

TR-BERT: Dynamic Token Reduction for Accelerating BERT Inference

Deming Ye^{1,2}, Yankai Lin⁴, Yufei Huang^{1,2}, Maosong Sun^{1,3*}

¹Department of Computer Science and Technology, Tsinghua University, Beijing, China
Institute for Artificial Intelligence, Tsinghua University, Beijing, China

Beijing National Research Center for Information Science and Technology

²State Key Lab on Intelligent Technology and Systems, Tsinghua University, Beijing, China

³Beijing Academy of Artificial Intelligence

⁴Pattern Recognition Center, WeChat AI, Tencent Inc.

ydm18@mails.tsinghua.edu.cn

Abstract

Existing pre-trained language models (PLMs) are often computationally expensive in inference, making them impractical in various resource-limited real-world applications. To address this issue, we propose a dynamic token reduction approach to accelerate PLMs' inference, named TR-BERT, which could flexibly adapt the layer number of each token in inference to avoid redundant calculation. Specially, TR-BERT formulates the token reduction process as a multi-step token selection problem and automatically learns the selection strategy via reinforcement learning. The experimental results on several downstream NLP tasks show that TR-BERT is able to speed up BERT by 2-5 times to satisfy various performance demands. Moreover, TR-BERT can also achieve better performance with less computation in a suite of long-text tasks since its token-level layer number adaption greatly accelerates the self-attention operation in PLMs. The source code and experiment details of this paper can be obtained from <https://github.com/thunlp/TR-BERT>.

1 Introduction

Large-scale pre-trained language models (PLMs) such as BERT (Devlin et al., 2019), XLNet (Yang et al., 2019) and RoBERTa (Liu et al., 2019) have shown great competence in learning contextual representation of text from large-scale corpora. With appropriate fine-tuning on labeled data, PLMs have achieved promising results on various NLP applications, such as natural language inference (Zhang et al., 2020b), text classification (Sun et al., 2019a) and question answering (Talmor and Berant, 2019).

Along with the significant performance improvements, PLMs usually have substantial computational cost and high inference latency, which presents challenges to their practicalities in resource-limited real-world applications, such as

real-time applications and hardware-constrained mobile applications. Even worse, these drawbacks become more severe in long-text scenarios because self-attention operation in PLMs scales quadratically with the sequence length. Therefore, researchers have made intensive efforts in PLM's inference acceleration recently. The mainstream approach is to reduce the layer number of PLMs such as knowledge distillation models (Sanh et al., 2019; Sun et al., 2019b), and adaptive inference models (Xin et al., 2020; Liu et al., 2020). Such layer-wise pruning reduces a tremendous amount of computation, but it sacrifices the models' capability in complex reasoning. Previous works (Sanh et al., 2019; Sun et al., 2019b) have found that the shallow model usually performs much worse on the relatively complicated question answering tasks than text classification tasks. It is straightforward that pruning the entire layer of PLMs may not be an optimal solution in all scenarios.

In this paper, we introduce a dynamic token reduction method TR-BERT to find out the well-encoded tokens in the layer-by-layer inference process, and save their computation in subsequent layers. The idea is inspired by recent findings that PLMs capture different information of words in different layers (e.g., BERT focuses on the word order information (Lin et al., 2019) in the bottom layers, obtains the syntactic information (Hewitt and Manning, 2019) in the middle layers, and computes the task-specific information in the top layers (Rogers et al., 2020)). Hence, we could adapt different tokens to different layers according to their specific roles in the context.

As shown in Figure 1, TR-BERT formulates the token reduction process as a multi-step selection problem. Specially, for each selection phase, TR-BERT finds out the words that require high-level semantic representations, and then selects them to higher layers. The main challenge in TR-BERT is how to determine each token's importance for text

*Corresponding author: M. Sun (sms@tsinghua.edu.cn)

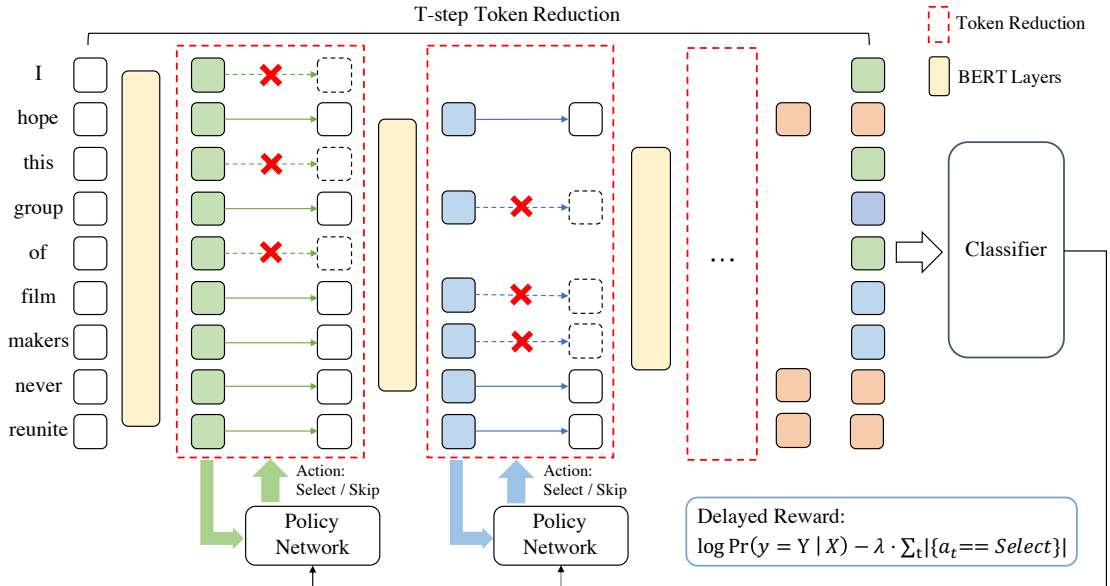


Figure 1: An illustration of the TR-BERT. TR-BERT gradually compresses the text sequence as the layer gets deeper. In RL training, we sample actions of *Select* or *Skip* for each token via the policy network.

understanding in the token selection. It is highly task-dependent and requires to consider the correlation and redundancy among various tokens. TR-BERT employs the reinforcement learning (RL) method to learn the dynamic token selection strategy automatically. After the token reduction, the RL reward involves the confidence of the classifier’s prediction based on the pruned network to reflect the quality of token selection. Moreover, we also add a penalty term about the number of selected tokens to the reward, by adjusting which, TR-BERT can utilize the different pruning intensities in response to various performance requirements. In TR-BERT, by selecting a few important tokens to go through the entire pipeline, the inference speed turns much faster and no longer grows quadratically with the sequence length.

