

SmartBERT: A Promotion of Dynamic Early Exiting Mechanism for Accelerating BERT Inference

Boren Hu , Yun Zhu , Jiacheng Li , Siliang Tang

Zhejiang University

{boren,zhuyun_dcd,lijiacheng,siliang}@zju.edu.cn

Abstract

Dynamic early exiting has been proven to improve the inference speed of the pre-trained language model like BERT. However, all samples must go through all consecutive layers before early exiting and more complex samples usually go through more layers, which still exists redundant computation. In this paper, we propose a novel dynamic early exiting combined with layer skipping for BERT inference named SmartBERT, which adds a skipping gate and an exiting operator into each layer of BERT. SmartBERT can adaptively skip some layers and adaptively choose whether to exit. Besides, we propose cross-layer contrastive learning and combine it into our training phases to boost the intermediate layers and classifiers which would be beneficial for early exiting. To keep the consistent usage of skipping gates between training and inference phases, we propose a hard weight mechanism during training phase. We conduct experiments on eight classification datasets of the GLUE benchmark. Experimental results show that SmartBERT achieves 2-3 \times computation reduction with minimal accuracy drops compared with BERT and our method outperforms previous methods in both efficiency and accuracy. Moreover, in some complex datasets like RTE and WNLI, we prove that the early exiting based on entropy hardly works, and the skipping mechanism is essential for reducing computation. Our codes are available at: <https://github.com/HuBoren99/SmartBert>.

1 Introduction

In recent years, large-scale pre-trained language models (PLMs) such as BERT [Devlin *et al.*, 2019], GPT [Radford *et al.*, 2018], ALBERT [Lan *et al.*, 2019], and XLNET [Yang *et al.*, 2019] RoBERTa [Liu *et al.*, 2019] have made significant progress in the field of natural language processing (NLP). However, these models usually require large computation resources. It is difficult to deploy these models in the case of limited resources. To reduce redundant computation, various approaches have been proposed, including network pruning [Li *et al.*, 2016; He *et al.*, 2017], weight

quantization [Jacob *et al.*, 2018], knowledge distillation [Hinton *et al.*, 2015; Sanh *et al.*, 2019; Jiao *et al.*, 2019] and dynamic early exiting [Xin *et al.*, 2020; Liu *et al.*, 2020; Schwartz *et al.*, 2020].

In this work, we mainly focus on dynamic early exiting methods. Such methods do not change the original network structure, and they only add some light plugins in the original network to decide whether early exit or not in each layer which can reduce plenty of computation while keeping comparable performance.

Although current dynamic early exiting technology has shown excellent characteristics, samples must go through all consecutive layers before early exiting, and complex samples have to go through nearly all layers. That is, complex samples have more redundant computation. In the field of computer vision, Highway Network [Srivastava *et al.*, 2015] and SkipNet [Wang *et al.*, 2018b] have proved that redundant computations exist in neural networks and some blocks (layers) can be skipped directly. Motivated by these works, we assume that some layers also can be skipped directly before early exiting for BERT to further reduce computation.

In this paper, firstly, we introduce a novel dynamic early exiting combined with layer skipping technique for BERT inference named SmartBERT, which plugs a skipping gate and an exiting operator into each layer of BERT and performs adaptive inference based on the principle of the higher priority of skipping than exiting. We use the widely used PLM BERT [Devlin *et al.*, 2019] as the backbone, and our method can be extended to other PLMs.

Secondly, in order to address the inconsistent usage of skipping gates between the train and inference stage, we design a hard weight mechanism. For adopting the transformation from soft to hard, we proffer a special training way. Thirdly, we propose cross-layer contrastive learning and combine it into our training phases to boost the intermediate layers and classifiers to achieve more efficient computation. Moreover, We conducted experiments on eight classification datasets of the GLUE benchmark [Wang *et al.*, 2018a], which shows that SmartBERT achieves 2-3 \times computation reduction with minimal accuracy drops compared with BERT and has a better performance than other dynamic early exiting models at the same computational cost. We also prove skipping mechanism is highly effective in some complex datasets like RTE and WNLI.

Our main contributions can be summarized as follows:

- We propose a novel dynamic early exiting method combined with layer skipping for BERT inference. And we design a hard weight mechanism and a special training way for consistent usage and better training respectively.
- We propose cross-layer contrastive learning into the training phase to boost the intermediate layers and classifiers, which is beneficial for early exiting.
- Our method can achieve 2-3x computation reduction with comparable performance and outperforms previous methods in both efficiency and accuracy. In some complex datasets like RTE and WNLI, our method can save the quantity of computation time (2-4x) compared to current early existing methods while achieving better performance.

2 Related Work

2.1 Model Compression

Many pre-trained language models (PLMs) have emerged in the last few years. Although these PLMs have achieved high scores in many NLP tasks, their inference time is slow, and the cost of calculation is expensive. One of the most representative models is BERT [Devlin *et al.*, 2019], which has made remarkable improvements in many NLP tasks. However, the inference speed of BERT is criticized. Therefore, a series of methods of model compression have been proposed to solve the above problems. Knowledge distillation [Hinton *et al.*, 2015] aims to transfer knowledge from the teacher model to the student model, which is applied in DistillBERT [Sanh *et al.*, 2019] and TinyBERT [Jiao *et al.*, 2019]. ALBERT [Lan *et al.*, 2019] uses sharing parameters to greatly reduce the number of parameters and memory consumption. Q8BERT [Zafri *et al.*, 2019] use symmetric linear quantization [Jacob *et al.*, 2018] to reduce the number of bits about the parameters of BERT. For Pruning, [Gordon *et al.*, 2020] mainly remove the unimportant part based on gradients of weights. Although these methods can improve the inference time of BERT, they can not adaptively change the architecture of the model according to the complexity of each sample. For example, simple samples must go through all layers for these methods. However, these samples may only go through early layers in early exiting models.

