

BERT-EMD: Many-to-Many Layer Mapping for BERT Compression with Earth Mover’s Distance

Jianquan Li^{1*}, Xiaokang Liu^{1*}, Honghong Zhao¹, Ruifeng Xu², Min Yang^{3†}, Yaohong Jin¹

¹Beijing Ultrapower Software Co.,Ltd., China

²Harbin Institute of Technology (Shenzhen), China

³Shenzhen Key Laboratory for High Performance Data Mining,

Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China

{lijianquan2, liuxiaokang1, zhaohonghong1, jinyaohong}@ultrapower.com.cn
xurui Feng@hit.edu.cn, min.yang@siat.ac.cn

Abstract

Pre-trained language models (e.g., BERT) have achieved significant success in various natural language processing (NLP) tasks. However, high storage and computational costs obstruct pre-trained language models to be effectively deployed on resource-constrained devices. In this paper, we propose a novel BERT distillation method based on many-to-many layer mapping, which allows each intermediate student layer to learn from any intermediate teacher layers. In this way, our model can learn from different teacher layers adaptively for various NLP tasks. In addition, we leverage Earth Mover’s Distance (EMD) to compute the minimum cumulative cost that must be paid to transform knowledge from teacher network to student network. EMD enables the effective matching for many-to-many layer mapping. Furthermore, we propose a cost attention mechanism to learn the layer weights used in EMD automatically, which is supposed to further improve the model’s performance and accelerate convergence time. Extensive experiments on GLUE benchmark demonstrate that our model achieves competitive performance compared to strong competitors in terms of both accuracy and model compression. For reproducibility, we release the code and data at <https://github.com/lxk00/BERT-EMD>.

1 Introduction

In recent years, pre-trained language models, such as GPT (Radford et al., 2018), BERT (Devlin et al., 2018), XL-Net (Yang et al., 2019), have been proposed and applied to many NLP tasks, yielding state-of-the-art performances. However, the promising results of the pre-trained language models come with the high costs of computation

and memory in inference, which obstruct these pre-trained language models to be deployed on resource-constrained devices and real-time applications. For example, the original BERT-base model, which achieved great success in many NLP tasks, has 12 layers and about 110 millions parameters.

It is therefore critical to effectively accelerate inference time and reduce the computational workload while maintaining accuracy. This research issue has attracted increasing attention (Wang et al., 2019; Shen et al., 2019; Tang et al., 2019), of which knowledge distillation (Tang et al., 2019) is considered to be able to provide a practical way. Typically, knowledge distillation techniques train a compact and shallow student network under the guidance of a complicated larger teacher network with a teacher-student strategy (Watanabe et al., 2017). Once trained, this compact student network can be directly deployed in real-life applications.

So far, there have been several studies, such as DistilBERT (Tang et al., 2019), BERT-PKD (Sun et al., 2019), TinyBERT (Jiao et al., 2019), which attempt to compress the original BERT into a lightweight student model without performance sacrifice based on knowledge distillation. For example, BERT-PKD (Sun et al., 2019) and TinyBERT (Jiao et al., 2019) are two representative BERT compression approaches, which encourage the student model to extract knowledge from both the last layer and the intermediate layers of the teacher network.

Despite the effectiveness of previous studies, there are still several challenges for distilling comprehensive knowledge from the teacher model, which are not addressed well in prior works. **First**, existing compression methods learn one-to-one layer mapping, where each student layer is guided by only one specific teacher layer. For example, BERT-PKD uses the 2, 4, 6, 8, 10 teacher layers to guide the 1 to 5 student layers, respectively. How-

* Equal contribution

† Min Yang is corresponding author

ever, these one-to-one layer mapping strategies are assigned based on empirical observations without theoretical guidance. **Second**, as revealed in (Clark et al., 2019), different BERT layers could learn different levels of linguistic knowledge. The one-to-one layer mapping strategy cannot learn an optimal, unified compressed model for different NLP tasks. In addition, most previous works do not consider the importance of each teacher layer and use the same layer weights among various tasks, which create a substantial barrier for generalizing the compressed model to different NLP tasks. Therefore, an adaptive compression model should be designed to transfer knowledge from all teacher layers dynamically and effectively for different NLP tasks.

To address the aforementioned issues, we propose a novel BERT compression approach based on many-to-many layer mapping and Earth Mover’s Distance (EMD) (Rubner et al., 2000), called BERT-EMD. **First**, we design a many-to-many layer mapping strategy, where each intermediate student layer has the chance to learn from all the intermediate teacher layers. In this way, BERT-EMD can learn from different intermediate teacher layers adaptively for different NLP tasks, motivated by the intuition that different NLP tasks require different levels of linguistic knowledge contained in the intermediate layers of BERT. **Second**, to learn an optimal many-to-many layer mapping strategy, we leverage EMD to compute the minimum cumulative cost that must be paid to transform knowledge from teacher network to student network. EMD is a well-studied optimization problem and provides a suitable solution to transfer knowledge from the teacher network in a holistic fashion.

