

Enabling Lightweight Fine-tuning for Pre-trained Language Model Compression based on Matrix Product Operators

Peiyu Liu^{1,4*}, Ze-Feng Gao^{2,1*}, Wayne Xin Zhao^{1,4,5†},
Z.Y. Xie², Zhong-Yi Lu^{2†} and Ji-Rong Wen^{1,3,4}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²Department of Physics, Renmin University of China

³School of Information, Renmin University of China

⁴Beijing Key Laboratory of Big Data Management and Analysis Methods

⁵Beijing Academy of Artificial Intelligence, Beijing, 100084, China

{liupeiyust,zfgao,qingtaoxie,zlu,jrwen}@ruc.edu.cn, batmanfly@gmail.com

Abstract

This paper presents a novel pre-trained language models (PLM) compression approach based on the matrix product operator (short as MPO) from quantum many-body physics. It can decompose an original matrix into central tensors (containing the core information) and auxiliary tensors (with only a small proportion of parameters). With the decomposed MPO structure, we propose a novel fine-tuning strategy by only updating the parameters from the auxiliary tensors, and design an optimization algorithm for MPO-based approximation over stacked network architectures. Our approach can be applied to the original or the compressed PLMs in a general way, which derives a lighter network and significantly reduces the parameters to be fine-tuned. Extensive experiments have demonstrated the effectiveness of the proposed approach in model compression, especially the reduction in fine-tuning parameters (91% reduction on average). The code to reproduce the results of this paper can be found at <https://github.com/RUCAIBox/MPOP>.

1 Introduction

Recently, pre-trained language models (PLMs) (Devlin et al., 2019; Peters et al., 2018; Radford et al., 2018) have made significant progress in various natural language processing tasks. Instead of training a model from scratch, one can fine-tune a PLM to solve some specific task through the paradigm of “pre-training and fine-tuning”.

Typically, PLMs are constructed with stacked Transformer layers (Vaswani et al., 2017), involving a huge number of parameters to be learned. Though effective, the large model size makes it impractical for resource-limited devices. Therefore, there is an increasing number of studies focused

on the parameter reduction or memory reduction of PLMs (Noach and Goldberg, 2020), including parameter sharing (Lan et al., 2020), knowledge distillation (Sanh et al., 2019), low-rank approximation (Ma et al., 2019) and data quantization (Hubara et al., 2017). However, these studies mainly apply parameter reduction techniques to PLM compression, which may not be intrinsically appropriate for the learning paradigm and architecture of PLMs. The compressed parameters are highly coupled so that it is difficult to directly manipulate different parts with specific strategies. For example, most PLM compression methods need to fine-tune the whole network architecture, although only a small proportion of parameters will significantly change during fine-tuning (Liu et al., 2020).

In this paper, we introduce a novel matrix product operator (MPO) technique from quantum many-body physics for compressing PLMs (Gao et al., 2020). The MPO is an algorithm that factorizes a matrix into a sequential product of local tensors (*i.e.*, a multi way array). Here, we call the tensor right in the middle as *central tensor* and the rest as *auxiliary tensors*. An important merit of the MPO decomposition is structural in terms of information distribution: the central tensor with most of the parameters encode the core information of the original matrix, while the auxiliary tensors with only a small proportion of parameters play the role of complementing the central tensor. Such a property motivates us to investigate whether such an MPO can be applied to derive a better PLM compression approach: can we compress the central tensor for parameter reduction and update the auxiliary tensors for lightweight fine-tuning? If this could be achieved, we can derive a lighter network meanwhile reduce the parameters to be fine-tuned.

To this end, we propose an MPO-based compression approach for PLMs, called *MPOP*. It is developed based on the MPO decomposition tech-

* Authors contributed equally.

† Corresponding author.

nique (Gao et al., 2020; Pirvu et al., 2010). We have made two critical technical contributions for compressing PLMs with MPO. First, we introduce a new fine-tuning strategy that only focuses on the parameters of auxiliary tensors, so the number of fine-tuning parameters can be largely reduced. We present both theoretical analysis and experimental verification for the effectiveness of the proposed fine-tuning strategy. Second, we propose a new optimization algorithm, called *dimension squeezing*, tailored for stacked neural layers. Since mainstream PLMs usually consist of multiple Transformer layers, this will produce accumulated reconstruction error by directly applying low-rank approximation with MPO at each layer. The dimension squeezing algorithm is able to gradually perform the dimension truncation in a more stable way so that it can dramatically alleviate the accumulation error in the stacked architecture.

To our knowledge, it is the first time that MPO is applied to the PLM compression, which is well suited for both the learning paradigm and the architecture of PLMs. We construct experiments to evaluate the effectiveness of the proposed compression approach for ALBERT, BERT, DistilBERT and MobileBERT, respectively, on GLUE benchmark. Extensive experiments have demonstrated the effectiveness of the proposed approach in model compression, especially dramatically reducing the fine-tuning parameters (91% reduction on average).

2 Related Work

We review the related works in three aspects.