2.2 Adaptive Inference

Adaptive inference can adaptively change the architecture of the model according to the complexity of samples. DeeBERT [Xin *et al.*, 2020] chooses whether to early exit according to the entropy of the output distribution on each layer. FastBERT divides the early exiting classifiers into student classifiers and the teacher classifier, and then trains all student classifiers through self-distillation. FastBERT also uses the early exit based on entropy. FastBERT and DeeBERT can adaptively adjust the model size according to different samples. However, samples need to go through all layers before early exiting, and if they are complex enough, they need to go through the whole model. Hence, we introduce cross-layer contrastive learning to obtain more powerful classifiers and use layer skipping mechanism to reduce redundancy further.

2.3 Contrastive Learning

Recently, contrastive learning (CL) has made significant progress in various domains. Especially, unsupervised CL can exploit a bulk of unlabelled data to train a model with generalization which can even surpass the model trained under supervised training in some situations. The aim of CL is to maximize the agreement between positive views which are jointly sampled and disparting negative views in the representation space. Pioneering works were mainly proposed in the CV domain [He *et al.*, 2020; Chen *et al.*, 2020]. Then in the graph domain, there are plentiful follow-up works [You *et al.*, 2020; Zhu *et al.*, 2022] which borrow ideas from pioneering works. However, in the field of natural language processing, there are few inventive works [Gao *et al.*, 2021] that have shown up in recent years. The main difficulty in NLP domain is that augmentation is hard to design. Besides, previous works mainly focus on the quality of the final representations which may neglect the intermediate layers. But in our work, we hope each layer can be competent for good classification to early exit. So, we propose *cross-layer* contrastive learning in our method for both the first stage and the second training stage. In the *cross-layer* contrastive learning, we do not rely on data augmentation and assume that each token should have similar semantics across the consecutive layers because their representations do not change drastically. The same tokens will be treated as positive pairs in the consecutive layers, otherwise they shall be negative pairs.

3 Methodology

In this section, we first introduce the architecture of SmartBERT, which combines the early exiting mechanism and layer skipping mechanism. Secondly, we will illustrate how to combine cross-layer contrastive learning into the training phase to obtain powerful classifiers. Lastly, we introduce a hard weight mechanism and a special training way for skipping gates, further keeping the consistent usage of skipping gates between the training and inference phases.

3.1 Preliminaries

Given an input sentence $\mathbf{S} = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N\} \in \mathbb{R}^{N \times W}$, N and L stand for the number of tokens and encoder layers respectively. $\mathbf{X}^i \in \mathbb{R}^{N \times D}$ represents the output of the i^{th} encoder layer $\mathcal{E}^i(\cdot)$. \mathbf{X}^0 stands for the output of embedding layer. $\mathcal{G}^i(\cdot)$ and $\mathcal{C}^i(\cdot)$ represent the skipping gate and classifier, and the superscript means they are inserted in i^{th} layer.

3.2 Model Design

Different to [Liu *et al.*, 2020; Xin *et al.*, 2020] whose each layer will introduce one exiting classifier, we additionally plug a skipping gate into each layer. In order to understand SmartBERT better, we will first introduce the skipping gates and then introduce exiting classifiers.

The skipping gate aims to decide whether to execute or bypass the current layer like Figure 1.

With the skipping gate $\mathcal{G}^i(\cdot)$, the output of i^{th} encoder block can be defined as

$$\mathbf{X}^i = \mathcal{G}^i(\mathbf{X}^{i-1})\mathbf{X}^{i-1} + (1 - \mathcal{G}^i(\mathbf{X}^{i-1}))\mathcal{E}^i(\mathbf{X}^{i-1}), \quad (1)$$

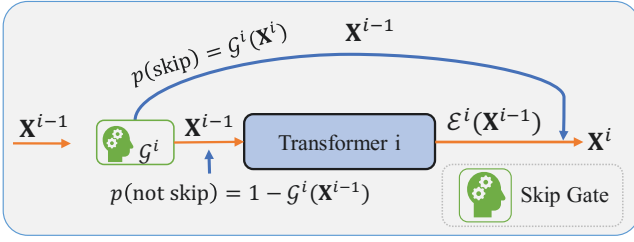


Figure 1: The skipping gate has $p(\text{skip})$ probability to skip the current transformer layer and has $p(\text{not skip})$ probability not to skip. In the inference phase, when $p(\text{skip})$ is greater than $p(\text{not skip})$, the gate will choose to skip, otherwise not to skip.

where $G^i(\mathbf{X}^{i-1}) \subseteq [0, 1]$ is the output of the i^{th} skipping gate which represents the probability of skipping the current layer. In the training phase, $G^i(\mathbf{X}^{i-1})$ is a continuous value in the range 0 to 1. However, in the inference phase, $G^i(\mathbf{X}^{i-1})$ is a discrete value that is 0 or 1, and the inconsistent usage of skipping gates between the training and inference phase will be discussed in Section 3.5.

