

COSEE: Consistency-Oriented Signal-Based Early Exiting via Calibrated Sample Weighting Mechanism

Jianing He¹, Qi Zhang¹, Hongyun Zhang¹, Xuanjing Huang², Usman Naseem³, Duoqian Miao^{1*}

¹Tongji University, China

²Fudan University, China

³Macquarie University, Australia

{jnhe, zhangqi_cs, zhanghongyun, dqmiao}@tongji.edu.cn, xjhuang@fudan.edu.cn, engr.usmannaseem87@gmail.com

Abstract

Early exiting is an effective paradigm for improving the inference efficiency of pre-trained language models (PLMs) by dynamically adjusting the number of executed layers for each sample. However, in most existing works, easy and hard samples are treated equally by each classifier during training, which neglects the test-time early exiting behavior, leading to inconsistency between training and testing. Although some methods have tackled this issue under a fixed speed-up ratio, the challenge of flexibly adjusting the speed-up ratio while maintaining consistency between training and testing is still under-explored. To bridge the gap, we propose a novel Consistency-Oriented Signal-based Early Exiting (COSEE) framework, which leverages a calibrated sample weighting mechanism to enable each classifier to emphasize the samples that are more likely to exit at that classifier under various acceleration scenarios. Extensive experiments on the GLUE benchmark demonstrate the effectiveness of our COSEE across multiple exiting signals and backbones, yielding a better trade-off between performance and efficiency.

Code — <https://github.com/He-Jianing/COSEE>

Extended version — <https://arxiv.org/abs/2412.13236>

1 Introduction

Although tremendous improvements have been achieved by pre-trained language models (PLMs) in natural language processing tasks (Devlin et al. 2019; Lan et al. 2020; Radford et al. 2019; Liu et al. 2019), high computational costs of PLMs in both training and inference still hinder their deployment in resource-constrained devices and real-time scenarios. Besides, overthinking problem (Kaya, Hong, and Dumitras 2019) also restricts the application of PLMs. Precisely, for easy samples, PLMs can generate correct predictions according to the representations indicated by shallow layers. However, high-level representations may focus on more intricate or unrelated details, leading to incorrect answers.

To address these issues, early exiting (Xin et al. 2020; Zhou et al. 2020; Xin et al. 2021; Liao et al. 2021; Sun et al. 2022; Zeng et al. 2024), a kind of adaptive inference strategy, has been proposed to accelerate the inference of

PLMs. As illustrated in Figure 2, each intermediate layer of the PLM is coupled with an internal classifier to give an early prediction. This enables the early exiting of samples once the early predictions are sufficiently reliable, eliminating the need for passing them through the entire model. This method employs a sample-wise inference strategy to deal with easy samples with shallow classifiers and process hard samples with deeper classifiers, significantly improving inference efficiency without sacrificing accuracy and alleviating the overthinking problem.

Signal-based early exiting methods (Xin et al. 2020; Liu et al. 2020; Zhou et al. 2020; Schwartz et al. 2020; Li et al. 2021; Liao et al. 2021; Ji et al. 2023; Zhu 2021; Zhu et al. 2021; Gao et al. 2023; Zhang et al. 2023; He et al. 2024; Akbari, Banitalebi-Dehkordi, and Zhang 2022; Xin et al. 2021; Balagansky and Gavrilo 2022) are typical implementations of early exiting, which rely on carefully designed exiting signals (e.g. entropy, energy score, softmax score, and patience) to dynamically adjust the number of executed layers for each sample. The inference process is terminated once the exiting signal meets a certain condition. These methods can easily adapt to various acceleration requirements during inference by simply adjusting the threshold, without incurring additional training costs. However, existing works simply use the (weighted) sum of cross-entropy losses from all classifiers as the training objective, where each classifier treats the loss of both easy and hard samples equally. This treatment ignores the dynamic early exiting behavior during inference (as shown in Figure 1), leading to a gap between training and testing.

To bridge the gap, router-based early exiting methods (Sun et al. 2022; Mangrulkar, MS, and Sembium 2022; Zeng et al. 2024) have been successively proposed. These methods employ a router (e.g. a hash function or a network) to determine the exiting layer of samples during both training and inference, and each sample only incurs a cross-entropy loss at its exiting classifier, ensuring consistency between training and testing. However, router-based early exiting methods fail to meet various acceleration requirements during inference, as a router can only generate a fixed exiting strategy, leading to unadjustable speed-up ratios.