Pre-trained Language Model Compression. Since the advent of large-scale PLMs, several variants were proposed to alleviate its memory consumption. For example, DistilBERT (Sanh et al., 2019) and MobileBERT (Sun et al., 2020c) leveraged knowledge distillation to reduce the BERT network size. SqueezeBERT (Iandola et al., 2020) and Q8BERT (Zafrir et al., 2019) adopted special techniques to substitute the operations or quantize both weights and activations. ALBERT (Lan et al., 2020) introduced cross-layer parameter sharing and low-rank approximation to reduce the number of parameters. More studies (Jiao et al., 2020; Hou et al., 2020; Liu et al., 2020; Wang et al., 2020; Khetan and Karnin, 2020; Xin et al., 2020; Pappas et al., 2020; Sun et al., 2020a) can be found in the comprehensive survey (Ganesh et al., 2020).

Tensor-based Network Compression. Tensor-based methods have been successfully applied to neural network compression. For example, MPO has been utilized to compress linear layers of deep neural network (Gao et al., 2020). Sun et al. (2020b) used MPO to compress the LSTM model on acoustic data. Novikov et al. (2015) coined the idea of reshaping weights of fully-connected layers into high-dimensional tensors and representing them in Tensor Train (TT) (Oseledets, 2011) format, which was extended to other network architectures (Garipov et al., 2016; Yu et al., 2017; Tjandra et al., 2017; Khruikov et al., 2019). Ma et al. (2019) adopted block-term tensor decomposition to compress Transformer layers in PLMs.

Lightweight Fine-tuning. In the past, lightweight fine-tuning was performed without considering parameter compression. As a typical approach, trainable modules are inserted into PLMs. For example, a “side” network is fused with PLM via summation in (Zhang et al., 2020), and adapter-tuning inserts task-specific layers (adapters) between each layer of PLMs (Houlsby et al., 2019; Lin et al., 2020; Rebuffi et al., 2017). On the contrary, several studies consider removing parameters from PLMs. For example, several model weights are ablated away by training a binary parameter mask (Zhao et al., 2020; Radiya-Dixit and Wang, 2020).

Our work is highly built on these studies, while we have a new perspective by designing the PLM compression algorithm, which enables lightweight fine-tuning. It is the first time that MPO is applied to PLM compression, and we make two major technical contributions for achieving lightweight fine-tuning and stable optimization.

3 Preliminary

In this paper, scalars are denoted by lowercase letters (*e.g.*, a), vectors are denoted by boldface lowercase letters (*e.g.*, v), matrices are denoted by boldface capital letters (*e.g.*, M), and high-order (order three or higher) tensors are denoted by boldface Euler script letters (*e.g.*, \mathcal{T}). An n -order tensor $\mathcal{T}_{i_1, i_2, \dots, i_n}$ can be considered as a multidimensional array with n indices $\{i_1, i_2, \dots, i_n\}$.

Matrix Product Operator. Originating from quantum many-body physics, matrix product operator (MPO) is a standard algorithm to factorize a matrix into a sequential product of multiple local tensors (Gao et al., 2020; Pirvu et al., 2010).

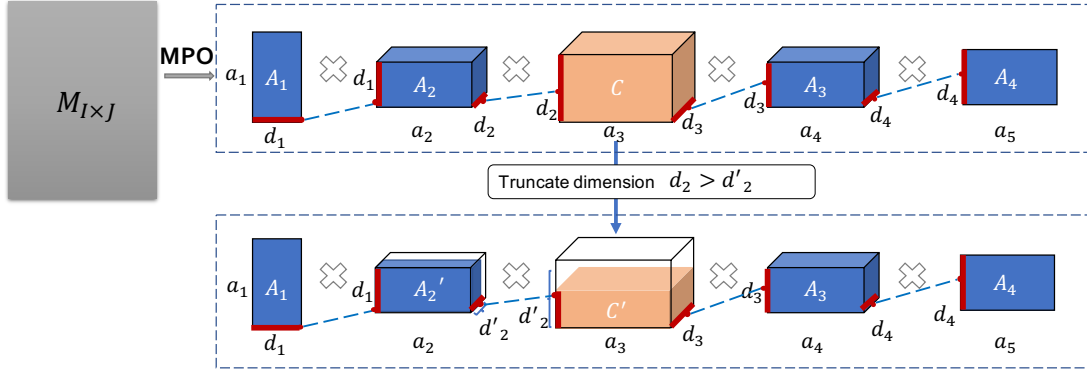


Figure 1: MPO decomposition for matrix $M_{I \times J}$ with five local tensors, where $\prod_{k=1}^n i_k = I$, $\prod_{k=1}^n j_k = J$, and $a_k = i_k \times j_k$ ($n = 5$ here). Auxiliary tensors ($\{A_i\}_{i=1}^4$) and central tensor (C) are marked in blue and orange, respectively. Dash line linking adjacent tensors denotes virtual bonds.