Next, we will introduce early exiting classifiers. All early exiting classifiers use the early exiting mechanism based on the entropy, and we define the output of classifiers as

$$\mathbf{z}^i = C^i(\mathbf{X}^i), \quad (2)$$

where $\mathbf{z}^i \in \mathbb{R}^{1 \times C}$ represents the probability of each category.

3.3 Training

SmartBERT is composed of pre-trained BERT, skipping gates, and early exiting classifiers. For stability, we use two stages training strategy for the model, and different stages will train different parts of the model. In this section, we will illustrate two stages separately.

First Training Stage

In the first stage, we train skipping gates and fine-tune pre-trained BERT in a joint way. Other components will be fixed except the last classifier, as shown in Figure 2a. The total loss of the first stage consists of three main parts: classification error, the sloth of skipping, and cross-layer contrastive loss. It can be formalized as

$$\mathcal{L}_{\text{first}} = \mathcal{L}_{\text{CE}}(\mathbf{z}^L, \mathbf{y}) + \frac{\lambda}{\sum_{i=1}^n G^i(\mathbf{X}^i)} + \eta \mathcal{L}_{\text{contra.1}}, \quad (3)$$

where the first term represents the classification error (e.g., cross-entropy loss) like Equation 2. The second term indicates the sloth of skipping. In order to encourage the gates to skip to some extent rather than only considering the performance, we introduce the second term as a regular term. To enlarge the output of skipping gates. And $\lambda, \eta \in [0, 1]$ are scaling factors that are used to balance the influence of the regular terms on the loss. The third term represents a *cross-layer* contrastive loss, which mainly aims to obtain a more powerful BERT in this stage.

In cross-layer contrastive learning (CCL), we assume that each token should have similar semantics across the consec-

utive layers because their representations do not change drastically. The experiments in Appendix C can support our motivation. The same tokens will be considered as positive pairs across consecutive layers. Otherwise, they shall be deemed to be negative pairs. The CCL loss in the first training stage can be formulated as

$$\mathcal{L}_{\text{contra.1}} = \frac{1}{N} \frac{1}{L-1} \sum_{i=1}^{L-1} \sum_{m=1}^N -\log \frac{e^{s(\mathbf{p}_m^i, \mathbf{p}_m^{i+1})/\tau}}{\sum_{k=1}^N e^{s(\mathbf{p}_m^i, \mathbf{p}_k^{i+1})/\tau}}, \quad (4)$$

where $\mathbf{p}_m^i = f(\mathbf{x}_m^i)$, \mathbf{x}_m^i represents the m^{th} word representation in the i^{th} layer. $f(\cdot)$ is a projection function that maps representations to another latent space where the contrastive loss is calculated. $s(x, y)$ represents a score function (e.g., cosine similarity), and τ is temperature parameter. In this training stage, we only boost the transformer layer by the cross-layer contrastive loss.

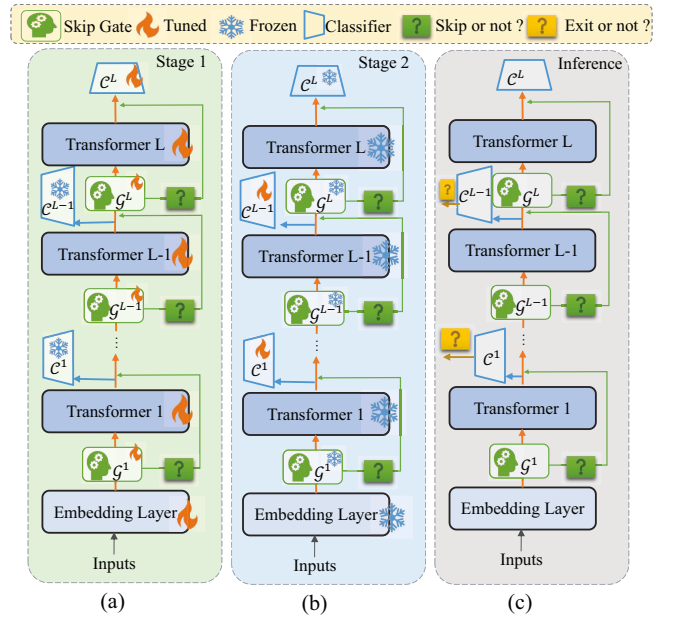


Figure 2: The processes of two-stage training stages and inference phase. (a) and (b) show the first and second stages of training, respectively. Note that different components will be trained in different stages. (c) shows the fast inference process. In each layer, the skipping will decide whether to skip the current layer. If it chooses to skip, then the input directly enters the next gate. Otherwise, the input will be fed into the current layer, and the corresponding classifier will decide whether to exit.