In this paper, we aim to bridge the gap between training and testing while enabling flexible adjustments of the speed-up ratio. To this end, building upon the signal-based

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

early exiting framework, we propose to assign sample-wise weights on the cross-entropy loss of all classifiers, such that each classifier is encouraged to emphasize samples that are more likely to exit at that classifier. Unfortunately, samples exit at different classifiers under various acceleration scenarios, bringing extreme challenges to weight assignment.

To address the challenges, we propose a novel framework of Consistency-Oriented Signal-based Early Exiting (COSEE). Specifically, at each training step, we mimic the test-time early exiting process at multiple randomly selected thresholds to find where the samples tend to exit under different accelerations. Subsequently, we adopt a heuristic sample weighting mechanism (SWM) to assign weights on the cross-entropy loss of each sample across all classifiers, where each sample is emphasized by the classifiers near its exiting layer. Accordingly, we minimize the mean of cross-entropy losses across different thresholds to ensure the model’s generalization ability in various acceleration scenarios. In addition, we further devise an online signal calibration (OSC) objective to generate highly discriminative exiting signals for more reliable exiting decisions, thus encouraging more proper loss weights based on exiting layers.

Our method is simple yet effective. Extensive experiments on the GLUE benchmark demonstrate that our COSEE framework with energy score consistently outperforms the state-of-the-art methods across all tasks, yielding a better trade-off between performance and efficiency with faster convergence speed and negligible additional storage overhead. In addition, an in-depth analysis further confirms the generalization of the COSEE framework on different exiting signals and backbones. Our main contributions can be summarized as follows:

- We disclose that the performance bottleneck of current early exiting methods primarily stems from the challenge of ensuring consistency between training and testing while flexibly adjusting the speed-up ratios.
- We propose a novel Consistency-Oriented Signal-based Early Exiting (COSEE) framework to bridge the gap, which incorporates a sample weighting mechanism (SWM) and an online signal calibration (OSC) objective.
- Extensive experiments verify the effectiveness of our COSEE across multiple exiting signals and backbones.

2 Preliminaries

In this section, we provide the necessary background for signal-based early exiting. Related works are detailed in the extended version.

2.1 Problem Definition

Per Figure 2, given a BERT-style PLM with M layers, we denote the hidden states at the m th layer as $h^{(m)}$. To enable early exiting during inference on a classification task involving C classes, each intermediate layer is equipped with an internal classifier $F_m, m \in \{1, 2, \dots, M-1\}$ to produce an early prediction $p^{(m)} = F_m(h^{(m)})$, i.e., a probability distribution over the C classes. Classifiers in different layers do not share parameters.




Figure 1: Exiting layer distribution on the QNLI development set with entropy-based exiting signal (Threshold = 0.4). Neg and Pos denote negative and positive samples, respectively. Samples near the classification boundary (hard samples) tend to exit at deep classifiers, while samples far from the classification boundary (easy samples) typically exit at shallow classifiers.

2.2 Signal-based Early Exiting

For a given sample x , the inference process is terminated once the exiting signal at the current layer meets a certain condition. For exiting signals that exhibit a positive correlation with sample difficulty (e.g. entropy and energy score), early exiting is triggered once the exiting signal falls below a predefined threshold. A higher threshold leads to a higher speed-up ratio and potentially some performance degradation. Conversely, for exiting signals negatively correlated with sample difficulty (e.g. patience and softmax score), the exiting condition is met when the exiting signal surpasses the threshold. A higher threshold leads to a lower speed-up ratio and performance improvements.

2.3 Conventional Training Methods

In current signal-based early exiting methods, a widely used training objective involves the (weighted) sum of cross-entropy losses across all classifiers:

$$L = \sum_{m=1}^M w_m L^{(m)}, \quad (1)$$

where $L^{(m)}$ denotes the cross-entropy loss of the m th classifier and w_m denotes the corresponding loss weight. Under Eq.(1), each classifier treats the loss of both easy and hard samples equally, which is inconsistent with the dynamic early exiting behavior during inference.

3 The COSEE Framework

3.1 Framework Overview

We propose a novel Consistency-Oriented Signal-based Early Exiting (COSEE) framework for PLMs, aiming to ensure consistency between training and testing while maintaining flexible adjustments of the speed-up ratio. Figure 2 provides an overview of our framework. We first propose a sample weighting mechanism (SWM) that identifies the potential exiting layer of samples by simulating the test-time