Formally, given a matrix $M \in \mathbb{R}^{I \times J}$, its MPO decomposition into a product of n local tensors can be represented as:

$$\text{MPO}(M) = \prod_{k=1}^n \mathcal{T}_{(k)}[d_{k-1}, i_k, j_k, d_k], \quad (1)$$

where the $\mathcal{T}_{(k)}[d_{k-1}, i_k, j_k, d_k]$ is a 4-order tensor with size $d_{k-1} \times i_k \times j_k \times d_k$ in which $\prod_{k=1}^n i_k = I$, $\prod_{k=1}^n j_k = J$ and $d_0 = d_n = 1$. We use the concept of *bond* to connect two adjacent tensors (Pirvu et al., 2010). The bond dimension d_k is defined by:

$$d_k = \min \left(\prod_{m=1}^k i_m \times j_m, \prod_{m=k+1}^n i_m \times j_m \right). \quad (2)$$

From Eq. (2), we can see that d_k is going to be large in the middle and small on both sides. We present a detailed algorithm for MPO decomposition in Algorithm 1. In this case, we refer to the tensor right in the middle as *central tensor*, and the rest as *auxiliary tensor*. Figure 1 presents the illustration of MPO decomposition, and we use $n = 5$ in this paper.

Algorithm 1 MPO decomposition for a matrix.

Input: matrix M , the number of local tensors n

Output : MPO tensor list $\{\mathcal{T}_{(k)}\}_{k=1}^n$

- 1: **for** $k = 1 \rightarrow n - 1$ **do**
 - 2: $M[I, J] \rightarrow M[d_{k-1} \times i_k \times j_k, -1]$
 - 3: $U\lambda V^T = \text{SVD}(M)$
 - 4: $U[d_{k-1} \times i_k \times j_k, d_k] \rightarrow U[d_{k-1}, i_k, j_k, d_k]$
 - 5: $\mathcal{T}^{(k)} := U$
 - 6: $M := \lambda V^T$
 - 7: **end for**
 - 8: $\mathcal{T}^{(n)} := M$
 - 9: Normalization
 - 10: **return** $\{\mathcal{T}_{(k)}\}_{k=1}^n$
-

MPO-based Low-Rank Approximation. With the standard MPO decomposition in Eq. (1), we can exactly reconstruct the original matrix M through the product of the derived local tensors. Following (Gao et al., 2020), we can truncate the k -th bond dimension d_k (see Eq. (1)) of local tensors to d'_k for low-rank approximation: $d_k > d'_k$. We can set different values for $\{d'_k\}_{k=1}^n$ to control the expressive capacity of MPO-based reconstruction. The truncation error induced by the k -th bond dimension d_k is denoted by ϵ_k (called *local truncation error*) which can be efficiently computed as:

$$\epsilon_k = \sum_{i=d_k-d'_k}^{d_k} \lambda_i, \quad (3)$$

where $\{\lambda_i\}_{i=1}^{d_k}$ are the singular values of $M[i_1 j_1 \dots i_k j_k, i_{k+1} j_{k+1} \dots i_n j_n]$.

Then the total truncation error satisfies:

$$\|M - \text{MPO}(M)\|_F \leq \sqrt{\sum_{k=1}^{n-1} \epsilon_k^2}. \quad (4)$$

The proof can be found in the supplementary materials¹. Eq. (1) indicates that the reconstruction error is bounded by the sum of the squared local truncation errors, which is easy to estimate in practice.

Suppose that we have truncated the dimensions of local tensors from $\{d_k\}_{k=1}^n$ to $\{d'_k\}_{k=1}^n$, the compression ratio introduced by quantum many-body physics (Gao et al., 2020) can be computed as follows:

$$\rho = \frac{\sum_{k=1}^n d'_{k-1} i_k j_k d'_k}{\prod_{k=1}^n i_k j_k}. \quad (5)$$

¹<https://github.com/RUCAIBox/MPOP>

Layers	(0,1e-4]	(1e-4,1e-3]	(1e-3,∞)
Word embedding	0.66	0.26	0.09
Feed-forward	0.09	0.64	0.27
Self-attention	0.09	0.64	0.27

Table 1: Distribution of parameter variations for BERT when fine-tuned on SST-2 task.

The smaller the compression ratio is, the fewer parameters are kept in the MPO representation. On the contrary, the larger the compression ratio ρ is, and the more parameters there are, and the smaller the reconstruction error is. When $\rho > 1$, it indicates the decomposed tensors have more parameters than the original matrix.

4 Approach

So far, most of pre-trained language models (PLM) are developed based on stacked Transformer layers (Vaswani et al., 2017). Based on such an architecture, it has become a paradigm to first pre-train PLMs and then fine-tunes them on task-specific data. The involved parameters of PLMs can be generally represented in the matrix format. Hence, it would be natural to apply MPO-based approximation for compressing the parameter matrices in PLMs by truncating tensor dimensions.

In particular, we propose two major improvements for MPO-based PLM compression, which can largely reduce the fine-tuning parameters and effectively improve the optimization of stacked architecture, respectively.

4.1 Lightweight Fine-tuning with Auxiliary Tensors

Due to the high coupling of parameters, previous PLM compression methods usually need to fine-tune all the parameters. As a comparison, the MPO approach decomposes a matrix into a list of local tensors, which makes it potentially possible to consider fine-tuning different parts with specific strategies. Next, we study how to perform lightweight fine-tuning based on MPO properties.