Second Training Stage

In the second stage, the trained parameters of the first stage will be frozen, and all classifiers except the last one will be trained as shown in Figure 2b. In order to obtain more powerful classifiers, we also adopt CCL into the second training phase. Specifically, the token representations obtained by frozen transformer layers will be fed into classifiers firstly. Secondly, we extract the hidden states that only go through the first layer (*i.e.*, self-attention layer) of classifiers, and

these hidden states will be fed into the projection head to map them into the space where contrastive loss is occupied. Lastly, the contrastive loss (*i.e.*, InfoNCE loss) is employed. Figure 3 shows the overall pipeline. The CCL loss in the second training stage can be formalized as

$$\ell(\hat{\mathbf{h}}_m^i, \hat{\mathbf{h}}_m^{i+1}) = -\log \frac{e^{s(\hat{\mathbf{h}}_m^i, \hat{\mathbf{h}}_m^{i+1})/\tau}}{\sum_{k=1}^N e^{s(\hat{\mathbf{h}}_m^i, \hat{\mathbf{h}}_k^{i+1})/\tau}}, \quad (5)$$

where $\hat{\mathbf{h}}_m^i = f(\mathbf{h}_m^i)$, $\mathbf{h}_m^i = \text{Self-Attention}(\mathbf{X}^i)[m]$, $\text{Self-Attention}(\cdot)$ represents a self-attention function and $[m]$ means we only select m -th token. Other notations have the same meanings as in Equation 4. Enumerating all tokens in all layers, the overall loss is:

$$\mathcal{L}_{\text{contra.2}} = \frac{1}{N} \frac{1}{L-1} \sum_{i=1}^{L-1} \sum_{m=1}^N \ell(\hat{\mathbf{h}}_m^i, \hat{\mathbf{h}}_m^{i+1}). \quad (6)$$

In this way, the exiting classifiers can help each other in the consecutive layers, which will assist our model early exiting. In detail, our classifier is similar to [Liu *et al.*, 2020], which consists of one attention layer and one linear layer. Each token in the i -th layer will acquire its high-level representation \mathbf{H}^i through the attention layer. Through Equation 5, the self-attention layer will be boosted so that we can empower the classifier to produce more confident results which are helpful for early exiting based on entropy.

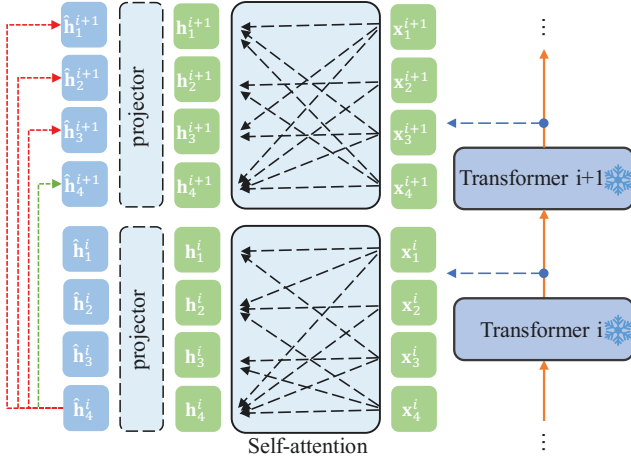


Figure 3: The pipeline of cross-layer contrastive learning in the second stage. The green line (*i.e.*, the same token in the consecutive layers) indicates positive pairs, and red lines (*i.e.*, otherwise) represent negative pairs.

For keeping consistent with the inference phase, SmartBERT uses the output of skipping gates to update the hidden states of the current layer like Equation 1. Next, we define the loss of each classifier as

$$\mathcal{L}_{\text{CE}}^i = \mathcal{L}_{\text{CE}}(\mathbf{z}^i, \mathbf{y}). \quad (7)$$

In the second stage, the loss function is the sum of the cross-entropy loss of classifiers with cross-layer contrastive loss:

$$\mathcal{L}_{\text{second}} = \sum_{i=1}^{L-1} \mathcal{L}_{\text{CE}}^i + \eta \mathcal{L}_{\text{contra.2}}. \quad (8)$$

3.4 Inference

After finishing the two-stage training, SmartBERT can take advantage of skipping gates and early exiting to accelerate inference considering the complexity of samples. Only when the output of skipping gates is greater than or equal to 0.5, the model will decide to skip. The specific formula follows:

$$\mathbf{X}^i = \begin{cases} \mathbf{X}^{i-1} & , \text{ if } \mathcal{G}^i(\mathbf{X}^{i-1}) \geq 0.5 \\ \mathcal{E}^i(\mathbf{X}^{i-1}) & , \text{ if } \mathcal{G}^i(\mathbf{X}^{i-1}) < 0.5 \end{cases} \quad (9)$$

When skipping gates choose not to skip, the current early exiting classifier will be used, and all classifiers use an early exiting mechanism based on the entropy:

$$\begin{cases} \text{Exit} & , \text{ if Entropy}(\mathbf{z}^i) < S \\ \text{Continue} & , \text{ otherwise} \end{cases} \quad (10)$$

where \mathbf{z}^i represents a categorical distribution of the i^{th} classifier, and S is the entropy threshold for early exiting, which is set manually. The entire inference algorithm is in Appendix A.

3.5 Hard Weight Mechanism

In the previous sections, we have introduced SmartBERT training and inference. However, we find the inconsistent usage of skipping gates between the training and inference phase. Specifically, in the training phase, the output of skipping gates is a continuous probability value between 0 and 1 like Equation 1, and in the inference phase, we will discretize the output of skipping gates like Equation 9 for acceleration. To reduce the impact of inconsistent usage of skipping gates, we introduce a hard-weight mechanism for skipping gates in the training phase. And the hard weight mechanism is differentiable, we use the trick:

$$G_{\text{hard}} = \mathbb{1}_{G_{\text{soft}} \geq 0.5} - \tilde{G}_{\text{soft}} + G_{\text{soft}}, \quad (11)$$