We summarize our main contributions as follows. (1) We propose a novel many-to-many layer mapping strategy for compressing the intermediate layers of BERT in an adaptive and holistic fashion. (2) We leverage EMD to formulate the distance between the teacher and student networks, and learn an optimal many-to-many layer mapping based on a solution to the well-known transportation problem. (3) We propose a cost attention mechanism to learn the layer weights used in EMD automatically, which can further improve the model’s performance and accelerate convergence time. (4) Extensive experiments on GLUE tasks show that BERT-EMD achieves better performance than the state-of-the-art BERT distillation methods.

2 Related Work

Language models pre-trained on large-scale corpora can learn universal language representations, which have proven to be effective in many NLP tasks (Mikolov et al., 2013; Pennington et al., 2014; Joulin et al., 2016). Early efforts mainly focus on learning good word embeddings, such as word2vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014). Although these pre-trained embeddings can capture semantic meanings of words, they are context-free and fail to capture higher-level concepts in context, such as syntactic structures and polysemous disambiguation. Subsequently, researchers have shifted attention to contextual word embeddings learning, such as ELMo (Peters et al., 2018), ULMFit (Howard and Ruder, 2018), GPT (Radford et al., 2018), BERT (Devlin et al., 2018), ENRIE (Zhang et al., 2019), XL-Net (Yang et al., 2019), RoBERTa (Liu et al., 2019). For example, Devlin et al. (2018) released the BERT-base of 110 million parameters and BERT-large of 330 million parameters, which achieved significantly better results than previous methods on GLUE tasks.

However, along with high-performance, the pre-trained language models (e.g., BERT) usually have a large number of parameters, which require a high cost of computation and memory in inference. Recently, many attempts have been made to reduce the computation overhead and model storage of pre-trained language models without performance sacrifice. Existing compression techniques can be divided into three categories: low-rank matrix factorization (Wang et al., 2019), quantization (Shen et al., 2019), and knowledge distillation (Tang et al., 2019). Next, we mainly review the related works that use knowledge distillation to compress the BERT model.

Knowledge distillation using the teacher-student strategy learns a lightweight student network under the guidance of a large and complicated teacher network. Mukherjee and Awadallah (2019) distilled BERT into an LSTM network via both hard and soft distilling methods. Sun et al. (2019) proposed the BERT-PKD model to transfer the knowledge from both the final layer and the intermediate layers of the teacher network. Jiao et al. (2019) proposed the TinyBERT model, which performed the Transformer distillation at both pre-training and fine-tuning processes. Xu et al. (2020) proposed the BERT-of-Theseus model to learn a compact student network by replacing the teacher layers with

their substitutes. Sun et al. (2020) introduced the MobileBERT model, which has the same number of layers with the teacher network, but was much narrower via adopting bottleneck structures. Wang et al. (2020) distilled the self-attention module of the last Transformer layer of the teacher network.

However, the aforementioned BERT compression approaches struggle to find an optimal layer mapping between the teacher and student networks. Each student layer merely learns from a single teacher layer, which may lose rich linguistic knowledge contained in the teacher network. Different from previous methods, we propose a many-to-many layer mapping method for BERT distillation, where each intermediate student layer can learn from any intermediate teacher layers adaptively. In addition, an Earth Mover’s Distance is applied to learn the optimal many-to-many layer mapping solution.

3 Methodology

In this section, we propose a novel BERT compression method based on many-to-many layer mapping and Earth Mover’s Distance (called BERT-EMD). In addition, we also propose a cost attention mechanism to learn the layer weights used in EMD automatically.

3.1 Overview of BERT-EMD

The main idea behind BERT-EMD is to transfer knowledge from a large teacher network T (large BERT) to a small student network S (BERT-EMD). Both the student and teacher networks are implemented with an embedding layer, several Transformer layers, and a prediction layer. We assume that the teacher network has M Transformer layers and the student network has N Transformer layers. Each Transformer layer contains an attention layer and a hidden layer.

Similar to TinyBERT (Jiao et al., 2019), our method also includes three primary distillation components: the embedding-layer distillation, the Transformer distillation, and the prediction-layer distillation. Concretely, both the embedding-layer distillation and the prediction-layer distillation employ the one-to-one layer mapping as in TinyBERT and BERT-PKD, where the two student layers are guided by the corresponding teacher layers, respectively. However, different from the previous works, we propose to exploit the many-to-many layer mapping for Transformer (intermediate lay-

ers) distillation (attention-based distillation and hidden states based distillation), where each student attention layer (resp. hidden layer) can learn from any teacher attention layers (resp. hidden layers). In this way, BERT-EMD can learn from different intermediate teacher layers adaptively for different NLP tasks, motivated by the intuition that different NLP tasks require different levels of linguistic knowledge contained in the attention and hidden layers of BERT. Next, we will describe the four distillation strategies of BERT-EMD in detail.