Parameter Variation from Pre-Training. To apply our solution to lightweight fine-tuning, we first conduct an empirical experiment to check the variation degree of the parameters *before* and *after* fine-tuning. Here, we adopt the standard pre-trained BERT (Devlin et al., 2019) and then fine-tune it on the SST-2 task (Socher et al., 2013). We first compute the absolute difference of the variation for

each parameter value and then compute the ratio of parameters with different variation levels. The statistical results are reported in Table 1. As we can see, most of parameters vary little, especially for the word embedding layer. This finding has also been reported in a previous studies (Khetan and Karnin, 2020). As discussed in Section 3, after MPO decomposition, the central tensor contains the majority of the parameters, while the auxiliary tensors only contain a small proportion of the parameters. Such merit inspires us to consider only fine-tuning the parameters in the auxiliary tensors while keeping the central tensor fixed during fine-tuning. If this approach was feasible, this will largely reduce the parameters to be fine-tuned.

Theoretical Analysis. Here we introduce entanglement entropy from quantum mechanics (Calabrese and Cardy, 2004) as the metric to measure the information contained in MPO bonds, which is similar to the entropy in information theory but replaces probabilities by normalized singular values produced by SVD. This will be more suitable for measuring the information of a matrix as singular values often correspond to the important information implicitly encoded in the matrix, and the importance is positively correlated with the magnitude of the singular values. Following (Calabrese and Cardy, 2004), the entanglement entropy S_k corresponding to the k -th bond can be calculated by:

$$S_k = - \sum_{j=1}^{d_k} v_j \ln v_j, \quad k = 1, 2, \dots, n-1, \quad (6)$$

where $\{v_j\}_{j=1}^{d_k}$ denote the normalized SVD eigenvalues of $\mathbf{M}[i_1 j_1 \dots i_k j_k, i_{k+1} j_{k+1} \dots i_n j_n]$. The entanglement entropy S_k is an increasing function of dimension d_k as described in (Gao et al., 2020). Based on Eq. (2), the central tensor has the largest bond dimension, corresponding to the largest entanglement entropy. This indicates that most of the information in an original matrix will be concentrated in the central tensor. Furthermore, the larger a dimension is, the larger the updating effect will be. According to (Pirvu et al., 2010), it is also guaranteed in principle that any change on some tensor will be transmitted to the whole local tensor set. Thus, it would have almost the same effect after convergence by optimizing the central tensor or the auxiliary tensors for PLMs.

Based on the above analysis, we speculate

that the affected information during fine-tuning is mainly encoded on the auxiliary tensors so that the overall variations are small. Therefore, for lightweight fine-tuning, we first perform the MPO decomposition for a parameter matrix, and then only update its auxiliary tensors according to the downstream task with the central tensor fixed. Experimental results in Section 5.2 will demonstrate that such an approach is indeed effective.

4.2 Dimension Squeezing for Stacked Architecture Optimization

Most of PLMs are stacked with multiple Transformer layers. Hence, a major problem with directly applying MPO to compressing PLMs is that the reconstruction error tends to be accumulated and amplified exponentially by the number of layers. It is thus urgent to develop a more stable optimization algorithm tailored to the stacked architecture.

Fast Reconstruction Error Estimation. Without loss of generality, we can consider a simple case in which each layer contains exactly one parameter matrix to be compressed. Assume that there are L layers, so we have L parameter matrices in total, denoted by $\{\mathbf{M}^{(l)}\}_{l=1}^L$. Let $\mathcal{C}^{(l)}$ denote the corresponding central tensor with a specific dimension $d^{(l)}$ after decomposing $\mathbf{M}^{(l)}$ with MPO. Our idea is to select a central tensor to reduce its dimension by one at each time, given the selection criterion that this truncation will lead to the least reconstruction error. However, it is time-consuming to evaluate the reconstruction error of the original matrix. According to Eq. (3), we can utilize the error bound $\sqrt{\sum_{k=1}^{n-1} \epsilon_k^2}$ for a fast estimation of the yielded reconstruction error. In this case, only one ϵ_k changes, and it can be efficiently computed via the pre-computed eigenvalues.

Fast Performance Gap Computation. At each time, we compute the performance gap before and after the dimension reduction ($d^{(l)} \rightarrow d^{(l)} - 1$) with the stop criterion. To obtain the performance \tilde{p} after dimension reduction, we need to fine-tune the truncated model on the downstream task. We can also utilize the lightweight fine-tuning strategy in Section 4.1 to obtain \tilde{p} by only tuning the auxiliary tensors. If the performance gap $\|p - \tilde{p}\|$ is smaller than a threshold Δ or the iteration number exceeds the predefined limit, the algorithm will end. Such an optimization algorithm is more stable to

optimize stacked architectures since it gradually reduces the dimension considering the reconstruction error and the performance gap. Actually, it is similar to the learning of variable matrix product states (Iblisdir et al., 2007) in physics, which optimizes the tensors one by one according to the sequence. As a comparison, our algorithm dynamically selects the matrix to truncate and is more suitable to PLMs.