where $\mathbb{1}_{G_{\text{soft}} \geq 0.5}$ is an indicator function which returns 1 if $G_{\text{soft}} \geq 0.5$ otherwise 0, and \tilde{G}_{soft} is the result of truncating the gradient of G_{soft} . With this trick, we achieve two things: (1) makes the output value exactly one-hot (since we add then subtract soft value), (2) makes the gradient equal to the soft gradient (since we strip all other gradients). In the training phase, we define the output of skipping gates as the soft weight G_{soft} . To keep the consistent usage of skipping gates between training and inference phases, we turn the soft weight G_{soft} into discrete values (0 or 1), and keep the gradient equal to Equation 12. In this way, the output of the skipping gate becomes a discrete skipping decision in the training phase, which is consistent with the inference phase:

$$G_{\text{hard}} = \begin{cases} 1 & , \text{ fp, } G_{\text{soft}} \geq 0.5 \\ 0 & , \text{ fp, } G_{\text{soft}} < 0.5 \\ G_{\text{soft}} & , \text{ backward pass} \end{cases} \quad (12)$$

where fp represents forward pass, G_{soft} is the soft weight and G_{hard} is the hard weight.

3.6 A Special Training Way for Skipping Gates

Although the hard weight mechanism can reduce the impact of inconsistent usage of skipping gates, directly using the hard weight mechanism will lead to some layers can not be well fine-tuned because they may be skipped. However, these layers may be used in the inference phase, which will result in unsatisfactory performance. To ensure all layers can be trained and keep the consistent usage of skipping gates in the training and inference phase, we introduce a special training way from soft weight mechanism to hard weight mechanism for skipping gates.

We use the transition from the soft weight mechanism to the hard weight mechanism to address the above problems. First, we use the soft weight mechanism to warm up all skipping gates and transformer layers and then use the hard weight mechanism to train them. Moreover, the soft weight mechanism aims to fine-tune these layers, and the hard weight mechanism is used for keeping the consistent usage of skipping gates between training and inference phases.

4 Experiments

4.1 Baselines

We compare our method with three baselines:

- **BERT.** A large-scale pre-trained language model based on Transformer, which is used as the backbone of our methods. In experiments, we only use the BERT-base model pre-trained by Google. [Devlin *et al.*, 2019]
- **DistillBERT.** It is a smaller transformer-based model by distilling the BERT. [Sanh *et al.*, 2019]
- **Early exiting model.** The dynamic early exiting method for BERT is an effective method to accelerate BERT inference. We choose some classical methods like **DeeBERT** [Xin *et al.*, 2020] and **FastBERT** [Liu *et al.*, 2020] as our baselines, which use the early exiting mechanism based on entropy.

4.2 Datasets

To verify the effectiveness of our methods, We conduct experiments on eight classification datasets of the GLUE benchmark [Wang *et al.*, 2018a], including SST-2 [Socher *et al.*, 2013], CoLA [Warstadt *et al.*, 2019], MRPC [Dolan and Brockett, 2005], MNLI [Williams *et al.*, 2017], QQP, QNLI [Rajpurkar *et al.*, 2016], RTE [Bentivogli *et al.*, 2009], and WNLI [Levesque *et al.*, 2012].

4.3 Experimental Setup

The experiments are done on an NVIDIA 2080Ti GPU. We adopt the same parameters for BERT, DeeBERT, FastBERT, and SmartBERT. In experiment, these models use the pre-trained parameters (**bert-base-uncased**) released by the Hugging Face Transformer Library [Wolf *et al.*, 2019]. In the pre-trained parameters, the number of transformer layers, the dimension of hidden states, and the max length of the input sentence are set to 12, 768, and 128. We use AdamW [Kingma and Ba, 2014] to train these models with a default batch size of 32. For each task, we select the best fine-tuning learning rate (among 1e-5, 2e-5, 5e-5).

In the first stage of SmartBERT, we train the model with 5 epochs and select one with the best accuracy for the second stage. In the second stage, we adopt cross-layer contrastive learning to train the classifiers and train these classifiers for 4 epochs. We slightly tune the hyper-parameters across the different tasks.

Following prior work, our batch size of inference is set to 1 in the inference phase.

In inference, We followed FastBERT and used FLOPs as an indicator to evaluate the computational cost. Generally speaking, the size of the FLOPs reflects the inference speed of the model, and the smaller the FLOPs of the model are, the shorter the inference time will be. Table 2 presents the computational cost of each operation within the SmartBERT, which shows that the computational cost of the Skipping Gate and Classifier is much lower than the Transformer Layer.

4.4 Main Results

We evaluate these models in eight classification datasets of the GLUE benchmark and select different entropy thresholds to test the performance and sample-averaged FLOPs for DeeBERT, FastBERT and SmartBERT. We set three different entropy thresholds for early exiting models and compare the results with other baselines in Table 1. The results of the experiment show that our model achieves 2-3 \times computation reduction with minimal accuracy drops and even has a better performance than BERT. In the results, due to DistillBert cannot adaptively change the model architecture according to the sample complexity, it has a fixed computational complexity. In other words, the computational cost of DistillBert is independent of sample difficulty. In fact, the difficulty of each dataset is different, and in more datasets, our method has a better performance than DistillBert.

In Figure 4, we set several entropy thresholds and compare different models' tradeoffs in accuracy and computational cost on some GLUE datasets. From Figure 4, we conclude that:

1. Compared with FastBERT and DeeBERT, SmartBERT have a better performance at the same computational cost, which proves our model is more effective than other approaches.
2. In all datasets, the scores (accuracy/F1 score/FLOPs) of these models start to decline when the computational costs reach certain values which verify the redundant computation in original BERT, and SmartBERT dominates the performance under different computation costs comparing with FastBERT and DeeBERT.

