2019b). Here we attentively combine the decomposed knowledge for all layers as:

$$\mathcal{L}_{KD} = \sum_{m=0}^M \sum_{i=1}^K w_{i,m} \cdot \mathcal{L}_{KD}^{i,m}, \quad (4)$$

$$w_{i,m} = \frac{\exp[y_m \cdot \log P_j^T(C_{j,m}^S)]}{\sum_{i'} \exp[y_m \cdot \log P_{j'}^T(C_{j',m}^S)]},$$

where M is the total number of training instances, y_m is the label of instance m . $w_{i,m}$ is set as the normalized weight according to the negative cross-entropy loss of the teacher probe P_j^T , so that probe classifiers making preciser predictions (smaller loss) gain higher weights. Besides, to enrich task-useful knowledge, we perform data augmentation on target task datasets with the augmentation process used in (Jiao et al., 2019), which leverages BERT and GloVe (Pennington et al., 2014) to replace words in original texts.

2.4 Efficiency-Aware Loss

Recall that we aim to compress the original BERT model into efficient compressed models. To achieve this, we further incorporate model efficiency into loss function from two aspects, i.e., parameter size and inference time. To be specific, for searched architecture α_c and K , we define the efficiency-aware loss in Equation (1) as:

$$\mathcal{L}_E = \frac{K}{K_{max}} \sum_{o_{i,j} \in \alpha_c} SIZE(o_{i,j}) + FLOPs(o_{i,j}), \quad (5)$$

where K_{max} is the pre-defined maximum number of layers, $SIZE(\cdot)$ and $FLOPs(\cdot)$ are the normalized parameter size and the number of floating point operations (FLOPs) for each operation. The sum of FLOPs of searched operations serves as an approximation to the actual inference time of the compressed model.

2.5 Differentiable Architecture Searching

A major difference between the proposed method and existing BERT compression methods is that the proposed AdaBERT method seeks to find task-adaptive structures for different tasks. Now we will discuss the task-adaptive structure searching via a differentiable structure search method with the aforementioned loss terms.

2.5.1 Search Space Setting

Before diving into the details of the search method, we first illustrate the search space for our method in Figure 2. To make it easy to stack cells and search for network layers, we keep the same shapes for input and output nodes of each layer. In the first layer, we adopt different input settings for single text tasks such as sentiment classification and text pair tasks such as textual entailment. As shown in Figure 2, the inputs c_{k-2} and c_{k-1} are set as the same text input for single text tasks, and set as two input texts for text pair tasks. This setting helps to explore self-attention or self-interactions for single text task, and pair-wise cross-interactions for text pair tasks.

For the candidate operations in the cell, we adopt lightweight CNN-based operations as they are effective in NLP tasks (Kim, 2014; Bai et al., 2018), and CNN-based operations have shown inference speed superiority over RNN-based models and self-attention based models (Shen et al., 2018; Chia et al., 2018) due to the fact that they are parallel-friendly operations. Specifically, the candidates

operation set \mathcal{O} include *convolution*, *pooling*, *skip* (identity) connection and *zero* (discard) operation. For *convolution* operations, both 1D standard convolution and dilated convolution with kernel size $\{3, 5, 7\}$ are included, among which the dilated convolution can be used to enhance the capability of capturing long-dependency information and each convolution is applied as a Relu-Conv-BatchNorm structure. *Pooling* operations include averaging pooling and max pooling with kernel size 3. The *skip* and *zero* operations are used to build residual connection and discard operation respectively, which are helpful to reduce network redundancy. Besides, for both convolution and pooling operations, we apply the ‘‘SAME’’ padding (Dumoulin and Visin, 2016) to make the output length as the same as the input.