3.2 Embedding-layer Distillation

Word embeddings are vital in NLP tasks and have been extensively studied in recent years. Better representations of words have come at the cost of huge memory footprints. Compressing embedding matrices without sacrificing model performance is essential for real-world applications. To this end, we minimize the mean squared error (MSE) between the embedding layers of the teacher and student networks:

$$\mathcal{L}_{\text{emb}} = \text{MSE}(\mathbf{E}^S \mathbf{W}_e, \mathbf{E}^T) \quad (1)$$

where the matrices \mathbf{E}^S and \mathbf{E}^T represent the embeddings of student and teacher networks, which have the same shape. \mathbf{W}_e is a projection parameter to be learned.

3.3 Prediction-layer Distillation

The student network also learns from the probability logits provided by teacher network. We minimize the prediction-layer distillation function as:

$$\mathcal{L}_{\text{pred}} = -\text{softmax}(\mathbf{z}^T) \cdot \log_{\text{softmax}}(\mathbf{z}^S/t) \quad (2)$$

where \mathbf{z}^T and \mathbf{z}^S represent the probability logits predicted by the teacher and student, respectively. t indicates a temperature value.

3.4 Transformer Distillation with Earth Mover’s Distance

Instead of imposing one-to-one layer mapping as in previous works (Sun et al., 2019; Jiao et al., 2019), our Transformer distillation approach allows many-to-many layer mapping and is capable of generalizing to various NLP tasks. The Earth Mover’s Distance (EMD) is proposed to measure the dissimilarity (distance) between the teacher and student networks as the minimum cumulative cost

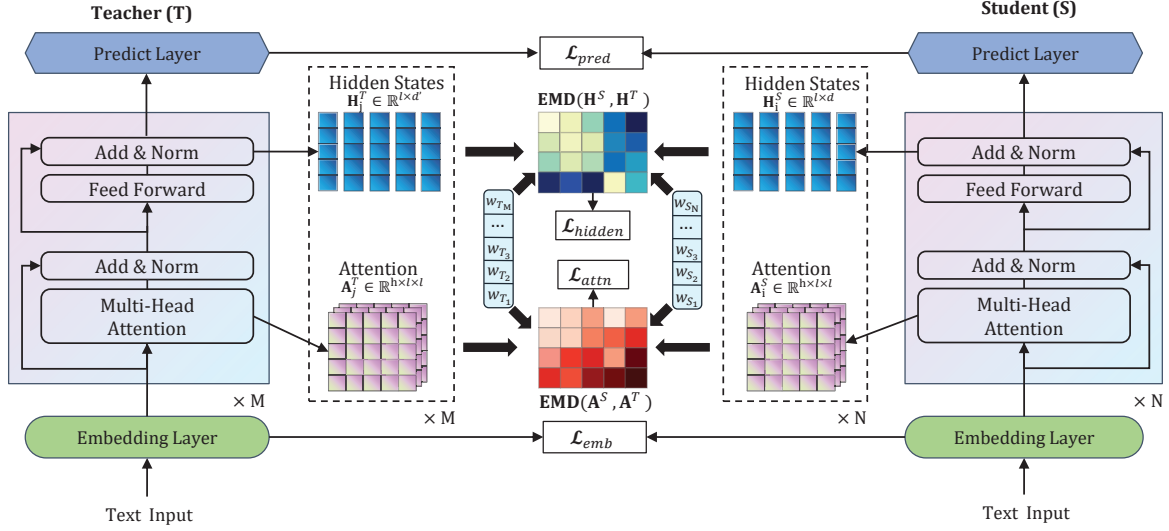


Figure 1: An overview of the proposed BERT-EMD method, which distills comprehensive knowledge from a large teacher (T) with M -layer Transformer to a small student (S) with N -layer Transformer. w_{T_i} and w_{S_j} are the weight of i -th teacher layer and j -th student layer used in EMD. Here, l denotes the length of the input sequence. h denotes the head number. d and d' are the hidden sizes of student and teacher Transformers, respectively.

of transforming knowledge from the teacher network to student network. The key insight is to view network layers as distributions, and the desired transformation should make the two distributions (teacher and student layers) close.

Attention-based Distillation We use the attention-based distillation to transform the linguistic knowledge from the teacher network to the student network based on EMD. Formally, let $\mathbf{A}^T = \{(\mathbf{A}_1^T, w_{T_1}^A), \dots, (\mathbf{A}_M^T, w_{T_M}^A)\}$ be the teacher attention layers and $\mathbf{A}^S = \{(\mathbf{A}_1^S, w_{S_1}^A), \dots, (\mathbf{A}_N^S, w_{S_N}^A)\}$ be the student attention layers, where M and N represent the numbers of the attention layers in the teacher and student networks, respectively. Each \mathbf{A}_i^T (resp. \mathbf{A}_i^S) represents the i -th teacher (resp. student) attention layer and $w_{T_i}^A$ (resp. $w_{S_i}^A$) indicates corresponding layer weight that is initialized as $\frac{1}{M}$ (resp. $\frac{1}{N}$). We also define a ‘‘ground’’ distance matrix $\mathbf{D}^A = [d_{ij}^A]$, where d_{ij}^A represents the cost of transferring the attention knowledge from \mathbf{A}_i^T to \mathbf{A}_j^S . Here, we use MSE to calculate the distance d_{ij}^A as:

$$d_{ij}^A = \text{MSE}(\mathbf{A}_i^S, \mathbf{A}_j^T) \quad (3)$$

Then, we attempt to find a mapping flow $\mathbf{F}^A = [f_{ij}^A]$, with f_{ij}^A the mapping flow between \mathbf{A}_i^T and \mathbf{A}_j^S , that minimizes the cumulative cost required to transform knowledge from the teacher attention

layers \mathbf{A}^T to the student attention layers \mathbf{A}^S :

$$\text{WORK}(\mathbf{A}^T, \mathbf{A}^S, \mathbf{F}^A) = \sum_{i=1}^M \sum_{j=1}^N f_{ij}^A d_{ij}^A \quad (4)$$

subject to the following constraints:

$$f_{ij}^A \geq 0 \quad 1 \leq i \leq M, 1 \leq j \leq N \quad (5)$$

$$\sum_{j=1}^N f_{ij}^A \leq w_{T_i}^A \quad 1 \leq i \leq M \quad (6)$$

$$\sum_{i=1}^M f_{ij}^A \leq w_{S_j}^A \quad 1 \leq j \leq N \quad (7)$$

$$\sum_{i=1}^M \sum_{j=1}^N f_{ij}^A = \min\left(\sum_{i=1}^M w_{T_i}^A, \sum_{j=1}^N w_{S_j}^A\right) \quad (8)$$

where the first constraint forces the mapping flow to be positive. The second constraint limits the amount of attention information that can be sent by \mathbf{A}^T to their weights. The third constraint limits the attention information that can be received by \mathbf{A}^S . The fourth constraint limits the amount of total flow.

The above optimization is a well-studied transportation problem (Hitchcock, 1941), which can be solved by previously developed methods (Rachev, 1985). Once the optimal mapping flow \mathbf{F}^A is learned, we can define the Earth Mover’s Distance as the work normalized by the total flow:

$$\text{EMD}(\mathbf{A}^S, \mathbf{A}^T) = \frac{\sum_{i=1}^M \sum_{j=1}^N f_{ij}^A d_{ij}^A}{\sum_{i=1}^M \sum_{j=1}^N f_{ij}^A} \quad (9)$$

Finally, the objective function for the attention-based distillation can be defined by the EMD between \mathbf{A}^T and \mathbf{A}^S :

$$\mathcal{L}_{\text{attn}} = \text{EMD}(\mathbf{A}^S, \mathbf{A}^T) \quad (10)$$

Hidden States-based Distillation Similar to attention-based distillation, we also learn the hidden layer mapping based on EMD. Formally, let $\mathbf{H}^T = \{(\mathbf{H}_1^T, w_{T_1}^{\mathbf{H}}), \dots, (\mathbf{H}_M^T, w_{T_M}^{\mathbf{H}})\}$ be the teacher hidden layers and $\mathbf{H}^S = \{(\mathbf{H}_1^S, w_{S_1}^{\mathbf{H}}), \dots, (\mathbf{H}_N^S, w_{S_N}^{\mathbf{H}})\}$ be the student hidden layers, where M and N represent the numbers of the hidden layers in the teacher and student networks, respectively. Each \mathbf{H}_i^T represents the i -th hidden layer and $w_{T_i}^{\mathbf{H}}$ indicates corresponding layer weight that is initialized as $\frac{1}{M}$. We also define a ‘‘ground’’ distance matrix $\mathbf{D}^{\mathbf{H}} = [d_{ij}^{\mathbf{H}}]$, where $d_{ij}^{\mathbf{H}}$ represents the cost of transferring the hidden states knowledge from \mathbf{H}_i^T to \mathbf{H}_j^S and we use a learnable projection parameter as \mathbf{W}_h . MSE is applied to calculate the distance $d_{ij}^{\mathbf{H}}$:

$$d_{ij}^{\mathbf{H}} = \text{MSE}(\mathbf{H}_i^S \mathbf{W}_h, \mathbf{H}_j^T) \quad (11)$$

Then, a mapping flow $\mathbf{F}^{\mathbf{H}} = [f_{ij}^{\mathbf{H}}]$, with $f_{ij}^{\mathbf{H}}$ the mapping flow between \mathbf{H}_i^T and \mathbf{H}_j^S , is learned by minimizing the cumulative cost required to transform knowledge from \mathbf{H}^T to \mathbf{H}^S :

$$\text{WORK}(\mathbf{H}^T, \mathbf{H}^S, \mathbf{F}^{\mathbf{H}}) = \sum_{i=1}^M \sum_{j=1}^N f_{ij}^{\mathbf{H}} d_{ij}^{\mathbf{H}} \quad (12)$$

subject to the following constraints:

$$f_{ij}^{\mathbf{H}} \geq 0 \quad 1 \leq i \leq M, 1 \leq j \leq N \quad (13)$$

$$\sum_{j=1}^N f_{ij}^{\mathbf{H}} \leq w_{T_i}^{\mathbf{H}} \quad 1 \leq i \leq M \quad (14)$$

$$\sum_{i=1}^M f_{ij}^{\mathbf{H}} \leq w_{S_j}^{\mathbf{H}} \quad 1 \leq j \leq N \quad (15)$$

$$\sum_{i=1}^M \sum_{j=1}^N f_{ij}^{\mathbf{H}} = \min\left(\sum_i w_{T_i}^{\mathbf{H}}, \sum_i w_{S_i}^{\mathbf{H}}\right) \quad (16)$$