4.5 Effectiveness of Skipping Mechanism

The early exiting mechanism is based on the entropy of the corresponding classifier output. However, previous methods hardly reduce computation cost when samples are complex and difficult enough. Because the entropy of classifier output will be extremely high that leads to most exiting classifiers do not exit early. In Table 1, we can see that even if the threshold S is set as 0.5, the DeeBERT and FastBERT still have an extremely high computational cost for RTE and WNLI dataset,

Dataset/ Model	MRPC		SST-2		QNLI		QQP		MNLI		MNLI-mm		CoLA		RTE		WNLI	
	F1	FLOPs (cost)	Acc	FLOPs (cost)	Acc	FLOPs (cost)	F1	FLOPs (cost)	Acc	FLOPs (cost)	Acc	FLOPs (cost)	Mcc	FLOPs (cost)	Acc	FLOPs (cost)	Acc	FLOPs (cost)
BERT	88.5	21744M (100%)	93.2	21744M (100%)	88.1	21744M (100%)	84.4	21744M (100%)	84.6	21744M (100%)	84.8	21744M (100%)	50.8	21744M (100%)	61.0	21744M (100%)	56.3	21744M (100%)
DistillBERT	87.0	10872M (50%)	91.5	10872M (50%)	86.8	10872M (50%)	82.4	10872M (50%)	81.9	10872M (50%)	81.8	10872M (50%)	47.5	10872M (50%)	59.2	10872M (50%)	55.3	10872M (50%)
DeeBERT (S=0.1)	88.4	21744M (100%)	92.4	13900M (63%)	87.5	20008M (92%)	84.2	17578M (80%)	84.2	20796M (95%)	84.3	20786M (95%)	49.9	21840M (100%)	59.9	21758M (100%)	56.3	21758M (100%)
DeeBERT (S=0.3)	88.2	19726M (90%)	90.3	10874M (50%)	87.1	15770M (72%)	83.1	14750M (67%)	84.2	19698M (90%)	84.1	18148M (83%)	49.9	16850M (77%)	59.9	21718M (99%)	56.3	21758M (100%)
DeeBERT (S=0.5)	87.2	13540M (62%)	85.2	7468M (34%)	84.9	10502M (48%)	74.5	8744M (40%)	83.9	18728M (86%)	83.1	16362M (75%)	49.9	15130M (70%)	59.9	21332M (98%)	56.3	21758M (100%)
FastBERT (S=0.1)	88.5	22196M (102%)	92.4	11094M (51%)	87.5	19684M (90%)	83.9	15808M (72%)	83.4	19980M (91%)	84.0	19116M (87%)	49.8	20846M (95%)	59.4	22196M (102%)	56.3	22196M (102%)
FastBERT (S=0.3)	88.5	20278M (93%)	90.7	6704M (30%)	86.8	13096M (60%)	82.9	10586M (48%)	83.1	15468M (71%)	83.6	15262M (70%)	49.5	15262M (70%)	59.4	22196M (102%)	56.3	22196M (102%)
FastBERT (S=0.5)	88.3	16478M (75%)	86.8	4076M (18%)	84.5	7964M (36%)	78.2	6064M (27%)	81.3	11958M (54%)	81.3	11020M (50%)	44.4	12700M (58%)	59.4	22196M (102%)	56.3	22196M (102%)
Ours (S=0.1)	89.8	21390M (98%)	93.1	9040M (41%)	87.9	16512M (75%)	84.1	13128M (60%)	84.4	18436M (84%)	84.6	18196M (83%)	51.5	20568M (94%)	60.6	12684M (58%)	56.3	6768M (31%)
Ours (S=0.3)	89.6	17884M (82%)	91.9	5706M (26%)	87.0	10878M (50%)	82.5	8610M (39%)	83.6	14442M (66%)	83.9	13936M (64%)	51.3	15072M (69%)	60.6	12684M (58%)	56.3	6768M (31%)
Ours (S=0.5)	88.9	14348M (65%)	87.7	3608M (16%)	84.9	7020M (32%)	78.0	5130M (23%)	82.4	10450M (49%)	82.0	10570M (48%)	49.9	10468M (49%)	60.6	12684M (58%)	56.3	6768M (31%)

Table 1: Comparison between baselines(BERT,DistillBERT,FastBERT,DeeBERT) and SmartBERT on the GLUE benchmark. *FLOPs* are multiply-accumulate operations which represent computational complexity. *S* represents the entropy threshold, and *cost* is the computational cost.

which means early exiting mechanism is invalid on such situations. But SmartBERT achieves a relatively large computation reduction with minimal accuracy drop, which shows the skipping mechanism is effective.

To further prove the effectiveness of the layer skipping mechanism, we disable the early exiting mechanism and only use the layer skipping mechanism during the inference phase. The empirical results are shown in Table 3, which proves that the skipping mechanism is able to reduce the computational cost compared with BERT. Moreover, combined Table 1 and Table 3, FLOPs and Accuracy are hardly changed on the RTE and WNLI datasets, which further proves that the early exiting mechanism is invalid in difficult datasets and layer skipping mechanism plays an important role in our method on such situations.