2.5.2 Search Algorithm

Directly optimizing the overall loss \mathcal{L} in Equation (1) by brute-force enumeration of all the candidate operations is impossible, due to the huge search space with combinatorial operations and the time-consuming training on its w_α w.r.t. α . Here we solve this problem by modeling searched architecture $\{K, o_{i,j}\}$ as discrete variables that obey discrete probability distributions $P_K = [\theta_1^K, \dots, \theta_{K_{max}}^K]$ and $P_o = [\theta_1^o, \dots, \theta_{|\mathcal{O}|}^o]$. K and $o_{i,j}$ are thus modeled as one-hot variables and sampled from layer range $[1, K_{max}]$ and candidate operation set \mathcal{O} respectively. However, \mathcal{L} is non-differentiable as the discrete sampling process makes the gradients cannot propagate back to the learnable parameters P_K and P_o . Inspired by (Xie et al., 2019; Wu et al., 2019), we leverage Gumbel Softmax technique (Jang et al., 2017; Maddison et al., 2017) to relax the categorical samples into continuous sample vectors $y^K \in R^{K_{max}}$ and $y^o \in R^{|\mathcal{O}|}$ as:

$$y_i^K = \frac{\exp[(\log(\theta_i^K) + g_i)/\tau]}{\sum_{j=1}^{K_{max}} \exp[(\log(\theta_j^K) + g_j)/\tau]}, \quad (6)$$

$$y_i^o = \frac{\exp[(\log(\theta_i^o) + g_i)/\tau]}{\sum_{j=1}^{|\mathcal{O}|} \exp[(\log(\theta_j^o) + g_j)/\tau]},$$

where g_i is a random noise drawn from Gumbel(0, 1) distribution, τ is the temperature parameter to control the degree of approximating Gumbel Softmax to *argmax*, i.e., as τ approaches 0, the samples become one-hot. By this way, y^K and y^o are differentiable proxy variables to discrete samples, and then we can efficiently optimize \mathcal{L}

directly using gradient-based optimizers. Specifically, we use one-hot sample vectors $\text{argmax}(y^K)$ and $\text{argmax}(y^o)$ in the forward stage while use continuous y^K and y^o in the back-propagation stage, which is called Straight-Through Estimator (Bengio et al., 2013) and makes the forward process in training consistent with testing.

Note that the optimization of the overall loss \mathcal{L} can be regarded as learning a parent graph defined by the search space α as described in Section 2.2, whose weights of candidate operations w_α and architecture distribution P_α are trained simultaneously. In the training stage, the randomness introduced by P_α enhances the exploring for suitable contextual encoders that mimic task-specific teacher $BERT_t$ under resource constraints. After the training of the parent graph, we can derive an efficient and task-adaptive child graph by applying argmax on P_α as the compressed model. Also note that in \mathcal{L} , the knowledge distillation loss \mathcal{L}_{KD} provides regularization for the architecture sampling on P_α , while the efficiency-aware loss \mathcal{L}_E promotes sparse structures that make the model compact and efficient.

3 Experiments

In this section, we evaluate the proposed AdaBERT method from the following aspects:

- (1) How does AdaBERT perform comparing to state-of-the-art BERT compression methods? (Section 3.2.1)
- (2) Can AdaBERT search task-adaptive network structures? (Section 3.2.2)
- (3) How do the key components of AdaBERT such as knowledge losses and efficiency-aware loss affect its performance? (Section 3.3)

3.1 Setup

Datasets We evaluate the proposed AdaBERT method on six datasets from GLUE (Wang et al., 2019a) benchmark. Specifically, we consider three types of NLP tasks, namely sentiment classification, semantic equivalence classification, and entailment recognition. *SST-2* (Socher et al., 2013) is adopted for sentiment classification, whose goal is to label movie reviews as positive or negative. *MRPC* (Dolan and Brockett, 2005) and *QQP* (Chen et al., 2018) are adopted for semantic equivalence classification, whose sentence pairs are extracted from news sources and online question-answering website respectively. *MNLI* (Williams et al., 2018),

QNLI (Wang et al., 2019a) and *RTE* (Bentivogli et al., 2009) are adopted for textual entailment recognition, whose premise-hypothesis pairs vary in domains and scales.