After solving the above optimization problem, we obtain the optimal mapping flow $\mathbf{F}^{\mathbf{H}}$. The earth mover’s distance can be then defined as the work normalized by the total flow:

$$\text{EMD}(\mathbf{H}^S, \mathbf{H}^T) = \frac{\sum_{i=1}^M \sum_{j=1}^N f_{ij}^{\mathbf{H}} d_{ij}^{\mathbf{H}}}{\sum_{i=1}^M \sum_{j=1}^N f_{ij}^{\mathbf{H}}} \quad (17)$$

Finally, the objective function for the hidden states-based distillation can be defined by the earth mover’s distance between \mathbf{H}^T and \mathbf{H}^S :

$$\mathcal{L}_{\text{hidden}} = \text{EMD}(\mathbf{H}^S, \mathbf{H}^T) \quad (18)$$

3.5 Weight Update with Cost Attention

In the EMD defined in Section 3.4, each teacher layer (resp. student layer) is assigned an equal weight $w_T = \frac{1}{M}$ (resp. $w_S = \frac{1}{N}$). Since different attention and hidden layers of BERT can learn different levels of linguistic knowledge, these layers should have different weights for various NLP tasks. Therefore, we propose a cost attention mechanism to assign weights for each attention and hidden layers automatically.

The main idea behind the cost attention is to make the teacher and student Transformer networks be as close as possible. That is, we could reduce the overall cost of EMD by increasing the weights of the layers with low flow cost, while the weights of the layers with high flow cost should be decreased adaptively.

We take the weight updating process of the teacher network as an example. The cost attention mechanism can be performed by three steps after learning the optimal solution (flow matrices $\mathbf{F}^{\mathbf{A}}$ and $\mathbf{F}^{\mathbf{H}}$ in EMD). First, we learn the transferring cost between each teacher and student layers (unit transferring cost). Formally, let $C_{T_i}^{\mathbf{A}}$ and $C_{T_i}^{\mathbf{H}}$ be the unit transferring cost of each attention and hidden layers respectively, which can be computed as:

$$\bar{C}_{T_i}^{\mathbf{A}} = \frac{\sum_{j=1}^N d_{ij}^{\mathbf{A}} f_{ij}^{\mathbf{A}}}{w_{T_i}} \quad (19)$$

$$\bar{C}_{T_i}^{\mathbf{H}} = \frac{\sum_{j=1}^N d_{ij}^{\mathbf{H}} f_{ij}^{\mathbf{H}}}{w_{T_i}} \quad (20)$$

Second, we update the weights ($w_{T_i}^{\mathbf{A}}$ and $w_{T_i}^{\mathbf{H}}$) of the teacher attention and hidden layers based on the learned unit transferring cost. Specifically, we compute the updated weights $w_{T_i}^{\mathbf{A}}$ and $w_{T_i}^{\mathbf{H}}$ as the inverse ratio of the transferring costs:

$$w_{T_i}^{\mathbf{A}} = \frac{\sum_{j=1}^M \bar{C}_j^{\mathbf{A}}}{\bar{C}_{T_i}^{\mathbf{A}}} \quad (21)$$

$$w_{T_i}^{\mathbf{H}} = \frac{\sum_{j=1}^M \bar{C}_{T_j}^{\mathbf{H}}}{\bar{C}_{T_i}^{\mathbf{H}}} \quad (22)$$

Finally, we normalize the updated layer weights used in EMD via softmax, and introduce a temperature coefficient τ to smooth the results. In particular, we update weight w_{T_i} of the i -th Transformer layer used in EMD by averaging the corresponding weights of attention and hidden layers:

$$w_{T_i} = \frac{1}{2}(\text{softmax}(w_{T_i}^{\bar{A}}/\tau) + \text{softmax}(w_{T_i}^{\bar{H}}/\tau)) \quad (23)$$

It is noteworthy that the learned new weights are leveraged as the constraints to optimize the EMD problem in the next batch. Specifically, we initialize the i -th teacher attention and hidden layer weights ($w_{T_i}^A$ and $w_{T_i}^H$) in the η -th batch with the updated weight w_{T_i} learned in the $\eta - 1$ -th batch. In this way, we can further improve the performance of BERT-EMD and accelerate convergence time.

3.6 Overall Learning Objective

Finally, we combine the embedding-layer distillation, attention-based distillation, hidden states-based distillation, prediction-layer distillation objectives to form the overall knowledge distillation objective as follows:

$$\mathcal{L}_{\text{distill}} = \beta(\mathcal{L}_{\text{emb}} + \mathcal{L}_{\text{attn}} + \mathcal{L}_{\text{hidden}}) + \mathcal{L}_{\text{pred}} \quad (24)$$

where β is a factor that controls the weights of the three distillation objectives (\mathcal{L}_{emb} , $\mathcal{L}_{\text{attn}}$, $\mathcal{L}_{\text{hidden}}$).