4.6 Effectiveness of Cross-Layer Contrastive Learning

We have introduced the cross-layer contrastive learning, which takes the same tokens across consecutive layers as positive pairs and the different tokens across consecutive layers as negative pairs. To further evaluate the effectiveness of contrastive learning, we compare the two cases of using contrastive learning and not using contrastive learning. Em-

Operation	FLOPs
Each Transformer Layer	1811.8M
Each Classifier	37.6M
Each Skipping Gate	37.4M

Table 2: This experiment evaluates each operation in the SmartBERT, using FLOPs as the indicator.

Dataset/ Model	SST-2		RTE		WNLI	
	Acc	FLOPs (cost)	Acc	FLOPs (cost)	Acc	FLOPs (cost)
BERT	93.2	21744 (100%)	61.0	21744 (100%)	56.3	21744 (100%)
Skip	93.1	14500 (66%)	60.6	12190 (56%)	56.3	6360 (29%)

Table 3: *Skip* represents our methods only use the skipping mechanism and disable the early exiting mechanism. This experiment is conducted on the SST-2, RTE and WNLI datasets.

pirical results are shown in Figure 5, which shows cross-layer contrastive learning is effective for further reducing the computation.

4.7 Effectiveness of Our Special Training Way

In the previous section, we introduce the soft weight mechanism and hard weight mechanism. Meanwhile, we propose a special training way from soft weight mechanism to hard weight mechanism. In the experiment, we compare three different training strategies: (1) soft weight mechanism: using the continuous probability value of skipping gates for training but using the discrete value of skipping gates for inference, (2) hard weight mechanism: using the discrete value of skipping gates for training and inference and (3) the special training way: firstly using continuous probability value for training, and then using the discrete value for training.

We compare these methods on the SST-2 and QNLI datasets, and the results are shown in Figure 6. The results prove that the special training way is effective for balancing the inconsistent usage of skipping gates between training and

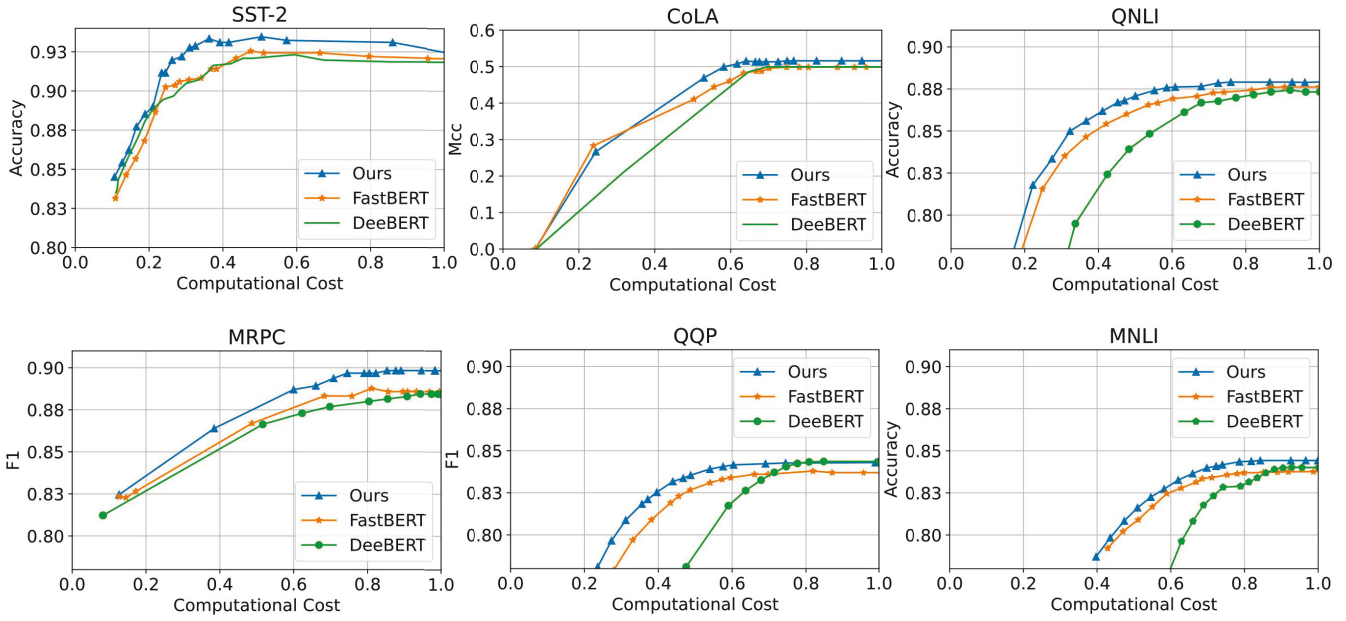


Figure 4: Comparison between baselines and SmartBERT on SST-2, QNLI, QQP, MNLI, CoLA, MRPC, showing the influence of computational cost on performance(Accuracy/F1 score/MCC).

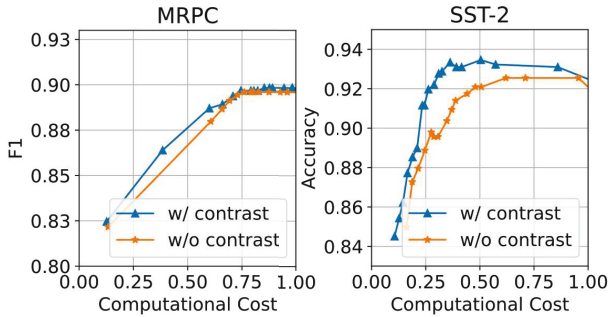


Figure 5: *contrast* represents using the cross-layer contrastive learning, *w/o contrast* represents not using the cross-layer contrastive learning.