We conduct experiments on eleven NLP benchmarks. Experimental results show that TR-BERT can accelerate BERT inference by 2-5 times to meet various performance demands, and significantly outperform previous baseline methods on question answering tasks. It verifies the effectiveness of the dynamic token reduction strategy. Moreover, benefiting from the long-distance token interaction, TR-BERT with 1,024 input length reaches higher performance with less inference time compared to the vanilla BERT in a suite of long-text tasks.

2 Background and Pilot Analysis

To investigate the potential impact of the token reduction in PLMs, we first introduce the Trans-

former architecture. After that, we conduct pilot experiments as well as empirical analyses for the lower and upper bound of the token reduction in this section.

The Transformer architecture (Vaswani et al., 2017) has been widely adopted by the pre-trained language models (PLMs) for inheriting its high capacity. Basically, each Transformer layer wraps a Self-Attention module (Self-ATT) and a Feed-Forward-Network module (FFN) by the residual connection and layer normalization. Formally, given a sequence of n words, the hidden state of the i -th layer, $\mathbf{H}_i = (h_1, h_2, \dots, h_n)$, is computed from the previous layer state:

$$\begin{aligned} \mathbf{M}_{i-1} &= \text{LN}(\mathbf{H}_{i-1} + \text{Self-ATT}(\mathbf{H}_{i-1})), \\ \mathbf{H}_i &= \text{LN}(\mathbf{M}_{i-1} + \text{FFN}(\mathbf{M}_{i-1})), \end{aligned} \quad (1)$$

where $i \in [1, L]$, L is the number of stacked Transformer layers, LN denotes the LayerNorm layer. For each Transformer layer, the complexity of the Self-Attention module scales quadratically with the sequence length. Therefore, the speed of Transformer architecture will decline heavily when the sequences become longer.

Previous findings (Rogers et al., 2020) reveal that some words, such as function words, do not require high-layer modeling, since they store little information and have been well handled by PLMs in bottom layers. Hence, selecting only the important words for high-layer computation may be a possible way to accelerate the PLMs’ inference.

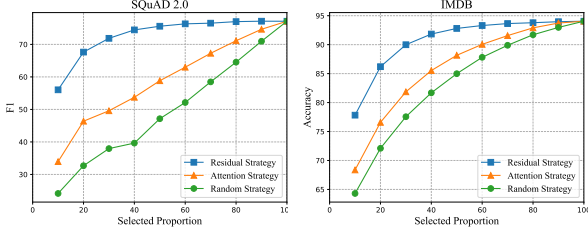


Figure 2: Performance under different selected proportion of tokens and different heuristic strategies.

To verify this assumption, we conduct a theoretical token elimination experiment in question answering (on SQuAD 2.0 (Rajpurkar et al., 2018)) and text classification (on IMDB (Maas et al., 2011)). We use the full-layer representations for the selected tokens and the early-layer representation of the deleted tokens for the prediction. To be specific, we eliminate tokens immediately after the l -th layer and adopt the following three strategies to select the retained tokens:

Random Strategy (Lower Bound) selects tokens randomly, assuming that all tokens are equivalent for understanding.

Residual Strategy (Upper Bound) directly utilizes the model prediction of the original model to guide the token selection. Specially, we define a token’s importance according to the influence on the model prediction when it’s not selected. When substituting the r -th layer representation \mathbf{H}_r with the l -th layer representation \mathbf{H}_l ($r > l$), we define the approximate variation to model loss as the token importance: $\mathbf{I} = \frac{\partial \text{loss}}{\partial \mathbf{H}_r}(\mathbf{H}_r - \mathbf{H}_l)$. Here, we set $r = 9$ since other values get a little worse results. Note that we could not obtain the model loss in the prediction stage. Hence, the Residual Strategy could be viewed as an upper bound of token selection to some extent when we ignore the correlation and redundancy among the selected tokens.

Attention Strategy is adopted by PoWER-BERT (Goyal et al., 2020) and L-Adaptive (Kim and Cho, 2020). It accumulates the attention values from other tokens to a given token. It selects the tokens receiving the greatest attentions, considering them responsible for retaining and disseminating the primary information of the context.

As shown in Figure 2, both Attention Strategy and Residual Strategy achieve considerable results, which demonstrates that to select important tokens is feasible for accelerating the inference of PLMs. Besides, the Residual Strategy outperforms the Attention strategies by a margin, especially at the

low token remaining proportion (+31.8% F1 on SQuAD 2.0 and +9.5% accuracy on IMDB when selecting 10% tokens). It suggests that the accumulated attention values still cannot well reflect tokens’ importance in text understanding, which requires further explorations.

3 Methodology

In this section, we present TR-BERT, which adopts a cascade token reduction to prune the BERT model at token-level granularity dynamically. In a one-step token reduction process, TR-BERT estimates the importance of each token, reserves the important ones, and delivers them to the higher layer. To better select important tokens for text understanding while satisfying various acceleration requirements, we employ the reinforcement learning (RL) method to automatically learn a dynamic token selection strategy.

3.1 Model Architecture

Figure 1 shows the model architecture of TR-BERT. To inherit the high capacity from the PLMs, TR-BERT keeps the same architecture as BERT. Differently, as the layer gets deeper, TR-BERT gradually shortens the sequence length via token reduction modules, aiming to reduce the computational redundancy of unimportant tokens.