Algorithm 2 presents a complete procedure for our algorithm. In practice, there are usually multiple parameter matrices to be optimized at each layer. This can be processed in a similar way: we select some matrices from one layer to optimize among all the considered matrices.

Algorithm 2 Training with dimension squeezing.

Input: : L layers with corresponding central tensor $\mathcal{C}^{(l)}$ and dimension $d^{(l)}$, threshold Δ and iteration step $iter$

- 1: Evaluate loss $p = \text{model}(Inputs)$
- 2: Perform MPO decomposition for each layer
- 3: **for** $step = 1 \rightarrow iter$ **do**
- 4: Find the layer (l^*) with the least reconstruction error
- 5: Compress MPO tensor by truncating $d^{(l^*)}$
- 6: Fine-tuning auxiliary tensors with $\{\mathcal{C}^{(l)}\}_{l=1}^L$ fixed
- 7: Evaluate loss $\tilde{p} = \text{model}(Inputs)$
- 8: **if** $\|p - \tilde{p}\| > \Delta$ **then**
- 9: **break**
- 10: **end if**
- 11: **end for**
- 12: **return** Compressed model

4.3 Overall Compression Procedure

Generally speaking, our approach can compress any PLMs with stacked architectures consisting of parameter matrices, even the compressed PLMs. In other words, it can work with the existing PLM compression methods to further achieve a better compression performance. Here, we select ALBERT (Lan et al., 2020) as a representative compressed PLM and apply our algorithm to ALBERT.

The procedure can be simply summarized as follows. First, we obtain the learned ALBERT model (complete) and perform the MPO-decomposition to the three major parameter matrices, namely word embedding matrix, self-attention matrix and feed-forward matrix². Each matrix will be decomposed into a central tensor and auxiliary tensors. Next, we perform the lightweight fine-tuning to update auxiliary tensors until convergence on downstream tasks. Then, we apply the dimension squeezing

²It introduces a parameter sharing mechanism to keep only one copy for both self-attention and feed-forward matrices.

Category	Method	Inference Time
Tucker	Tucker _(d=1) (CP)	$\mathcal{O}(nmd^2)$
	Tucker _(d>1)	$\mathcal{O}(nmd + d^n)$
MPO	MPO _(n=2) (SVD)	$\mathcal{O}(2md^3)$
	MPO _(n>2)	$\mathcal{O}(nmd^3)$

Table 2: Inference time complexities of different low-rank approximation methods. Here, n denotes the number of the tensors, m denotes $\max(\{i_k\}_{k=1}^n)$ means the largest i_k in input list, and d denotes $\max(\{d'_k\}_{k=0}^n)$ means the largest dimension d'_k in the truncated dimension list.

optimization algorithm to the three central tensors, *i.e.*, we select one matrix for truncation each time. After each truncation, we fine-tune the compressed model and further stabilize its performance. This process will repeat until the performance gap or the iteration number exceeds the pre-defined threshold.

In this way, we expect that ALBERT can be further compressed. In particular, it can be fine-tuned in a more efficient way, with only a small amount of parameters to be updated. Section 5.2 will demonstrate this.

4.4 Discussion

In mathematics, MPO-based approximation can be considered as a special low-rank approximation method. Now, we compare it with other low-rank approximation methods, including SVD (Henry and Hofrichter, 1992), CPD (Hitchcock, 1927) and Tucker decomposition (Tucker, 1966).

We present the categorization of these methods in Table 2. For PLM compression, low-rank decomposition is only performed once, while it repeatedly performs forward propagation computation. Hence, we compare their inference time complexities. Indeed, all the methods can be tensor-based decomposition (*i.e.*, a list of tensors for factorization) or matrix decomposition, and we characterize their time complexities with common parameters. Indeed, MPO and Tucker represent two categories of low-rank approximation methods. Generally, the algorithm capacity is larger with the increase of n (more tensors). When $n > 3$, MPO has smaller time complexity than Tucker decomposition. It can be seen that SVD can be considered as a special case of MPO when tensor dimension $n = 2$ and CPD is a special case of Tucker when the core tensor is the super-diagonal matrix.

In practice, we do not need to strictly follow

the original matrix size. Instead, it is easy to pad additional zero entries to enlarge matrix rows or columns, so that we can obtain different MPO decomposition results. It has demonstrated that different decomposition plans always lead to almost the same results (Gao et al., 2020). In our experiments, we adopt an odd number of local tensors for MPO decomposition, *i.e.*, five local tensors (see supplementary materials). Note that MPO decomposition can work with other compression methods: it can further reduce the parameters from the matrices compressed by other methods, and meanwhile largely reduce the parameters to be fine-tuned.

5 Experiments

In this section, we first set up the experiments, and then report the results and analysis.

5.1 Experimental Setup

Datasets. We evaluate the effectiveness of compressing and fine-tuning PLMs of our approach MPOP on the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2019). GLUE is a collection of 9 datasets for evaluating natural language understanding systems. Following (Sanh et al., 2019), we report macro-score (average of individual scores, which is slightly different from official GLUE score, since Spearman correlations are reported for STS-B and accuracy scores are reported for the other tasks) on the development sets for each task by fine-tuning MPOP.