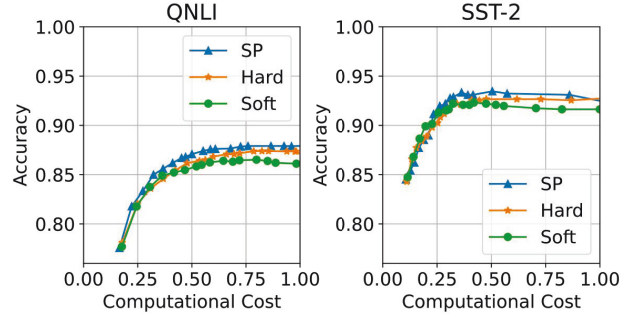


Figure 6: *Hard* represents the hard weight mechanism, *Soft* represents the soft weight mechanism, and *SP* uses the special training way from soft mechanism to hard mechanism.

inference phases.

5 Conclusion

In this paper, we propose a novel dynamic early exiting combined with layer skipping for BERT inference. To address the inconsistent usage of skipping gates in the inference and training phase, we propose hard weight mechanism and a special training way. In addition, we propose cross-layer contrastive learning and combine it into our training phases to boost the intermediate layers and classifiers which can further reduce computation cost. We evaluated our model on eight classification datasets of the GLUE benchmark. Empirical results show that SmartBERT could achieve performance comparable to BERT while significantly reducing computational cost. Compared to other dynamic early exiting models, SmartBERT obtain better accuracy with lower computation. Moreover, we conduct a series of ablation studies to demon-

strate that each component is beneficial. In the future, we will adopt our methods to other pre-trained language models. We also will further study the combination of other acceleration methods and layer skipping mechanism.

Acknowledgements

We would like to thank anonymous reviewers for their valuable comments and suggestions. This work has been supported in part by the NSFC (No. 62272411), the Zhejiang NSF (LR21F020004), Ant Group and Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies.

Contribution Statement

Boren Hu, Yun Zhu designed the research and wrote the paper. **Siliang Tang** was an early advocate for the research, and advised **Boren Hu** and **Yun Zhu**. **Jiacheng Li** provided opinions for the writing of the paper. Moreover, **Boren Hu**

and **Yun Zhu** have equal contributions for this paper, and **Siliang Tang** is the corresponding author.

References

- [Bentivogli *et al.*, 2009] Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. The fifth pascal recognizing textual entailment challenge. In *TAC*, 2009.
- [Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proc. of ICML*, 2020.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL*, 2019.
- [Dolan and Brockett, 2005] Bill Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Third International Workshop on Paraphrasing (IWP2005)*, 2005.
- [Gao *et al.*, 2021] Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence embeddings. In *Proc. of EMNLP*, 2021.
- [Gordon *et al.*, 2020] Mitchell A. Gordon, Kevin Duh, and Nicholas Andrews. Compressing BERT: studying the effects of weight pruning on transfer learning. *CoRR*, 2020.
- [He *et al.*, 2017] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proc. of ICCV*, 2017.
- [He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proc. of CVPR*, 2020.
- [Hinton *et al.*, 2015] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [Jacob *et al.*, 2018] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proc. of CVPR*, 2018.
- [Jiao *et al.*, 2019] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Lan *et al.*, 2019] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [Levesque *et al.*, 2012] Hector Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*, 2012.
- [Li *et al.*, 2016] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.
- [Liu *et al.*, 2019] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [Liu *et al.*, 2020] Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Haotang Deng, and Qi Ju. Fastbert: a self-distilling bert with adaptive inference time. *arXiv preprint arXiv:2004.02178*, 2020.
- [Radford *et al.*, 2018] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. In *Advances in neural information processing systems*, pages 8735–8745, 2018.
- [Rajpurkar *et al.*, 2016] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- [Sanh *et al.*, 2019] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [Schwartz *et al.*, 2020] Roy Schwartz, Gabriel Stanovsky, Swabha Swayamdipta, Jesse Dodge, and Noah A Smith. The right tool for the job: Matching model and instance complexities. *arXiv preprint arXiv:2004.07453*, 2020.
- [Socher *et al.*, 2013] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. of EMNLP*, 2013.
- [Srivastava *et al.*, 2015] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
- [Wang *et al.*, 2018a] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [Wang *et al.*, 2018b] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *Proc. of ECCV*, 2018.
- [Warstadt *et al.*, 2019] Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 2019.

- [Williams *et al.*, 2017] Adina Williams, Nikita Nangia, and Samuel R Bowman. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*, 2017.
- [Wolf *et al.*, 2019] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
- [Xin *et al.*, 2020] Ji Xin, Raphael Tang, Jaejun Lee, Yao-liang Yu, and Jimmy Lin. Deebert: Dynamic early exiting for accelerating bert inference. *arXiv preprint arXiv:2004.12993*, 2020.
- [Yang *et al.*, 2019] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Proc. of NeurIPS*, 2019.
- [You *et al.*, 2020] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. *Proc. of NeurIPS*, 2020.
- [Zafir *et al.*, 2019] Ofir Zafir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. Q8bert: Quantized 8bit bert. In *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*, 2019.
- [Zhu *et al.*, 2022] Yun Zhu, Jianhao Guo, Fei Wu, and Siliang Tang. Rosa: A robust self-aligned framework for node-node graph contrastive learning. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 3795–3801. International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track.