The token reduction modules are required to measure the importance of tokens and offer an integral selection scheme. Due to the lack of direct supervision, we employ the policy network for training the module, which adopts a stochastic policy and uses a delayed reward to guide the policy learning. In one-step reduction, we perform action sampling for the current sequence. The selected tokens are conveyed to the next Transformer layer for further computation. In contrast, the unselected tokens are terminated with their representation remaining unchanged. After all the actions are decided, we fetch each token’s representation from the layer where it terminated, and compute the golden label’s likelihood as a reward. To be specific, we introduce state, action, reward, and objective function as follows:

State State s_t consists of the token representations inherited from the previous layer before the t -th token reduction layer.

Action We adopt two alternative actions for each token, $\{Select, Skip\}$, where the token can be selected for further computation or be skipped to the

final layer. We implement the policy network as a two-layer feed-forward network with GeLU activation (Hendrycks and Gimpel, 2017):

$$\pi(a_t|s_t; \theta) = \sigma(\mathbf{W}_2(\text{GeLU}(\mathbf{W}_1\mathbf{H}_{s_t} + b_1)) + b_2), \quad (2)$$

where a_t denotes the action at state s_t for sequence representation $\mathbf{H}_{s_t} = \{h_1, h_2, \dots, h_n\}$ at t -th reduction, $\theta = \{\mathbf{W}_1, \mathbf{W}_2, b_1, b_2\}$ are trainable parameters, and $\sigma(\cdot)$ is sigmoid activation function. For the selected token set $\{t_1, t_2, \dots, t_{n^*}\}$, where $n^* \leq n$, we conduct a Transformer layer operation on their corresponding representations:

$$\mathbf{H}' = \text{Transformer}([h_{t_1}, h_{t_2}, \dots, h_{t_{n^*}}]). \quad (3)$$

For the selected tokens, their representation \mathbf{H}' is conveyed to the next layer for further feature extraction and information aggregation. For the other skipped tokens, their representations in the current layer are regarded as their final representations.

Reward Aiming to select significant tokens for making a precise decision in the prediction layer, we adopt the likelihood of predicting the golden label as a reward. For example, when classifying the input sequence \mathbf{X} , we use the models' predicting probability of the ground-truth label \mathbf{Y} to reflect the quality of the token selection. In addition, to encourage the model to delete more redundant tokens for accelerating, we include an additional punitive term by counting the number of selected tokens. Hence, the overall reward R is defined as:

$$R = \log \text{Pr}(y = \mathbf{Y}|\mathbf{X}) - \lambda \cdot \sum_t |\{a_t = \text{Select}\}|, \quad (4)$$

where $\sum_t |\{a_t = \text{Select}\}|$ denotes the total number of the selected tokens in all token reduction modules, and λ is a harmonic coefficient to balance two reward terms.

Objective Function We optimize the policy network to maximize the expected reward. Formally, our objective function is defined as:

$$\begin{aligned} J(\theta) &= \mathbb{E}_{(s_t, a_t) \sim \pi(a_t|s_t; \theta)} r[(s_1, a_1) \dots (s_T, a_T)] \\ &= \sum_{(s_1, a_1) \dots (s_T, a_T)} \prod_t \pi_\theta(a_t|s_t) \cdot R, \end{aligned} \quad (5)$$

where T is the number of states. According to the REINFORCE algorithm (Williams, 1992) and policy gradient method (Sutton et al., 1999), we update network with the policy gradient as below:

$$\nabla_\theta J(\theta) = \sum_{t=1}^T R \cdot \nabla_\theta \log \pi_\theta(a_t|s_t). \quad (6)$$

3.2 Model Training

Our policy network is integrated into the original Transformer network, and we train both of them simultaneously. The entire training process involves three steps:

- (1) Fine-tune the PLM model for downstream tasks with the task-specific objective;
- (2) Freeze all the parameters except that of the policy network, conduct reinforcement learning (RL), and update the policy network to learn token reduction strategy;
- (3) Unfreeze all parameters and train the entire network with the task-specific objective and RL objective simultaneously.

Due to the large searching space, RL learning is difficult to converge. We adopt imitation learning (Hussein et al., 2017) for warming up the training of the policy network. To be specific, in the RL training, we sample several action sequences via the policy network to compute rewards. And we guide the optimization direction by providing heuristic action sequences sampled by the Residual Strategy during the early training period, which could roughly select the most important tokens. The heuristic action sequence is defined as selecting the top K important tokens and skipping the others, where K is defined as the expected selected number of the current policy network. In our preliminary experiment, both the heuristic action sequence and expected selected number mechanism are beneficial to the stable training.

To further improve the performance of our pruned model, we also adopt Knowledge Distillation (KD) (Hinton et al., 2015) to transfer knowledge from the intact original fine-tuned model.

3.3 Complexity Analysis

For a Transformer layer with a hidden size of d and an input sequence of n tokens, the Self-Attention module consumes $O(n^2d)$ time and memory complexity while the Feed-Forward Network takes $O(nd^2)$. That is, our token reduction gains near-linear speedup when n is relatively smaller than d . Therefore, when the input sequence gets longer, such as up to 1,024 tokens, our method can enjoy a more effective speedup.

In the RL training, we compute loss on the pruned model, so the acceleration is still valid for this stage. Since we focus on accelerating BERT inference, we consider the extra training consumption on the pruned model is acceptable.

4 Experiment

In this section, we first introduce the baseline models and the evaluation datasets. After that, we verify the effectiveness of TR-BERT on eleven NLP benchmarks. Finally, we conduct a detailed analysis and case study on TR-BERT to investigate the selected tokens’ characteristics.

4.1 Baselines

We adopt two pre-trained models and three pruned networks as our baselines for comparison:

BERT (Devlin et al., 2019) is a Transformer-based pre-trained model. We use the BERT_{BASE} model¹, which consists of 12 Transformer layers and supports a maximum sequence length of 512.