4 Experimental Setup

4.1 Experimental Data

We evaluate our BERT-EMD model on the General Language Understanding Evaluation (GLUE) (Wang et al., 2018) benchmark, which is a collection of nine diverse sentence-level classification tasks. Concretely, GLUE consists of (i) Microsoft Research Paraphrase Matching (MRPC), Quora Question Pairs (QQP) and Semantic Textual Similarity Benchmark (STS-B) for paraphrase similarity matching; (ii) Stanford Sentiment Treebank (SST-2) for sentiment classification; (iii) Multi-Genre Natural Language Inference Matched (MNLI-m), Multi-Genre Natural Language Inference Mismatched (MNLI-mm), Question Natural Language Inference (QNLI) and Recognizing Textual Entailment (RTE) for natural language inference task; and (iv) the Corpus of Linguistic Acceptability (CoLA) for linguistic acceptability.

4.2 Evaluation Metrics

Following previous works (Sun et al., 2019; Jiao et al., 2019), we use classification accuracy as the evaluation metric for SST-2, MNLI-m, MNLI-mm, QNLI, and RTE datasets. For a fair comparison with TinyBERT (Jiao et al., 2019), the F1 metric is adopted for MRPC and QQP datasets, the Spearman correlation is adopted for STS-B, and the Matthew’s correlation is adopted for CoLA. The results reported for the test set of GLUE are in the same format as on the official leaderboard.

4.3 Implementation Details

Similar to TinyBERT, our BERT-EMD method also contains a general distillation and a task-specific distillation. In particular, we initialize our student model with the general distillation model provided by TinyBERT¹. The teacher model is implemented as a 12-layer BERT model (BERT_{BASE12}), which is fine-tuned for each task to perform knowledge distillation.

We employ the grid search algorithm on the validation set to tune the hyper-parameters. Since there are many hyper-parameter combinations, we first do the grid search on β and the learning rate. Then, we fix the values of these two hyper-parameters and tune the values of the other hyper-parameters. Specifically, the batch size is 32, the learning rate is tuned from $\{5e - 5, 2e - 5, 1e - 5\}$, the parameter t defined in Eq. (2) is tuned from $\{1, 3, 7, 10\}$, the temperature coefficient τ is tuned from $\{1, 2, 5, 10\}$, and β is tuned from $\{0.01, 0.001, 0.005\}$.

4.4 Baseline Methods

In this paper, we compare our BERT-EMD with several state-of-the-art BERT compression approaches, including the original 4/6-layer BERT models (Devlin et al., 2018), DistilBERT (Tang et al., 2019), BERT-PKD (Sun et al., 2019), TinyBERT (Jiao et al., 2019), BERT-of-Theseus (Xu et al., 2020). However, the original TinyBERT employs a data augmentation strategy in the training process, which is different from the other baseline models. For a fair comparison, we re-implement the TinyBERT model by eliminating the data augmentation strategy.

It is noteworthy that we do not compare BERT-EMD with the recent MobileBERT (Sun et al., 2020) and MiniLM (Wang et al., 2020), since

¹<https://github.com/TinyBERT/TinyBERT>

Model	Params Inference		MNLI-m	MNLI-mm	QQP	SST-2	CoLA	QNLI	MRPC	RTE	STS-b	AVE
	Num	Time	(393k)	(393k)	(364k)	(67k)	(8.5k)	(108k)	(3.5k)	(2.5k)	(5.7k)	
BERT _{BASE12} -G	110M	×1	84.6	83.4	71.2	93.5	52.1	90.5	88.9	66.4	85.8	79.60
BERT _{BASE12} -T	110M	×1	84.4	83.3	71.6	93.4	52.8	90.5	88.1	66.9	85.2	79.58
BERT _{SMALL4}	14.5M	-	75.4	74.9	66.5	87.6	19.5	84.8	83.2	62.6	77.1	70.18
DistilBERT ₄	52.2M	×3.0	78.9	78.0	68.5	91.4	32.8	85.2	82.4	54.1	76.1	71.93
BERT-PKD ₄	52.2M	×3.0	79.9	79.3	70.2	89.4	24.8	85.1	82.6	62.3	79.8	72.60
TinyBERT ₄	14.5M	×9.4	81.2	80.3	68.9	90.0	25.3	86.2	85.4	63.9	80.4	73.51
BERT-EMD₄	14.5M	×9.4	82.1	80.6	69.3	91.0	25.6	87.2	87.6	66.2	82.3	74.66
BERT-PKD ₆	66.0M	×1.9	81.5	81.0	70.7	92.0	43.5	89.0	85.0	65.5	81.6	76.61
BERT-of-Theseus ₆	66.0M	-	82.4	82.1	71.6	92.2	47.8	89.6	87.6	66.2	84.1	78.18
TinyBERT ₆	66.0M	×1.9	84.4	83.1	71.3	92.6	46.1	89.8	88.0	69.7	83.9	78.77
BERT-EMD₆	66.0M	×1.9	84.7	83.5	72.0	93.3	47.5	90.7	89.8	71.7	86.8	80.00

Table 1: Experimental results on the GLUE test set. The subscript within each model name represents the number of Transformer layers. AVE represents the average score over all tasks. BERT_{BASE12}-G and BERT_{BASE12}-T indicate the results of the fine-tuned BERT-base from (Devlin et al., 2018) and in our implementation, respectively.