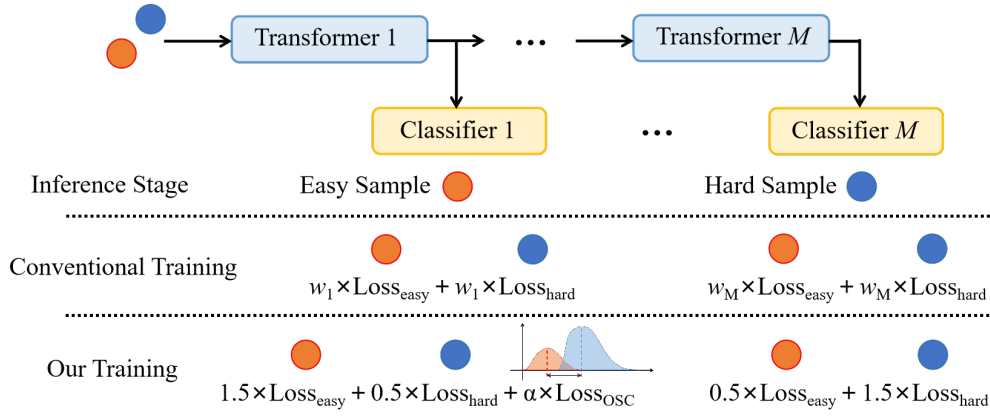


Figure 2: Comparison between the conventional signal-based early exiting framework and our COSEE. The conventional framework simply minimizes the (weighted) sum of cross-entropy losses from all classifiers, where each classifier treats all samples equally during training. Instead, our COSEE enables each classifier to emphasize samples that are more likely to exit at that classifier, ensuring consistency between training and testing. We also incorporate an online signal calibration objective Loss_{OSC} for each internal classifier to encourage highly discriminative exiting signals for more reliable exiting decisions and loss weights.

early exiting process during training and then uses this information to produce sample-wise loss weights across all classifiers. Additionally, we further devise an online signal calibration (OSC) objective to encourage highly discriminative exiting signals for more reliable exiting decisions, thus ensuring more proper loss weights based on exiting layers. Finally, regarding the exiting signal, we introduce a normalized energy score to align energy distributions across different layers for easy threshold selection. We primarily use it to implement the COSEE framework.

3.2 Sample Weighting Mechanism

Our goal is to identify the potential exiting layer of samples in various acceleration scenarios, and then assign greater weights to the cross-entropy loss of each sample on classifiers closer to its exiting layer. Accordingly, at each training step, all samples are passed through the entire model to generate predictions and exiting signals at all classifiers. Subsequently, we randomly select K thresholds and simulate the early exiting process based on exiting signals at each threshold to find where the samples exit. This information is used to produce sample-wise loss weights across all classifiers.

Range for Threshold Selection. For threshold selection, we collect the maximum and minimum values of exiting signals across all layers for training samples within each epoch and use them to create the selection range for the next epoch. We start with the thresholds randomly selected between 0 and 1 in the first epoch.

Weight Assignment. For a given threshold τ , we impose sample-wise loss weights across all classifiers based on the exiting layer of samples and then compute the classification loss at threshold τ :

$$L_{\text{CE},\tau} = \frac{1}{N} \sum_{n=1}^N \sum_{m=1}^M w_n^{(m)} \cdot \text{CE}(\hat{y}_n^{(m)}, y_n), \quad (2)$$

$$w_n^{(m)} = \frac{e^{-\beta_t \cdot |m - m_n^*|}}{\sum_{m=1}^M e^{-\beta_t \cdot |m - m_n^*|}}, \quad (3)$$

where $\text{CE}(\hat{y}_n^{(m)}, y_n)$ and $w_n^{(m)}$ denote the cross-entropy loss and the loss weight for the n th sample at the m th classifier respectively, and $w_n^{(m)}$ satisfies $\sum_{m=1}^M w_n^{(m)} = 1$. N denotes the number of samples. m_n^* denotes the index of exiting layer for the n th sample at threshold τ , and β_t denotes the decay factor at the t th training step. According to Equation 3, classifiers closer to the exiting layer are assigned greater weights compared to those further away, i.e., each sample is emphasized by the classifiers near its exiting layer. Note that the loss weights of classifiers are symmetrical around the exiting layer for easy parameter selection. Different from router-based early exiting methods, which employ one-hot sample-wise loss weights such that each sample only incurs a cross-entropy loss on its exiting classifier, we employ a softer sample weighting mechanism to enable the generality of our COSEE on unseen thresholds.

During the early training stage, unstable exiting layers often lead to fluctuating loss weights, consequently impacting the model's convergence. To mitigate this problem, we conduct a warm-up operation for the decay factor β_t to gradually increase the impact of the sample's exiting layers on the loss weights during training:

$$\beta_t = \gamma_t \cdot \beta_0, \quad (4)$$


where β_0 is positive, and γ_t is the ratio of the current training step to the total training steps.