BERT_L is our implemented BERT, which can support input sequences with up to 1,024 tokens. We initialize the parameters of BERT_L with that of BERT, where the additional position embedding is initialized with the first 512 ones. After that, we continue to train it on Wikipedia² for 22k steps.

DistilBERT (Sanh et al., 2019) is the most popular distilled version of BERT, which leverages the knowledge distillation to learn knowledge from the BERT model. We use the 6-layer DistilBERT released by Hugging Face³. In addition, we use the same method to distill BERT with 3 layers to obtain DistilBERT₃.

DeFormer (Cao et al., 2020) is designed for question answering, which encodes questions and passages separately in lower layers. It pre-computes all the passage representation and reuses them to speed up the inference. In our experiments, we do not count DeFormer’s pre-computation.

PoWER-BERT (Goyal et al., 2020) is mainly designed for text classification, which also decreases the length of a sequence as layer increases. It adopts the Attention Strategy to measure the significance of each token and always selects tokens with the highest attention. Given a length penalty, PoWER-BERT searches a fixed length pruning configuration for all examples.

DynaBERT (Hou et al., 2020) can not only adjust model’s width by varying the number of attention heads, but also provide an adaptive layer depth to satisfy different requirements. For a given speed demand, we report its best performance with all the feasible width and depth combination options.

¹<https://github.com/google-research/bert>

²<https://en.wikipedia.org/>

³<https://github.com/huggingface/transformers>

4.2 Datasets

To verify the effectiveness of reducing the sequence length, we evaluate TR-BERT on several tasks with relatively long context, including question answering and text classification. Table 1 shows the context length of these datasets. We adopt seven question-answering datasets, including SQuAD 2.0 (Rajpurkar et al., 2018), NewsQA (Trischler et al., 2017), NaturalQA (Kwiatkowski et al., 2019), RACE (Lai et al., 2017), HotpotQA (Yang et al., 2018), TriviaQA (Joshi et al., 2017) and WikiHop (Welbl et al., 2018). And we also evaluate models on four text classification datasets, including YELP.F (Zhang et al., 2015), IMDB (Maas et al., 2011), 20NewsGroups (20News.) (Lang, 1995), and Hyperpartisan (Hyperp.) (Kiesel et al., 2019). Among them, HotpotQA, TriviaQA and WikiHop possess abundant contexts for reading, while the performance of question answering (QA) models heavily relies on the amount of text they read. To fairly compare BERT and BERT_L, we split the context into slices and apply a shared-normalization training objective (Clark and Gardner, 2018) to produce a global answer candidate comparison across different slices for the former two extractive QA datasets. And we average the candidate scores in all slices for WikiHop. Details of all datasets are shown in the Appendix.

4.3 Experimental Settings

We adopt a maximum input sequence length of 384 for SQuAD 2.0, 1,024 for long-text tasks and 512 for others. We use the Adam optimizer (Kingma and Ba, 2015) to train all models. The detailed training configuration is shown in the Appendix.

For the RL training, we sample 8 action sequences each time and average their rewards as the reward baseline. In the second training process which aims to warm up the policy network, we employ 20% imitation learning steps for question answering tasks and 50% steps for text classification tasks. We search the number of token reduction module $T \in [1, 2, 3]$. And we find the models with $T = 2$ gets similar quality and speed trade-offs as the models with $T = 3$, and both of them perform better than models with $T = 1$. Thus we adopt $T = 2$ for simplification. We denote the pruned models from BERT, BERT_L and DistilBERT₆ as TR-BERT₁₂, TR-BERT_L, TR-BERT₆, respectively. For BERT and BERT_L, we attach the token reduction modules before the second and the sixth layers.

# Tokens	SQuAD 2.0	NewsQA	NaturalQA	RACE	HotpotQA	TriviaQA	WikiHop	YELP.F	20News.	IMDB	Hyperp.
Average	152	656	248	381	1,988	3,117	1,499	179	551	264	755
95th percentile	275	878	1,088	542	2,737	4,004	2,137	498	1,769	679	2,005

Table 1: Average and 95th percentile of context length of datasets in wordpieces.

Model	SQuAD 2.0		NewsQA		NaturalQA		RACE		YELPF		20NewsGroups		IMDB	
	F1	FLOPs	F1	FLOPs	F1	FLOPs	Acc.	FLOPs	Acc.	FLOPs	Acc.	FLOPs	Acc.	FLOPs
BERT	77.12	1.00x	66.82	1.00x	78.32	1.00x	66.30	1.00x	69.93	1.00x	86.66	1.00x	94.05	1.00x
DistilBERT ₆	68.17	2.00x	63.56	2.00x	76.21	2.00x	52.63	2.00x	69.42	2.00x	85.85	2.00x	93.11	2.00x
PoWER-BERT	—	—	—	—	—	—	—	—	67.37	2.75x	86.51	2.91x	92.02	3.05x
DeFormer	71.41	2.19x	60.68	2.01x	74.34	2.34x	64.27	2.19x	—	—	—	—	—	—
DynaBERT	74.83	2.00x	66.13	2.00x	78.14	2.00x	65.38	2.00x	69.96	2.00x	86.03	2.00x	94.00	2.00x
TR-BERT₁₂	75.66	2.08x	67.18	2.05x	79.03	2.51x	65.00	2.16x	69.97	2.19x	87.44	4.22x	93.60	2.26x
TR-BERT₆	71.75	3.07x	65.36	2.96x	78.11	3.74x	53.40	4.10x	70.04	2.83x	86.58	5.90x	92.64	4.13x

Table 2: Comparison of performance and FLOPs (speedup) between TR-BERT and baselines.

For DistilBERT₆, we insert the token reduction modules before the second and the fourth layers.