Method	MNLI-m	QQP	RTE	STS-b
BERT-EMD ₄	82.1	69.3	66.2	82.3
w/o CA ₄	81.6	69.0	65.1	81.6
w/o EMD ₄	80.7	67.7	64.1	80.7
BERT-EMD ₆	84.7	72.0	71.7	86.8
w/o CA ₆	84.5	71.6	71.0	85.3
w/o EMD ₆	84.2	71.2	70.4	84.7

Table 2: Ablation test results in terms of removing EMD (w/o EMD) and cost attention (w/o CA).

MiniLM does not report the results on the GLUE test set and the MobileBERT model employs the Transformer block with different architectures.

5 Experimental Results

5.1 Main Results

We summarize the experimental results on the GLUE test sets in Table 1. The number below each task denotes the number of training instances. Following previous works (Sun et al., 2019,?), we also report the average values of these nine tasks (the ‘‘AVE’’ column). From the results, we can observe that BERT-EMD substantially outperforms state-of-the-art baseline methods by a noticeable margin on most tasks.

Among all the 4-layer BERT approaches, our BERT-EMD₄ method achieves the best results on almost all the tasks except SST-2 and CoLA. First, BERT-EMD₄ achieves significantly better results than BERT_{SMALL4} on all the GLUE tasks with a large improvement of 4.48% on average. Second, BERT-EMD₄ also outperforms DistilBERT₄ and BERT-PKD₄ by a substantial margin, even with only 30% parameters and inference time. Furthermore, BERT-EMD₄ exceeds the TinyBERT model (the best competitor) by 2.3% accuracy on RTE, 2.2% F1 on MRPC, and 1.9% Spearman correlation on STS-B. This verifies the

effectiveness of our BERT-EMD model in improving the performance of small BERT-based methods on various language understanding tasks.

We can observe similar trends in the 6-layer BERT models. Table 1 shows that the proposed BERT-EMD₆ method can effectively compress BERT_{BASE12} into a 6-layer BERT model without performance sacrifice. Specifically, BERT-EMD₆ performs better than the 12-layer BERT BERT_{BASE12} model on 7 out of 9 tasks, with only about 50% parameters and inference time of the original BERT_{BASE12} model. For example, BERT-EMD achieves a noticeable improvement of 5.3% accuracy on RTE and 1% Spearman correlation on STS-B, over the BERT_{BASE12} model.

5.2 Ablation Study

To verify the effectiveness of EMD and the cost attention mechanism, we perform ablation test of BERT-EMD on two large datasets (MNLI and QQP) and two small datasets (MRPC and RTE) in terms of removing EMD (denoted as w/o EMD) and cost attention (w/o CA), respectively. In particular, for the method of removing EMD, we retain the many-to-many layer mapping by simply replacing the EMD with the mean squared error when measuring the distance between the teacher and student layers.

The ablation test results are summarized in Table 2. Generally, both EMD and cost attention contribute noticeable improvement to our method. The performances decrease sharply, especially on the STS-B task, when removing the EMD module. This is within our expectation since the EMD module formulates the distance between the teacher and student networks as an optimal transport prob-