Classification Objective. To enable various acceleration ratios during inference, the classification objective is defined as the mean of classification losses across all K thresholds:

$$L_{\text{CE}} = \frac{1}{K} \sum_{\tau} L_{\text{CE},\tau}. \quad (5)$$

3.3 Online Signal Calibration

While SWM effectively facilitates the training of multi-exit networks, exiting signals may not consistently reflect sample difficulty, particularly during the early training stages. This affects the reliability of exiting decisions, leading to



(a) Original Energy Score (b) Normalized Energy Score

Figure 3: Energy distribution across layers 2, 6, and 10 for the SST-2 task. Normalization aligns the energy distribution across layers, facilitating threshold selection.

sub-optimal loss weights based on exiting layers. Therefore, we introduce an online signal calibration (OSC) objective to explicitly enlarge the distribution divergence of exiting signals between easy and hard samples. Specifically, for exiting signals that indicate the sample difficulty (e.g. entropy and energy score), our OSC objective is formulated as:

$$L_{\text{OSC}} = \frac{1}{M-1} \sum_{m=1}^{M-1} L_{\text{OSC}}^{(m)}, \quad (6)$$

$$L_{\text{OSC}}^{(m)} = \max(0, \bar{S}_{\text{easy}}^{(m)} - \bar{S}_{\text{hard}}^{(m)} + \epsilon), \quad (7)$$

where $L_{\text{OSC}}^{(m)}$ is the signal calibration loss at the m th layer. $\bar{S}_{\text{easy}}^{(m)}$ and $\bar{S}_{\text{hard}}^{(m)}$ are the mean of exiting signals on easy and hard samples at the m th layer, respectively, and ϵ is the margin parameter shared across layers. For exiting signals negatively correlated with sample difficulty (e.g. softmax score), the calculation for $L_{\text{OSC}}^{(m)}$ in Eq.(7) needs to be replaced with:

$$L_{\text{OSC}}^{(m)} = \max(0, \bar{S}_{\text{hard}}^{(m)} - \bar{S}_{\text{easy}}^{(m)} + \epsilon). \quad (8)$$

Note that we only minimize the signal calibration loss for the first $M-1$ layers, since there is no need to exit at the last layer. Additionally, we define samples as easy or hard depending on whether the internal classifier can predict them correctly, thus the partition may differ across layers.

3.4 Training Objective

The training objective of the COSEE is formulated as the weighted sum of the classification and OSC objective:

$$L = L_{\text{CE}} + \alpha \times L_{\text{OSC}}, \quad (9)$$

where α is a hyper-parameter used to balance the classification and OSC objectives. All internal classifiers are jointly trained with the backbone.

3.5 Exiting Signal

Following E-LANG (Akbari, Banitalebi-Dehkordi, and Zhang 2022), we primarily implement our COSEE with the energy-based exiting signal. The energy score is defined as:

$$E(x; F_m) = -\log \sum_{i=1}^C e^{f_i^{(m)}}, \quad (10)$$

where C is the number of classes, and $f_i^{(m)}$ denotes the logit value of sample x on class i suggested by the m th internal classifier F_m . A lower energy score indicates lower sample difficulty. The exiting criterion is met when the energy score falls below a predefined threshold. To align the energy distribution across different layers for threshold selection, we normalize the original energy scores to $(0, 1)$:

$$E_{\text{norm}}(x; F_m) = (1 + e^{-E(x; F_m)})^{-1}. \quad (11)$$

Figure 3 confirms the superiority of the normalized energy score over the original energy score. In this paper, we mainly conduct experiments with the normalized energy score. Nevertheless, we also verify the effectiveness of the COSEE framework on other exiting signals, i.e., entropy and softmax score (see Section 5.2).

4 Experiments

4.1 Tasks and Datasets

Following Li et al. (2021); Liao et al. (2021), we evaluate COSEE on six classification tasks from the GLUE benchmark (Wang et al. 2019), including SST-2, MRPC, QNLI, RTE, QQP, and MNLI. Data statistics are shown in Table 1.

4.2 Baselines

We compare our COSEE model with three groups of representative and state-of-the-art baselines.

Backbone. We adopt the widely used BERT-base (Devlin et al. 2019) as the backbone for convincing comparisons.

Budget Exiting. We directly train a BERT-base with 6 layers (BERT-6L) to obtain a speed-up ratio of $2.00\times$, establishing a lower bound for early exiting methods as no techniques are employed.