To avoid the pseudo improvement by pruning padding for TR-BERT, we evaluate all models with input sequences without padding to the maximum length. For each dataset, we report the F1 scores or accuracy (Acc.), and the FLOPs speedup ratio compared to the BERT model. The model’s FLOPs are consistent in the various operating environment. Therefore, it is convenient to estimate and compare the models’ inference time by FLOPs.

4.4 Overall Results

The comparison between TR-BERT and the baselines are shown in Table 2 and Figure 3. We adjust the length penalty coefficient of TR-BERT for an intuitional comparison. From the experimental results, we have the following observations:

(1) TR-BERT₁₂ achieves higher performance while using less computation on all span-extraction QA datasets compared to all the baselines. For example, TR-BERT₁₂ outperforms DynaBERT by 1.8 F1 with faster speed. TR-BERT₁₂ even achieves better performance than BERT at low speedup rate, which demonstrates that discarding some redundant information in the top layer helps to find the correct answer. For multiple-choice RACE, TR-BERT₁₂ achieves better performance than DeFormer while doesn’t need to pre-compute the passage representation.

(2) TR-BERT₆ performs better than PoWER-BERT by a margin in text classification tasks. It shows that the fixed pruning configuration and the attention-based selection strategy adopted by PoWER-BERT may not be flexible to accelerate inference for various input sequences. In contrast,

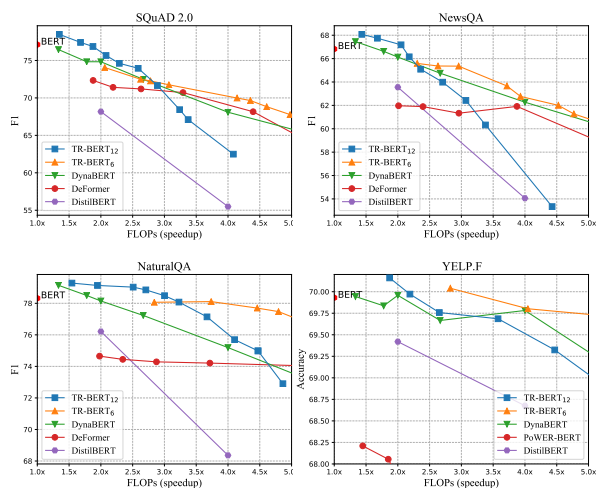


Figure 3: Quality and efficiency trade-offs for TR-BERT₁₂ and TR-BERT₆.

our dynamic token selection can automatically determine the proper pruning length and tokens for each example according to the actual situation, which leads to a more effective model acceleration.

Overall, TR-BERT retains most of BERT’s performance though it omits lots of token interactions in the top layers. It shows that TR-BERT learns a satisfactory token selection strategy through reinforcement learning, and could effectively reduce the redundant computation of tokens that have been extracted enough information in the bottom layers.

4.5 Fuse Layer-wise and Token-wise Pruning

Since layer-wise pruning and token-wise pruning are compatible, we also explore the incorporation of these two pruning strategies. We apply our dynamic token reduction on the 6-layer DistilBERT to obtain TR-BERT₆. The trade-off comparison of

Model	HotpotQA		TriviaQA	
	F1	FLOPs	F1	FLOPs
BERT	57.33	1.00x	68.75	1.00x
BERT _L	65.45	0.91x	69.69	0.92x
TR-BERT_L	65.57	1.56x	70.41	1.20x

Model	WikiHop		Hyperparisan	
	Acc.	FLOPs	Acc.	FLOPs
BERT	67.67	1.00x	74.39	1.00x
BERT _L	68.49	0.92x	76.83	0.92x
TR-BERT_L	67.78	4.65x	74.39	1.96x

Table 3: Comparison performance and FLOPs (speedup) between TR-BERT_L and BERTs with different maximum input sequence.

TR-BERT₁₂ and TR-BERT₆ is shown in Figure 3, from which we have the following findings:

(1) In general, as the speedup ratio increases, the performance of all models decrease, which indicates that retaining more token information usually results in a more potent model.

(2) TR-BERT₆ consistently outperforms TR-BERT₁₂ on all tasks at a high speedup ratio. In this situation, the budget doesn't allow enough tokens to go through the top layers. TR-BERT₆ makes a more elaborate pruning than TR-BERT₁₂ at bottom layers to obtain a better effectiveness.

(3) At low speedup ratio, TR-BERT₁₂ performs better than TR-BERT₆ on the question answering tasks, but worse on the text classification tasks. In general, a deep Transformer architecture can offer multi-turn feature extraction and information propagation, which can meet the complex reasoning requirements for question answering. In contrast, the result of text classification usually depends on the keywords in the context, for which a shallow model is an affordable solution. To obtain a better trade-off, we can flexibly employ a deep and narrow model for question answering and a shallow and wide model for text classification.

4.6 Results on Long-text Tasks

With token pruning, TR-BERT is able to process a longer sequence. We apply our dynamic token pruning strategy on BERT_L, which can process sequence with up to 1,024 tokens, to obtain TR-BERT_L, and conduct experiments on four datasets with longer documents, including HotpotQA, TriviaQA, WikiHop and Hyperparisan. Results on long-text tasks are shown in Table 3, from which we have the following observations:

(1) BERT_L achieves better performance than BERT, especially on HotpotQA and WikiHop,

which require the long-range multi-hop reasoning;

(2) Compared to the vanilla BERT, TR-BERT_L achieves 8.2% F1 improvement with 1.56x speedup on HotpotQA, obtains 1.7% F1 improvement with 1.24x speedup on TriviaQA, gains 4.65x speedup on WikiHop and 1.96x speedup on Hyperparisan without performance drops. Compared to BERT which can only deal with up to 512 tokens at a time, BERT_L considers a longer-range token interaction and obtains a more complete reasoning chain. However, the running time of BERT_L also increase as the input sequence's length extends, which poses a challenge to the utilization of longer text. TR-BERT_L inherits the broader view from BERT_L to get a better performance with a faster inference. Moreover, the inference acceleration effect of TR-BERT_L is relatively better than TR-BERT within 512 tokens, which is coincident to the above complexity analysis section. With a longer sequence, TR-BERT can achieve extra speedup, because it significantly saves the time of the Self-Attention module, which demonstrates that TR-BERT can be further applied to process much longer tokens with limited computation.