Baselines We compare the proposed AdaBERT method with several state-of-the-art BERT compression methods including BERT-PKD (Sun et al., 2019), DistilBERT (Sanh et al., 2019), TinyBERT (Jiao et al., 2019) and BiLSTM_{SOFT} (Tang et al., 2019). Since our approach searches architecture from a space including the number of network layers K , here we also compare with several different versions of these baselines with the different number of layers for a comprehensive comparison.

AdaBERT Setup We fine-tune the BERT-base model (Devlin et al., 2019) on the six adopted datasets respectively as teacher models for knowledge distillation. For input text, following (Lan et al., 2020), we factorize the WordPiece embedding of $BERT_t$ into smaller embeddings whose dimension is 128 and set the max input length as 128. For the optimization of operation parameters, we adopt SGD with momentum as 0.9 and learning rate from $2e-2$ to $5e-4$ scheduled by cosine annealing. For the optimization of architecture distribution parameters P_α , we use Adam with a learning rate of $3e-4$ and weight decay of $1e-3$. For AdaBERT, we set $\gamma = 0.8$, $\beta = 4$, $T = 1$, inner node $N = 3$ and search layer $K_{max} = 8$. We search P_α for 80 epochs and derive the searched structure with its trained operation weights. The searching process can be finished within 0.5h to 10h for different tasks using 4 V100 GPUs based on our PyTorch implementation.

3.2 Overall Results

3.2.1 Compression Results

The compression results on the six adopted datasets, including parameter size, inference speedup and classification accuracy, are summarized in Table 1. Detailed results of AdaBERT method for different tasks are reported in Table 2.

Overall speaking, on all the evaluated datasets, the proposed AdaBERT method achieves significant efficiency improvement while maintaining comparable performance. Compared to the BERT₁₂-T, the compressed models are 11.5x to 17.0x smaller in parameter size and 12.7x to 29.3x faster in inference speed with an average performance degradation of 2.79%. This demonstrates

Method	# Params	Inference Speedup	SST-2	MRPC	QQP	MNLI	QNLI	RTE	Average
BERT ₁₂	109M	1x	93.5	88.9	71.2	84.6	90.5	66.4	82.5
BERT ₁₂ -T			93.3	88.7	71.1	84.8	90.4	66.1	82.4
BERT ₆ -PKD	67.0M	1.9x	92.0	85.0	<u>70.7</u>	81.5	89.0	65.5	80.6
BERT ₃ -PKD	45.7M	3.7x	87.5	80.7	68.1	76.7	84.7	58.2	76.0
DistilBERT ₄	52.2M	3.0x	91.4	82.4	68.5	78.9	85.2	54.1	76.8
TinyBERT ₄	14.5M	<u>9.4x</u>	92.6	86.4	71.3	82.5	<u>87.7</u>	62.9	80.6
BiLSTM _{SOFT}	<u>10.1M</u>	7.6x	90.7	-	68.2	73.0	-	-	-
AdaBERT	6.4M ~ 9.5M	12.7x ~ 29.3x	91.8	<u>85.1</u>	<u>70.7</u>	<u>81.6</u>	86.8	<u>64.4</u>	<u>80.1</u>

Table 1: The compression results including model efficiency and accuracy from the GLUE test server, and the MNLI result is evaluated for matched-accuracy (MNLI-m). BERT₁₂ indicates the results of the fine-tuned BERT-base from (Devlin et al., 2019) and BERT₁₂-T indicates the results of the fine-tuned BERT-base in our implementation. The results of BERT-PKD are from (Sun et al., 2019), the results of DistilBERT₄ and TinyBERT₄ are from (Jiao et al., 2019), and the results of BiLSTM_{SOFT} is from (Tang et al., 2019). The number of model parameters includes the embedding size, and the inference time is tested with a batch size of 128 over 50,000 samples. The bold numbers and underlined numbers indicate the best and the second-best performance respectively.