Baselines. Our baseline methods include:

- **BERT** (Devlin et al., 2019): The 12-layer BERT-base model was pre-trained on Wikipedia corpus released by Google.
- **ALBERT** (Lan et al., 2020): It yields a highly compressed BERT variant with only 11.6M parameters, while maintains competitive performance, which serves as the major baseline.
- **DistilBERT** (Sanh et al., 2019): It is trained via knowledge distillation with 6 layers.
- **MobileBERT** (Sun et al., 2020c): It is equipped with bottleneck structures and a carefully designed balance between self-attentions and feed-forward networks.

All these models are released by Huggingface³. We select these baselines because they are widely adopted and have a diverse coverage of compression techniques. Note that we do not directly com-

³<https://huggingface.co/>

Experiments	Score	SST-2 (acc)	MNLI (m_cc)	QNLI (acc)	CoLA (mcc)	STS-B (ρ)	QQP (acc)	MRPC (acc)	RTE (acc)	WNLI (acc)	Avg. #Pr/#To(M)
ALBERT _{pub}	-	90.3	81.6	-	-	-	-	-	-	-	11.6/11.6
ALBERT _{rep}	78.9	90.6	84.5	89.4	53.4	88.2	89.1	88.5	71.1	54.9	11.6/11.6
MPOP	79.7	90.8	83.3	90.5	54.7	89.2	89.4	89.2	73.3	56.3	1.1/9
MPOP _{full}	80.3	92.2	84.4	91.4	55.7	89.2	89.6	87.3	76.9	56.3	12.7/12.7
MPOP _{full+LFA}	80.4	93.0	84.3	91.3	56.0	89.2	89.0	88.0	78.3	56.3	1.2/12.7
MPOP _{dir}	68.6	86.6	79.2	81.9	15.0	82.5	87.0	74.3	54.2	56.3	1.1/9

Table 3: Performance on GLUE benchmark obtained by fine-tuning ALBERT and MPOP. “ALBERT_{pub}” and “ALBERT_{rep}” denote the results from the original paper (Lan et al., 2020) and reproduced by ours, respectively. “#Pr” and “#To” denote the number (in millions) of pre-trained parameters and total parameters, respectively.

pare our approach with other competitive methods (Tambe et al., 2020) that require special optimization tricks or techniques (*e.g.*, hardware-level optimization).

Implementation. The original paper of ALBERT only reported the results of SST-2 and MNLI in GLUE. So we reproduce complete results denoted as “ALBERT_{rep}” with the Huggingface implementation (Wolf et al., 2020). Based on the pre-trained parameters provided by Huggingface, we also reproduce the results of BERT, DistilBERT and MobileBERT. To ensure a fair comparison, we adopt the same network architecture. For example, the number of self-attention heads, the hidden dimension of embedding vectors, and the max length of the input sentence are set to 12, 768 and 128, respectively.

5.2 Experimental Results

Note that our focus is to illustrate that our approach can improve either original (uncompressed) or compressed PLMs. In our main experiments, we adopt ALBERT as the major baseline, and report the comparison results in Table 3.

Comparison with ALBERT. As shown in Table 3, our approach MPOP is very competitive in the GLUE benchmark, and it outperforms ALBERT in all tasks (except MNLI) with a higher overall score of 79.7. Looking at the last column, compared with ALBERT, MPOP reduces total parameters by 22% (#To). In particular, it results in a significant reduction of pre-trained parameters by 91% (#Pr). Such a reduction is remarkable in lightweight fine-tuning, which dramatically improves the fine-tuning efficiency. By zooming in on specific tasks, the improvements over ALBERT are larger on CoLA, RTE and WNLI tasks. An interesting explanation is that RTE and WNLI tasks have small training sets (fewer than $4k$ samples). The lightweight fine-

tuning strategy seems to work better with limited training data, which enhances the capacity of PLMs and prevents overfitting on downstream tasks.

Ablation Results. Our approach has incorporated two novel improvements: lightweight fine-tuning with auxiliary tensors and optimization with dimension squeezing. We continue to study their effect on the final performance. Here we consider three variants for comparison: (1) MPOP_{full} and MPOP_{full+LFA} are full-rank MPO representation (without reconstruction error), and fine-tune *all the tensors* and *only auxiliary tensors*, respectively. This comparison is to examine whether only fine-tuning auxiliary tensors would lead to a performance decrease. (2) MPOP_{dir} directly optimizes the compressed model without the dimension squeezing algorithm. This variant is used to examine whether our optimization algorithm is more suitable for stacked architecture. Table 3 (last three rows) shows the results when we ablate these. In particular, the dimension squeezing algorithm plays a key role in improving our approach (a significant performance decrease for MPOP_{dir}), since it is tailored to stacked architecture. Comparing MPOP_{full} with MPOP_{full+LFA}, it is noted that fine-tuning all the parameters seems to have a negative effect on performance. Compared with ALBERT, we speculate that fine-tuning a large model is more likely to overfit on small datasets (*e.g.*, RTE and MRPC).

These results show that our approach is able to further compress ALBERT with fewer fine-tuning parameters. Especially, it is also helpful to improve the capacity and robustness of PLMs.