4.7 Case Study

To investigate the characteristics of the selected tokens, we conduct a detailed case study on various datasets. As shown in Table 4, TR-BERT chooses to abandon the function word, such as *the*, *and*, *with*, in the first token reduction module as the first module is placed at the bottom layer of BERT. The second token reduction module is placed at the middle layer of BERT, and we could observe that it is used to retaining task-specific tokens. In the first example about question answering, the second token reduction module maintains the whole question and the question-related tokens from the context for further propagating messages. In the second and third examples about movie review sentimental classification, the second token reduction module chooses to select sentimental words, such as *great*, *excited*, *disappointed* to determine whether the given sequence is positive or negative.

Although we train the token reduction module without direct human annotations, TR-BERT can remain the meaningful tokens in the bottom layer and select the higher layer's task-relevant tokens. It demonstrates that the pruned network's ground-truth probability is an effective signal to facilitate the reinforcement learning for token selection.

Dataset	Example
SQuAD 2.0	Question: How long did Western Europe control Cyprus? Paragraph: The conquest of Cyprus by the Anglo-Norman forces of the Third Crusade opened a new chapter in the history of the island , which would be under Western European domination for the following 380 years . Although not part of a planned operation , the conquest had much more permanent results than initially expected . Answer: 380 years
IMDB Positive	The Buddy Holly Story is a great biography with a super performance from Gary Busey . Busey did his own singing for this film and he does a great job .
IMDB Negative	I was really excited when I read “The Canterville Ghost” would be shown on TV . However , I was deeply disappointed . I loved the original story written by Oscar Wilde and sadly nothing of that was transferred by the movie .

Table 4: Case study of the token reduction. Skipped tokens in first token reduction module are colored with light blue. Skipped tokens in the second token reduction module are colored with blue. The final selected token are colored with dark blue.

5 Related Work

Researchers have made various attempts to accelerate the inference of PLMs, such as quantization (Shen et al., 2020; Zhang et al., 2020a), attention head pruning (Michel et al., 2019; Hou et al., 2020), dimension reduction (Sun et al., 2020; Chen et al., 2020), and layer reduction (Sanh et al., 2019; Sun et al., 2019b; Jiao et al., 2019). In current studies, one of the mainstream methods is to dynamically select the layer number of Transformer layers to make a on-demand lighter model (Fan et al., 2020; Xin et al., 2020; Liu et al., 2020). However, these methods operate at the whole text and they cannot perform pruning operations in a smaller granularity, such as the token-level granularity.

To consider the deficiencies of layer-level pruning methods, researchers decide to seek solutions from a more meticulous perspective by developing methods to extend or accelerate the self-attention mechanism of the Transformer. For example, Sparse Transformer (Child et al., 2019), LongFormer (Beltagy et al., 2020) and Big Bird (Zaheer et al., 2020) employ the sparse attention to allow model to handle long sequences. However, these methods only reduce the CUDA memory but cannot be not faster than the full attention. Besides, researchers also explore the feasibility of reducing the number of involved tokens. For example, Funnel-Transformer (Dai et al., 2020) reduces the sequence length with pooling for less computation, and finally up-samples it to the full-length representation. Universal Transformer (Dehghani et al., 2019) builds a self-attentive recurrent sequence model, where each token uses a dynamic halting layer. And DynSAN (Zhuang and Wang, 2019) ap-

plies a gate mechanism to measure the importance of tokens for selection. Spurred by these attempts and positive results, we introduce TR-BERT in this study, which can creatively prune the network at the token level. To be specific, our work aims to accelerate the Transformer by deleting tokens gradually as the layer gets deeper. Compared with these models, TR-BERT is easy to adapt to the current PLMs models without a significant amount of pre-training and is flexible to adjust the model speed according to different performance requirements.

The main idea of TR-BERT is to select essential elements and infuse more computation on them, which is widely adopted in various NLP tasks. ID-LSTM (Zhang et al., 2018) selects important and task-relevant words to build sentence representation for text classification. SR-MRS (Nie et al., 2019) retrieves the question-related sentences to reduce the size of reading materials for question answering. TR-BERT can be viewed as a unified framework on the Transformer for the important element selection, which can be easy to be applied in wide-range tasks.

6 Conclusion and Future Work

In this paper, we propose a novel method for accelerating BERT inference, called TR-BERT, which prunes BERT at token-level granularity. Specifically, TR-BERT utilizes reinforcement learning to learn a token selection policy, which is able to select general meaningful tokens in the bottom layers and select task-relevant tokens in the top layers. Experiments on eleven NLP tasks demonstrate the effectiveness of TR-BERT as it accelerates BERT inference by 2-5 times for various performance demand. Besides, TR-BERT achieves a better quality

and speed trade-off on long-text tasks, which shows its potential to process large amounts of information in the real-world applications.

In the future, we would like to attempt to apply TR-BERT in the pre-training process of PLMs. Through the automatically learned token reduction module, it is possible to reveal how BERT stores syntactic and semantic information in various tokens and different layers. And it's also worth speeding up the time-consuming pre-training process.

Acknowledgement

This research is mainly supported by Science & Tech Innovation 2030 Major Project "New Generation AI" (Grant no. 2020AAA0106500) as well as supported in part by a grant from the Institute for Guo Qiang, Tsinghua University.