Task	K	# Params	Inference Speedup
SST-2	3	6.4M	29.3x
MRPC	4	7.5M	19.2x
QQP	5	8.2M	16.4x
MNLI	7	9.5M	12.7x
QNLI	5	7.9M	18.1x
RTE	6	8.6M	15.5x

Table 2: The number of layers, parameters and inference speedups of searched structures by AdaBERT for different tasks.

the effectiveness of AdaBERT to compress BERT into task-dependent small models.

Comparing with different Transformers-based compression baselines, the proposed AdaBERT method is 1.35x to 3.12x faster than the fastest baseline, TinyBERT₄, and achieves comparable performance with the two baselines that have the best averaged accuracy, BERT₆-PKD and TinyBERT₄. Further, as shown in Table 2, AdaBERT searches suitable layers and structures for different tasks, e.g., the searched structure for SST-2 task is lightweight since this task is relatively easy and a low model capacity is enough to mimic task-useful knowledge from the original BERT. This observation confirms that AdaBERT can automatically search small compressed models that adapt to downstream tasks.

Comparing with another structure-heterogeneous method, BiLSTM_{SOFT}, AdaBERT searches CNN-based models and achieves much

better improvements, especially on the MNLI dataset. This is because AdaBERT searches different models for different downstream tasks (as Table 2 shows), and adopts a flexible searching manner to find suitable structures for different tasks while BiLSTM_{SOFT} uses a Siamese structure for all different tasks. This shows the flexibility of AdaBERT to derive task-oriented compressed models for different tasks, and we will investigate more about this in the following part.

3.2.2 Adaptiveness Study

Cross-Task Validation In order to further examine the adaptiveness of searched structures by AdaBERT, we apply the searched compression model structures across different downstream tasks. For example, the searched structure for task SST-2 (we denote this searched structure as AdaBERT-SST-2) is applied to all different tasks. For such cross-task validation, we randomly initialize the weights of each searched structure and re-train its weights using corresponding training data to ensure a fair comparison. The results of cross-task validation is summarized in Table 3.

From Table 3, we can observe that: The searched structures achieve the best performance on their original target tasks compared with other tasks, in other words, the performance numbers along the diagonal line of this table are the best. Further, the performance degradation is quite significant across different kinds of tasks (for example, applying the searched structures of sentiment classification tasks to entailment recognition task, or vice versa), while the performance degradations within the same kind

Structure \ Task	SST-2	MRPC	QQP	MNLI	QNLI	RTE
AdaBERT-SST-2	91.9	78.1	58.6	64.0	74.1	53.8
AdaBERT-MRPC	81.5	84.7	68.9	75.9	82.2	60.3
AdaBERT-QQP	81.9	84.1	70.5	76.3	82.5	60.5
AdaBERT-MNLI	82.1	81.5	66.8	81.3	86.1	63.2
AdaBERT-QNLI	81.6	82.3	67.7	79.2	87.2	62.9
AdaBERT-RTE	82.9	81.1	66.5	79.8	86.0	64.1
Random	80.4 ± 4.3	79.2 ± 2.8	61.8 ± 4.9	69.7 ± 6.7	78.2 ± 5.5	55.3 ± 4.1

Table 3: Accuracy comparison on the dev sets with the searched compression structures applying to different tasks. For Random, 5-times averaging results with standard deviations are reported.

of tasks (for example, MRPC and QQP for semantic equivalence classification) are relatively small, since they have the same input format (i.e., a pair of sentences) and similar targets. All these observations verify that AdaBERT can search task-adaptive structures for different tasks with the guidance of task-specific knowledge.

We also conduct another set of experiment: for each task, we randomly sample a model structure without searching, and then train such structures on the corresponding datasets. From the last row of Table 3, we can see that the randomly sampled structures perform worse than the searched structures and their performances are not stable. This shows the necessity of the proposed adaptive structure searching.