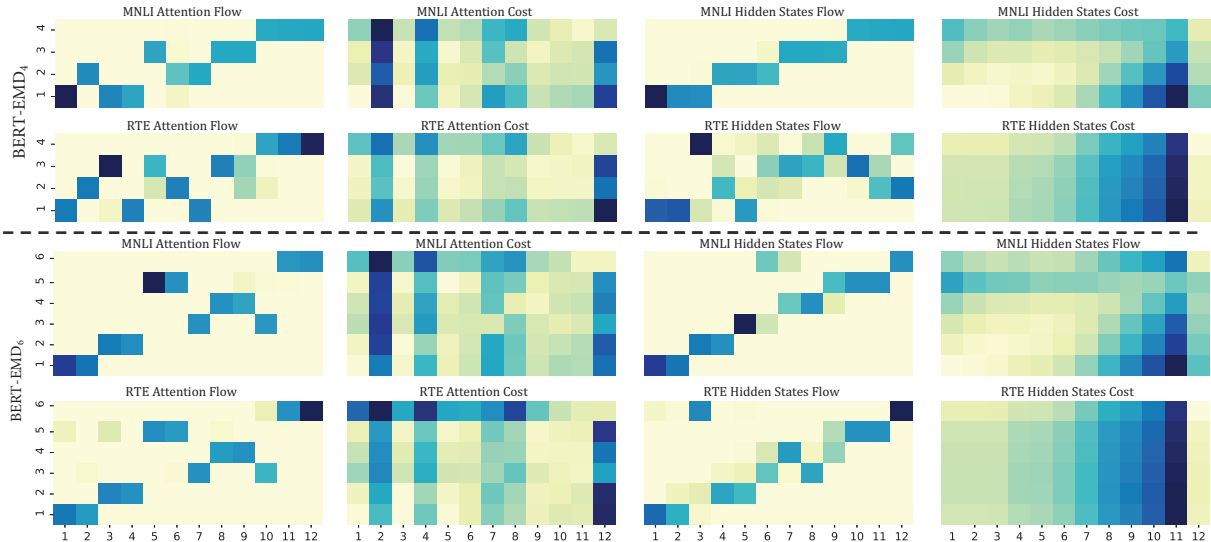


Figure 2: The visualization of flow matrices (\mathbf{F}) and distance matrices (\mathbf{D}) in developing BERT-EMD₄ (above) and BERT-EMD₆ (below) for two examples from MNLi and RTE tasks, respectively. The abscissa represents the Transformer layers of BERT_{BASE12}, and the ordinate represents the Transformer layers of BERT-EMD₄/BERT-EMD₆. The color depth represents the values (weights) of the layers.

lem, which helps to learn an optimal many-to-many layer mapping. The cost attention also contributes to the effectiveness of BERT-EMD. This verifies that the cost attention can further improve the many-to-many layer mapping by learning the importance of each teacher layer in guiding the student network. It is noteworthy that when removing the EMD module in the many-to-many layer mapping process, our w/o EMD₄ performs slightly worse than TinyBERT₄ on the MNLi and QQP tasks. This is because we cannot automatically control the information flow during the many-to-many layer mapping without using EMD, which further verifies the effectiveness of EMD in the many-to-many layer mapping process.

5.3 Visualization of Compression Process

To better understand the many-to-many layer mapping process, we illustrate the flow matrices \mathbf{F} and cost (distance) matrices \mathbf{D} in developing BERT-EMD₄ (above) and BERT-EMD₆ (below) for two examples from MNLi and RTE tasks, respectively. In Figure 2, we report the averaged values of the flow and cost matrices of the entire epoch that achieves the best performance on the validation set with heat maps.

From the results in Figure 2, we have several key observations. First, different tasks could emphasize different teacher layers in compressing the Transformer. The diagonal positions of the matrices are almost always important for the MNLi

task, which exhibits similar trends with TinyBERT with the one-to-one ‘‘Skip’’ layer mapping strategy. However, for the RTE task, each student Transformer layer can learn from any teacher Transformer layers. The previous one-to-one layer mapping methods cannot take full advantage of the teacher network. This argument can be verified by the quantitative results in Table 1, where our BERT-EMD has a much larger improvement on RTE than on MNLi over TinyBERT. Second, comparing BERT-EMD₄ and BERT-EMD₆, we can observe that BERT-EMD₄ usually needs to learn more comprehensive information from skipped teacher Transformer layers, resulting in more divergent many-to-many layer mappings.

6 Conclusion

In this paper, we propose a novel BERT compression method based on many-to-many layer mapping by Earth Mover’s Distance (EMD). To our knowledge, BERT-EMD is the first work that allows each intermediate student layer to learn from any intermediate teacher layers adaptively. In addition, a cost attention mechanism is designed to further improve the model’s performance and accelerate convergence time by learning the layer weights used in EMD automatically. Extensive experiments on GLUE tasks show that BERT-EMD can achieve competitive performances with the large BERT-Base model while significantly reducing the model size and inference time.

Acknowledgement

Min Yang was partially supported by National Natural Science Foundation of China (No. 61906185), the Natural Science Foundation of Guangdong Province of China (No. 2019A1515011705, 2018A030313943), the Youth Innovation Promotion Association of CAS.

References

- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Frank L Hitchcock. 1941. The distribution of a product from several sources to numerous localities. *Journal of mathematics and physics*, 20(1-4):224–230.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- 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*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119.
- Subhabrata Mukherjee and Ahmed Hassan Awadallah. 2019. Distilling transformers into simple neural networks with unlabeled transfer data. *arXiv preprint arXiv:1910.01769*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (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. *arXiv preprint arXiv:1802.05365*.
- Svetlozar T Rachev. 1985. The monge–kantorovich mass transference problem and its stochastic applications. *Theory of Probability & Its Applications*, 29(4):647–676.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf.
- Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. 2000. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121.
- 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*.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic bert for resource-limited devices. *arXiv preprint arXiv:2004.02984*.
- 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*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *arXiv preprint arXiv:2002.10957*.
- Ziheng Wang, Jeremy Wohlwend, and Tao Lei. 2019. Structured pruning of large language models. *arXiv preprint arXiv:1910.04732*.
- Shinji Watanabe, Takaaki Hori, Jonathan Le Roux, and John R Hershey. 2017. Student-teacher network learning with enhanced features. In *ICASSP*, pages 5275–5279. IEEE.
- Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. 2020. Bert-of-theseus: Compressing bert by progressive module replacing. *arXiv preprint arXiv:2002.02925*.
- 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 5754–5764.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. Ernie: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*.