5.3 Detailed Analysis

In this section, we perform a series of detailed analysis experiments for our approach.

Evaluation with Other BERT Variants. In general, our approach can be applied to either uncom-

Models	WNLI (acc)	MRPC (acc)	RTE (acc)	Avg. #Pr/#To(M)
BERT	56.3	85.5	70.0	110/110
MPOP _B	56.3	84.3	70.8	7.7/70.4
DistilBERT	56.3	84.1	61.4	66/66
MPOP _D	56.3	84.3	61.7	4.0/43.4
MobileBERT	56.2	86.0	63.5	25.3/25.3
MPOP _M	56.2	85.3	65.7	4.4/15.4

Table 4: Evaluation with different BERT variants.

pressed or compressed PLMs. We have evaluated its performance with ALBERT. Now, we continue to test it with other BERT variants, namely original BERT, DistilBERT and MobileBERT. The latter two BERT variants are knowledge distillation based methods, and the distilled models can also be represented in the format of parameter matrix. We apply our approach to the three variants. Table 4 presents the comparison of the three variants before and after the application of MPOP. As we can see, our approach can substantially reduce the network parameters, especially the parameters to be fine-tuned. Note that DistilBERT and MobileBERT are highly compressed models. These results show that our approach can further improve other compressed PLMs.

Evaluation on Different Fine-Tuning Strategies.

Experiments have shown that our approach is able to largely reduce the number of parameters to be fine-tuned. Here we consider a more simple method to reduce the fine-tuning parameters, *i.e.*, only fine-tune the last layers of BERT. This experiment reuses the settings of BERT (12 layers) and our approach on BERT (*i.e.*, MPOP_B in Table 4). We fine-tune the last 1-3 layers of BERT, and compare the performance with our approach MPOP_B. From Table 5, we can see that such a simple way is much worse than our approach, especially on the RTE task. Our approach provides a more principled way for lightweight fine-tuning. By updating auxiliary tensors, it can better adapt to task-specific loss, and thus achieve better performance.

Evaluation on Low-Rank Approximation. As introduced in Section 4.4, MPO is a special low-rank approximation method, and we first compare its compression capacity with other low-rank approximation methods. As shown in Table 2, MPO and Tucker decomposition represent two main categories of low-rank approximation methods. We select CPD (Henry and Hofrichter, 1992)

Models	SST-2	MRPC	RTE	Avg. #Pr(M)
BERT ₁₀₋₁₂	91.9	76.5	67.2	45.7
BERT ₁₁₋₁₂	91.7	75.3	62.8	38.6
BERT ₁₂	91.4	72.1	61.4	31.5
MPOP _B	92.6	84.3	70.8	10.1

Table 5: Comparison of different fine-tuning strategies on three GLUE tasks. The subscript number in BERT_(·) denotes the index of the layers to be fine-tuned.




Figure 2: Comparison of different low-rank approximation variants. x -axis denotes the compression ratio (ρ in Eq. (5)) and y -axis denotes the reconstruction error, measured in the Frobenius norm.

for comparison because general Tucker decomposition (Tucker, 1966) cannot obtain results with reasonable memory. Our evaluation task is to compress the word embedding matrix of the released “bert-base-uncased” model⁴. As shown in Figure 2(a), MPO achieves a smaller reconstruction error with all compression ratios, which shows that MPO is superior to CPD. Another hyper-parameter in our MPO decomposition is the number of local tensors (n). We further perform the same evaluation with different numbers of local tensors ($n = 3, 5, 7$). From Figure 2(b), it can be observed that our method is relatively stable with respect to the number of local tensors. Overall, a larger n requires a higher time complexity and can yield flexible decomposition. Thus, we set $n = 5$ for making a trade-off between flexibility and efficiency.

6 Conclusion

We proposed an MPO-based PLM compression method. With MPO decomposition, we were able to reorganize and aggregate information in central tensors effectively. Inspired by this, we designed a novel fine-tuning strategy that only needs to fine-tune the parameters in auxiliary tensors. We also developed a dimension squeezing training algorithm for optimizing low-rank approximation over

⁴<https://huggingface.co/bert-base-uncased>

stacked network architectures. Extensive experiments had demonstrated the effectiveness of our approach, especially on the reduction of fine-tuning parameters. We also empirically found that such a fine-tuning way was more robust to generalize on small training datasets. To our knowledge, it is the first time that MPO decomposition had been applied to compress PLMs. In future work, we will consider exploring more decomposition structures for MPO.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China under Grants No. 61872369, 61832017 and 11934020, Beijing Academy of Artificial Intelligence (BAAI) under Grant No. BAAI2020ZJ0301, Beijing Outstanding Young Scientist Program under Grant No. BJJWZYJH012019100020098, the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China under Grant No. 18XNLG22, 19XNQ047, 20XNLG19 and 21XNH027. Xin Zhao and Zhong-Yi Lu are the corresponding authors.