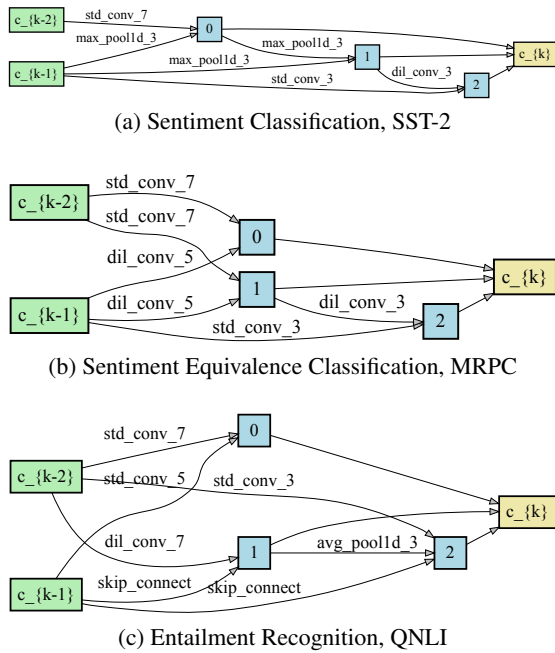


Figure 3: The basic cells of searched structures for three kinds of NLP tasks.

Architecture Study In order to examine the adaptiveness of searched structures, we also visualize the basic cell of searched structures for different tasks in Figure 3.

By comparing the structure searched for single-text task such as sentiment classification (Figure 3a) with the one for sentiment equivalence classification task (Figure 3b), we can find that the former searched structure has more aggregation operations (*max_pool1d_3*) and smaller feature filters (*std_conv_3* and *dil_conv_3*) since encoding local features are good enough for the binary sentiment classification task, while the latter searched structure has more interactions between the two input nodes as it deals with text pairs.

For the text-pair tasks, compared with sentiment equivalence classification task (Figure 3b), the searched structure for the entailment recognition task (Figure 3c) has more diverse operations such as *avg_pool_3* and *skip_connect*, and more early interactions among all the three subsequent nodes (node 0, 1 and 2). This may be justified by the fact that the textual entailment requires different degrees of reasoning and thus the searched structure has more complex and diverse interactions.

From the above comparisons among the searched structures for different tasks, it is confirmed that the proposed AdaBERT method can search task-adaptive structure for BERT compression. In the next part of this section, we conduct ablation studies to examine how knowledge losses and efficiency-aware loss affect the performance of AdaBERT.

3.3 Ablation Study

Knowledge Losses We first evaluate the effect of knowledge distillation (\mathcal{L}_{KD}) and supervised label knowledge (L_{CE}) by conducting experiments

on different tasks. The results are summarized in Table 4.

	SST-2	MRPC	QNLI	RTE
Base-KD	86.6	77.2	82.0	56.7
+ Probe	88.4	78.7	83.3	58.1
+ DA	91.4	83.9	86.5	63.2
+ \mathcal{L}_{CE} (All)	91.9	84.7	87.2	64.1

Table 4: The effect of knowledge loss terms.

The Base-KD is a naive knowledge distillation version in which only the logits of the last layer are distilled without considering hidden layer knowledge and supervised label knowledge. By incorporating the probe models, the performance (line 2 in Table 4) is consistently improved, indicating the benefits from hierarchically decomposed task-oriented knowledge. We then leverage Data Augmentation (DA) to enrich task-oriented knowledge and this technique also improves performance for all tasks, especially for tasks that have a limited scale of data (i.e., MRPC and RTE). DA is also adopted in existing KD-based compression studies (Tang et al., 2019; Jiao et al., 2019).

When taking the supervised label knowledge (\mathcal{L}_{CE}) into consideration, the performance is further boosted, showing that this term is also important for AdaBERT by providing focused search hints.

Efficiency-aware Loss Last, we test the effect of efficiency-aware loss \mathcal{L}_E by varying its corresponding coefficient, including the standard case ($\beta = 4$), without efficiency constraint ($\beta = 0$), and strong efficiency constraint ($\beta = 8$).