References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *CoRR*, abs/2004.05150.
- Qingqing Cao, Harsh Trivedi, Aruna Balasubramanian, and Niranjan Balasubramanian. 2020. Deformer: Decomposing pre-trained transformers for faster question answering. In *Proceedings of ACL*, pages 4487–4497.
- Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. 2020. The lottery ticket hypothesis for pre-trained BERT networks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *CoRR*, abs/1904.10509.
- Christopher Clark and Matt Gardner. 2018. Simple and effective multi-paragraph reading comprehension. In *Proceedings of ACL*, pages 845–855.
- Zihang Dai, Guokun Lai, Yiming Yang, and Quoc V. Le. 2020. Funnel-transformer: Filtering out sequential redundancy for efficient language processing. *CoRR*, abs/2006.03236.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. 2019. Universal transformers. In *Proceedings of ICLR*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Angela Fan, Edouard Grave, and Armand Joulin. 2020. Reducing transformer depth on demand with structured dropout. In *Proceedings of ICLR*. OpenReview.net.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. MRQA 2019 shared task: Evaluating generalization in reading comprehension. pages 1–13.
- Saurabh Goyal, Anamitra Roy Choudhury, Venkatesan T. Chakaravarthy, Saurabh ManishRaje, Yogish Sabharwal, and Ashish Verma. 2020. Powerbert: Accelerating BERT inference for classification tasks.
- Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proceedings of ICLR*.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. In *Proceedings of NAACL-HLT*, pages 4129–4138.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531.
- Lu Hou, Lifeng Shang, Xin Jiang, and Qun Liu. 2020. Dynabert: Dynamic BERT with adaptive width and depth. *CoRR*, abs/2004.04037.
- Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. 2017. Imitation learning: A survey of learning methods. *ACM Computing Survey*, 50(2):21:1–21:35.
- 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. *CoRR*, abs/1909.10351.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of ACL*, pages 1601–1611.
- Johannes Kiesel, Maria Mestre, Rishabh Shukla, Emmanuel Vincent, Payam Adineh, David Corney, Benno Stein, and Martin Potthast. 2019. Semeval-2019 task 4: Hyperpartisan news detection. In *Proceedings of Workshop on NAACL-HLT*, pages 829–839.
- Gyuwan Kim and Kyunghyun Cho. 2020. Length-adaptive transformer: Train once with length drop, use anytime with search. *CoRR*, abs/2010.07003.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.

- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: a benchmark for question answering research. *TACL*, 7:452–466.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard H. Hovy. 2017. RACE: large-scale reading comprehension dataset from examinations. In *Proceedings of EMNLP*, pages 785–794.
- Ken Lang. 1995. Newsweeder: Learning to filter news. In *Machine Learning, Proceedings of the Twelfth International Conference on Machine Learning, Tahoe City, California, USA, July 9-12, 1995*, pages 331–339. Morgan Kaufmann.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. Open sesame: Getting inside bert’s linguistic knowledge. *CoRR*, abs/1906.01698.
- 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 ACL*, pages 6035–6044.
- 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. *CoRR*, abs/1907.11692.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of ACL*, pages 142–150.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *Proceedings of NeurIPS*, pages 14014–14024.
- Yixin Nie, Songhe Wang, and Mohit Bansal. 2019. Revealing the importance of semantic retrieval for machine reading at scale. In *Proceedings of EMNLP*, pages 2553–2566.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of ACL*, pages 784–789.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100, 000+ questions for machine comprehension of text. In *Proceedings of EMNLP*, pages 2383–2392.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how BERT works. *Trans. Assoc. Comput. Linguistics*, 8:842–866.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108.
- Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2020. Q-BERT: hessian based ultra low precision quantization of BERT. In *Proceedings of AAAI*, pages 8815–8821.
- Chi Sun, Luyao Huang, and Xipeng Qiu. 2019a. Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence. In *Proceedings of NAACL-HLT*, pages 380–385.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019b. Patient knowledge distillation for BERT model compression. In *Proceedings of ACL*, pages 4322–4331.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic BERT for resource-limited devices. In *Proceedings of ACL*, pages 2158–2170.
- Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. 1999. Policy gradient methods for reinforcement learning with function approximation. In *Proceedings of NeurIPS*, pages 1057–1063.
- Alon Talmor and Jonathan Berant. 2019. MultiQA: An empirical investigation of generalization and transfer in reading comprehension. In *Proceedings of ACL*, pages 4911–4921.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. Newsqa: A machine comprehension dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP, Rep4NLP@ACL 2017, Vancouver, Canada, August 3, 2017*, pages 191–200.
- 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 *Proceeding of NeurIPS*, pages 5998–6008.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Trans. Assoc. Comput. Linguistics*, 6:287–302.
- Ronald J. Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8:229–256.
- Ji Xin, Raphael Tang, Jaehun Lee, Yaoliang Yu, and Jimmy Lin. 2020. DeeBERT: Dynamic early exiting for accelerating BERT inference. In *Proceedings of ACL*, pages 2246–2251.

- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Proceedings of NeurIPS*, pages 5754–5764.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of EMNLP*, pages 2369–2380.
- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big bird: Transformers for longer sequences. *CoRR*, abs/2007.14062.
- Tianyang Zhang, Minlie Huang, and Li Zhao. 2018. Learning structured representation for text classification via reinforcement learning. In *Proceedings of AAAI*, pages 6053–6060. AAAI Press.
- Wei Zhang, Lu Hou, Yichun Yin, Lifeng Shang, Xiao Chen, Xin Jiang, and Qun Liu. 2020a. Ternarybert: Distillation-aware ultra-low bit BERT. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 509–521. Association for Computational Linguistics.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceeding of NeurIPS*, pages 649–657.
- Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020b. Semantics-aware BERT for language understanding. In *Proceedings of AAAI*, pages 9628–9635.
- Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul N. Bennett, and Saurabh Tiwary. 2020. Transformer-xh: Multi-evidence reasoning with extra hop attention. In *Proceedings of ICLR*.
- Yimeng Zhuang and Huadong Wang. 2019. Token-level dynamic self-attention network for multi-passage reading comprehension. In *Proceedings of ACL*, pages 2252–2262.