References

- Pasquale Calabrese and John Cardy. 2004. Entanglement entropy and quantum field theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2004(06):P06002.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Prakhar Ganesh, Yao Chen, Xin Lou, Mohammad Ali Khan, Yin Yang, Deming Chen, Marianne Winslett, Hassan Sajjad, and Preslav Nakov. 2020. Compressing large-scale transformer-based models: A case study on bert. *arXiv preprint arXiv:2002.11985*.
- Ze-Feng Gao, Song Cheng, Rong-Qiang He, ZY Xie, Hui-Hai Zhao, Zhong-Yi Lu, and Tao Xiang. 2020. Compressing deep neural networks by matrix product operators. *Physical Review Research*, 2(2):023300.
- Timur Garipov, Dmitry Podoprikin, Alexander Novikov, and Dmitry Vetrov. 2016. Ultimate tensorization: compressing convolutional and fc layers alike. *arXiv preprint arXiv:1611.03214*.
- ER Henry and J Hofrichter. 1992. [8] singular value decomposition: Application to analysis of experimental data. *Methods in enzymology*, 210:129–192.
- Frank L Hitchcock. 1927. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189.
- Lu Hou, Zhiqi Huang, Lifeng Shang, Xin Jiang, Xiao Chen, and Qun Liu. 2020. [Dynabert: Dynamic BERT with adaptive width and depth](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2017. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898.
- Forrest N Iandola, Albert E Shaw, Ravi Krishna, and Kurt W Keutzer. 2020. Squeezebert: What can computer vision teach nlp about efficient neural networks? *arXiv preprint arXiv:2006.11316*.
- S Iblisdir, R Orus, and JI Latorre. 2007. Matrix product states algorithms and continuous systems. *Physical Review B*, 75(10):104305.
- 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.
- 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, Online. Association for Computational Linguistics.
- Valentin Khruikov, Oleksii Hrinchuk, Leyla Mirvakhabova, and Ivan Oseledets. 2019. Tensorized embedding layers for efficient model compression. *arXiv preprint arXiv:1901.10787*.
- 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 *8th International Conference on Learning Representations*,

- ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.
- Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2020. [Exploring versatile generative language model via parameter-efficient transfer learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 441–459. Association for Computational Linguistics.
- Weijie Liu, Peng Zhou, Zhiruo Wang, Zhe Zhao, Haotang Deng, and Qi Ju. 2020. [Fastbert: a self-distilling BERT with adaptive inference time](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 6035–6044. Association for Computational Linguistics.
- Xindian Ma, Peng Zhang, Shuai Zhang, Nan Duan, Yuexian Hou, Ming Zhou, and Dawei Song. 2019. [A tensorized transformer for language modeling](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 2229–2239.
- Matan Ben Noach and Yoav Goldberg. 2020. Compressing pre-trained language models by matrix decomposition. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 884–889.
- Alexander Novikov, Dmitry Podoprikin, Anton Osookin, and Dmitry P. Vetrov. 2015. [Tensorizing neural networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 442–450.
- Ivan V Oseledets. 2011. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317.
- Nikolaos Pappas, Phoebe Mulcaire, and Noah A Smith. 2020. Grounded compositional outputs for adaptive language modeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1252–1267.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Bogdan Pirvu, Valentin Murg, J Ignacio Cirac, and Frank Verstraete. 2010. Matrix product operator representations. *New Journal of Physics*, 12(2):025012.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Evani Radiya-Dixit and Xin Wang. 2020. How fine can fine-tuning be? learning efficient language models. In *International Conference on Artificial Intelligence and Statistics*, pages 2435–2443. PMLR.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. [Learning multiple visual domains with residual adapters](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 506–516.
- 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*.
- 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 *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Siqi Sun, Zhe Gan, Yuwei Fang, Yu Cheng, Shuohang Wang, and Jingjing Liu. 2020a. Contrastive distillation on intermediate representations for language model compression. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 498–508.
- Xingwei Sun, Ze-Feng Gao, Zhong-Yi Lu, Junfeng Li, and Yonghong Yan. 2020b. A model compression method with matrix product operators for speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2837–2847.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020c. [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.
- Thierry Tambe, Coleman Hooper, Lillian Pentecost, En-Yu Yang, Marco Donato, Victor Sanh, Alexander M Rush, David Brooks, and Gu-Yeon Wei. 2020. Edgebert: Optimizing on-chip inference for multi-task nlp. *arXiv preprint arXiv:2011.14203*.
- Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. 2017. Compressing recurrent neural network with tensor train. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 4451–4458. IEEE.

- Ledyard R Tucker. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45.
- Ji Xin, Raphael Tang, Jaejun Lee, Yaoliang Yu, and Jimmy Lin. 2020. Deebert: Dynamic early exiting for accelerating bert inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2246–2251.
- Rose Yu, Stephan Zheng, Anima Anandkumar, and Yisong Yue. 2017. Long-term forecasting using tensor-train rnns. *Arxiv*.
- Ofir Zafir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. Q8bert: Quantized 8bit bert. *arXiv preprint arXiv:1910.06188*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Mengjie Zhao, Tao Lin, Fei Mi, Martin Jaggi, and Hinrich Schütze. 2020. [Masking as an efficient alternative to finetuning for pretrained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2226–2241. Association for Computational Linguistics.