	SST-2	MRPC	QNLI	RTE
$\beta = 0$	91.8 (7.5M)	84.5 (7.8M)	87.1 (8.3M)	63.9 (9.1M)
$\beta = 4$	91.9 (6.4M)	84.7 (7.5M)	87.2 (7.9M)	64.1 (8.6M)
$\beta = 8$	91.3 (5.3M)	84.2 (6.4M)	86.4 (7.1M)	63.3 (7.8M)

Table 5: The effect of efficiency loss term.

The model performance and corresponding model size are reported in Table 5. On the one hand, removing the efficiency-aware loss ($\beta = 0$) leads to the increase in model parameter size, on

the other hand, a more aggressive efficiency preference ($\beta = 8$) results in the small model size but degraded performance, since a large β encourages the compressed model to adopt more lightweight operations such as *zero* and *skip* which hurt the performance. A moderate efficiency constraint ($\beta = 4$) provides a regularization, guiding the AdaBERT method to achieve a trade-off between the small parameter size and the good performance.

4 Related Work

Pretrained Language Model Compression Existing efforts to compress pre-trained language models such as BERT can be broadly categorized into four lines: knowledge distillation (Hinton et al., 2015), parameter sharing (Ullrich et al., 2017), pruning (Cheng et al., 2017) and quantization (Han et al., 2016).

For knowledge distillation based methods, in Tang et al. (2019), BERT is distilled into a simple BiLSTM and achieves comparable results with ELMo. A dual distillation is proposed to reduce the vocabulary size and the embedding size in Zhao et al. (2019). PKD-BERT (Sun et al., 2019) and DistilBERT (Sanh et al., 2019) distill BERT into shallow Transformers in the fine-tune stage and the pre-train stage respectively. A concurrent work, TinyBERT (Jiao et al., 2019), further distills BERT with a two-stage knowledge distillation for hidden attention matrices and embedding matrices. For parameter sharing based methods, the multi-head attention is compressed into a tensorized Transformer in Ma et al. (2019). Another concurrent work, AIBERT (Lan et al., 2020) leverages cross-layer parameter sharing to speed up the training and achieves new state-of-the-art results with 233M parameters. Different from these existing methods, the proposed method AdaBERT incorporates both task-oriented knowledge distillation and efficiency factor to achieve competing results with much faster inference time and smaller model size. More importantly, the proposed method AdaBERT automatically compresses BERT into task-adaptive small structures instead of a task-independent structure in existing methods.

For pruning and quantization based methods, the majority of attention heads in BERT are iteratively pruned without seriously affecting performance in Michel et al. (2019). In Wang et al. (2019b), low-rank factorization is adopted to prune a BERT variation, RoBERTa (Liu et al., 2019c), for sev-

eral classification tasks. Q8BERT (Zafrir et al., 2019) quantizes matrix multiplication operations in BERT into 8-bit operations, while Q-BERT (Shen et al., 2019) quantizes BERT with Hessian based mix-precision. These methods and the proposed method AdaBERT compress BERT from different aspects that are complementary, that is, one can first distill BERT into a small model, and then further prune or quantize the small model.

Neural Architecture Search Automatically discovering neural network architecture gains increasing attention recently. Early NAS methods search state-of-the-art architectures based on reinforcement learning (Zoph and Le, 2017) and evolution (Real et al., 2019) while they are computationally expensive. Recent NAS methods significantly speed up the search and evaluation stages by architecture parameter sharing such as ENAS (Pham et al., 2018), gradient descent on differentiable searching objectives such as DARTS (Liu et al., 2019a) and SNAS (Xie et al., 2019), and hardware-aware optimization such as AMC (He et al., 2018) and FBNet (Wu et al., 2019). Different from them, we incorporate knowledge distillation loss by probing task-useful knowledge as search hints. To the best of our knowledge, this is the first work to compress large language models with NAS.