A Details of Datasets

We evaluate models on seven question-answering datasets, including SQuAD 2.0 (Rajpurkar et al., 2018), NewsQA (Trischler et al., 2017), NaturalQA (Kwiatkowski et al., 2019), RACE (Lai et al., 2017), HotpotQA (Yang et al., 2018), TriviaQA (Joshi et al., 2017) and WikiHop (Welbl et al., 2018). Besides, we evaluate models on four long-text classification datasets, including YELPF (Zhang et al., 2015), 20News-Groups (Lang, 1995), IMDB (Maas et al., 2011),

and Hyperpartisan (Kiesel et al., 2019). We use the MRQA (Fisch et al., 2019) version of NewsQA and NaturalQA. Details of all evaluation datasets are shown below:

SQuAD 2.0 (Rajpurkar et al., 2018) is a large-scale reading comprehension dataset. Compared to its former SQuAD 1.0 (Rajpurkar et al., 2016), SQuAD 2.0 involves 54k unanswerable questions, which empirically makes the task much harder. For SQuAD 2.0, TR-BERT not only needs to find the question-relevant tokens, but also requires to check sufficient evidence to make a waiver decision when no answer is available. To predict the answer, we attach a span predictor on the top of BERTs and set the answer of the unanswerable question as a span of $[0, 0]$.

MRQA (Fisch et al., 2019) integrates several existing datasets to a unified format, which provides a single context within 800 tokens for each question, ensuring at least one answer could be accurately found in the context. We adopt the NewsQA and NaturalQA dataset from the MRQA benchmark.

RACE (Lai et al., 2017) is composed of 98k multiple-choice questions collected from English examinations. The model needs to figure out the correct answer from four options for a given question and passage. Passages in RACE cover a variety of topics, which can examine the generalization of our token selection.

HotpotQA (Yang et al., 2018) is an extractive question answering dataset, which requires multi-hop reasoning over multiple supporting documents for answering 113k questions. We adopt the full-wiki setting for HotpotQA, which requires models to find answers from a large-scale corpus. We employ the retriever and re-ranker in Transformer-XH (Zhao et al., 2020) to obtain question-related paragraphs and merge them into one document. In HotpotQA, models are required to reason over bridge entities or check multiple properties in different positions, which brings challenges to the token selection of TR-BERT in considering the global information. We concatenate two positive paragraphs and several negative paragraphs to make the reading material for training, which contains up to 1,024 tokens. And we concatenate the re-rank passages in order for evaluating. For the vanilla BERT, we apply a shared-normalization training objective (Clark and Gardner, 2018) to produce a global answer candidate comparison across two parts of the document.

TriviaQA (Joshi et al., 2017) has more than 95k question-answer pairs authored by Trivia enthusiasts. We use the Wikipedia setting of TriviaQA, which provides question-retrieved paragraphs from Wikipedia. We use the linear passage re-ranker in DocQA (Clark and Gardner, 2018) to re-rank these retrieved paragraphs and finally concatenate the first 1,024 tokens as a new reading document. We also employ the shared-normalization training objective (Clark and Gardner, 2018) for the vanilla BERT.

WikiHop (Welbl et al., 2018) consists of 51k questions, candidate answers, and supporting paragraphs triples. It requires models to find multi-hop reasoning chains for choosing the correct answer. Due to the long length of the concatenation of supporting paragraphs, we follow the processing tactic in Longformer (Beltagy et al., 2020), which splits the document into several parts and then averages their candidate scores.

YELP.F (Zhang et al., 2015) contains 1,569k samples with review texts, which are obtained from the YELP Dataset Challenge in 2015. Yelp Review Full (YELP.F) contains five star classes.

20NewsGroups (Lang, 1995) comprises around 18k newsgroups posts on 20 topics.

IMDB (Maas et al., 2011) consists of 50k informal movie reviews from the Internet Movie Database. Each review is annotated as *positive* or *negative*.

Hyperpartisan (Kiesel et al., 2019) aims to decide whether a news article text follows a hyperpartisan argumentation. Hyperparisan only contains 645 documents, which makes it a good testbed in a low-resource scenario. For the vanilla BERT, we adopt the max-pooling results of BERT sliding window. We split the data into training / validation / test set with a ratio of 9:1:1, run each model five times, and report the median performance.

B Training Configuration

We follow the configuration from previous work (Devlin et al., 2019; Beltagy et al., 2020; Liu et al., 2020) for fine-tuning BERT, DistilBERT and BERT_L. Hyperparameters are shown in Table 5.

To train TR-BERT, we first initialize TR-BERT with corresponding fine-tuned models, which are trained with a task-specific objective for N epochs. After that, we maintain the same learning rate, warmup proportion and batch size for the latter two-step training: (1) Freeze all the parameters

Dataset	Epoch	LR	WP	BSZ	Optimizer
SQuAD	2	$3e-5$	10%	32	Adam
NewsQA	5	$3e-5$	10%	32	Adam
NaturalQA	5	$3e-5$	10%	32	Adam
RACE	5	$5e-5$	10%	32	Adam
HotpotQA	6	$5e-5$	7.3%	32	Adam
TriviaQA	5	$3e-5$	10.6%	32	Adam
WikiHop	15	$3e-5$	1%	32	Adam
YELP.F	3	$3e-5$	10%	32	Adam
20NewsGroups	5	$3e-5$	10%	32	Adam
IMDB	5	$3e-5$	10%	32	Adam
Hyperparisan	15	$3e-5$	10%	32	Adam

Table 5: Hyperparameters of all the models in different datasets. LR: Learning rate; BSZ: Batch size; WP: Warmup proportion.

except that of the policy network and conduct reinforcement learning to update the policy network for $\lceil (N + 1)/2 \rceil$ epochs; (2) Unfreeze all parameters and train the entire network with the task-specific knowledge distillation objective and the reinforcement learning objective simultaneously for N epochs.

C Actual Wall Time

In practical applications, the wall time acceleration of TR-BERT is similar to the FLOPs acceleration. We evaluate our model on a single V100 GPU with 32 batch size on SQuAD. TR-BERT in Table 2 with 2.08x FLOPs speedup achieves 2.01x actual inference time speedup.