5 Conclusion

In this work, motivated by the strong need to compress BERT into small and fast models, we propose AdaBERT, an effective and efficient model that adaptively compresses BERT for various downstream tasks. By leveraging Neural Architecture Search, we incorporate two kinds of losses, task-useful knowledge distillation loss depending on the original BERT and efficiency-aware loss based on the searched structure, such that the task-suitable structures of compressed BERT can be automatically and efficiently found using gradient information. We evaluate the proposed AdaBERT on six datasets involving three kinds of NLP tasks. Extensive experiments demonstrate that AdaBERT achieves comparable performance while significantly improves the efficiency by 12.7x to 29.3x speedup in inference time and 11.5x to 17.0x compression ratio in parameter size. Further, the adaptiveness study confirms that the proposed AdaBERT can find different models varying in model efficiencies and architectures that are suitable for

different downstream tasks.

References

- Shaojie Bai, J Zico Kolter, and Vladlen Koltun. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2009. The fifth pascal recognizing textual entailment challenge. In *TAC*.
- Zihan Chen, Hongbo Zhang, Xiaoji Zhang, and Leqi Zhao. 2018. Quora question pairs.
- Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. 2017. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*.
- Yew Ken Chia, Sam Witteveen, and Martin Andrews. 2018. Transformer to cnn: Label-scarce distillation for efficient text classification. In *NIPS 2018 Workshop CDNNRIA*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186.
- William B. Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *IWP*.
- Vincent Dumoulin and Francesco Visin. 2016. A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*.
- Song Han, Huizi Mao, and William J Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *ICLR*.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. 2018. Amc: Automl for model compression and acceleration on mobile devices. In *ECCV*, pages 784–800.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *ICLR*.
- Ganesh Jawahar, Benoît Sagot, Djamé Seddah, et al. 2019. What does bert learn about the structure of language? In *ACL*, pages 3651–3657.

- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- 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 *ICLR*.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2019a. DARTS: Differentiable architecture search. In *ICLR*.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019b. Linguistic knowledge and transferability of contextual representations. In *NAACL*, pages 1073–1094.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Xindian Ma, Peng Zhang, Shuai Zhang, Nan Duan, Yuexian Hou, Ming Zhou, and Dawei Song. 2019. A tensorized transformer for language modeling. In *NeurIPS*, pages 2229–2239.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *NeurIPS*, pages 14014–14024.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL*.
- Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. 2018. Efficient neural architecture search via parameter sharing. In *ICML*, pages 4092–4101.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *arXiv*.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. 2019. Regularized evolution for image classifier architecture search. In *AAAI*, volume 33, pages 4780–4789.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *NeurIPS EMC² Workshop*.
- Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. 2019. Q-bert: Hessian based ultra low precision quantization of bert. *arXiv preprint arXiv:1909.05840*.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. 2018. Bi-directional block self-attention for fast and memory-efficient sequence modeling. In *ICLR*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. Patient knowledge distillation for bert model compression. In *EMNLP*.
- Raphael Tang, Yao Lu, Linqing Liu, Lili Mou, Olga Vechtomova, and Jimmy Lin. 2019. Distilling task-specific knowledge from bert into simple neural networks. *arXiv preprint arXiv:1903.12136*.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT rediscovers the classical NLP pipeline. In *ACL*, pages 4593–4601.
- Karen Ullrich, Edward Meeds, and Max Welling. 2017. Soft weight-sharing for neural network compression. *ICLR*.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *ACL*, pages 5797–5808.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*.
- Ziheng Wang, Jeremy Wohlwend, and Tao Lei. 2019b. Structured pruning of large language models. *arXiv preprint arXiv:1910.04732*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL*, pages 1112–1122.
- Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. 2019. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *CVPR*, pages 10734–10742.
- Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. 2019. SNAS: stochastic neural architecture search. In *ICLR*.

Ofir Zafir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. Q8bert: Quantized 8bit bert. In *NeurIPS EMC² Workshop*.

Sanqiang Zhao, Raghav Gupta, Yang Song, and Denny Zhou. 2019. Extreme language model compres-

sion with optimal subwords and shared projections. *arXiv preprint arXiv:1909.11687*.

Barret Zoph and Quoc V Le. 2017. Neural architecture search with reinforcement learning. *ICLR*.