Early Exiting. For signal-based early exiting methods, we choose DeeBERT (Xin et al. 2020), PABEE (Zhou et al. 2020), BERxiT (Xin et al. 2021), LeeBERT (Zhu 2021), GPFEE (Liao et al. 2021), GAML-BERT (Zhu et al. 2021), PALBERT (Balagansky and Gavrilov 2022), and DisentangledEE (Ji et al. 2023). For router-based early exiting methods, we choose state-of-the-art ConsistentEE (Zeng et al. 2024). Notably, some early exiting methods (Sun et al. 2022; Mangrulkar, MS, and Sembium 2022; Zhang et al. 2023; Zhu et al. 2023) are not included due to the difference in backbones. CascadeBERT (Li et al. 2021) and E-LANG (Akbari, Banitalebi-Dehkordi, and Zhang 2022) are excluded for fair comparisons since they implement early exiting within several complete networks instead of a multi-exit network. Refer to the extended version for more details.

4.3 Experimental Settings

Measurement. Since the runtime is unstable across different runs, following Zhang et al. (2022) and Liao et al. (2021), we utilize the saved layers to measure the speed-up ratio:

$$\text{Speed-up Ratio} = \frac{\sum_{m=1}^M M \times N^m}{\sum_{m=1}^M m \times N^m}, \quad (12)$$

where M is the total number of layers and N^m is the number of samples exiting from the m th layer. According to Xin et al. (2020), this metric is proportional to actual runtime.

Dataset	Classes	Train	Test	Task
SST-2	2	67k	1.8k	Sentiment
MRPC	2	3.7k	1.7k	Paraphrase
QQP	2	364k	391k	Paraphrase
MNLI	3	393k	20k	NLI
QNLI	2	105k	5.4k	QA/NLI
RTE	2	2.5k	3k	NLI

Table 1: Dataset Statistics. NLI is the Natural Language Inference task, and QA is the Question Answering task.




Figure 4: Impact of SWM and OSC on the trade-off between performance and efficiency for COSEE with energy.

Training. Our implementation is based on Hugging Face’s Transformers (Wolf et al. 2020). Each internal classifier consists of a single linear layer. We mainly implement our COSEE framework with the normalized energy score if not specified. We also conduct experiments with entropy and softmax scores for generalization analysis. Following the previous work (Zhou et al. 2020; Zhang et al. 2022; Liao et al. 2021), we perform a grid search over learning rates of $\{1e-5, 2e-5, 3e-5, 5e-5\}$, batch sizes of $\{16, 32, 128\}$, α values in Eq.(9) of $\{0.001, 0.01, 0.1, 1.0\}$, and β_0 values in Eq.(4) of $\{0.05, 0.2, 1.0, 10.0\}$. We set ϵ to 0.3 in Eq.(7) and K to 5 in Eq.(5). The maximum sequence length is fixed at 128. We employ a linear decay learning rate scheduler and the AdamW optimizer (Loshchilov and Hutter 2019). We conduct experiments on two RTX4090 GPUs with 24GB.

Inference. Following previous work (Zhang et al. 2022; Liao et al. 2021), we adopt a batch size of 1 for inference, emulating a typical industry scenario where requests from various users arrive one by one. For fair comparisons, we carefully adjust the threshold τ for each task to achieve a similar speed-up ratio as the baseline methods (approximately $2.00\times$) and further compare the trade-off between task performance and inference efficiency.




Figure 5: DIS heatmap of different models for layers 2, 6, and 10 on the SST-2 and QNLI development sets.

4.4 Overall Performance Comparison

Table 2 reports the test results of each early exiting method on the GLUE benchmark with BERT-base as the backbone model. The speed-up ratio is approximately $2.00\times$ ($\pm 38\%$). Overall, our COSEE framework with normalized energy score demonstrates a superior performance-efficiency trade-off across different tasks compared to the baseline methods, which verifies the effectiveness of our design. Notably, our COSEE can even outperform the original BERT-base on RTE and QQP tasks, indicating that our method can effectively alleviate the overthinking problem of PLMs. This suggests that, for easy samples, predictions from intermediate layers may outperform those from the final layer. Our method enables easy samples to exit at shallow classifiers, thereby reducing the inference time while maintaining or even improving the task performance. Besides, our method can save training costs (see Section 4.5) and introduce negligible additional storage overhead (see Section 5.3). We also explore the impact of hyperparameters and statistically analyze the failure cases in the extended version.

Although using energy scores and BERT-base in the primary experiments, we also verify the generality of COSEE on various exiting signals and backbones (see Section 5.2).

4.5 Ablation Studies

Performance-Efficiency Trade-Off. To investigate the effectiveness of SWM and OSC, we plot the performance-efficiency trade-off curves of models trained using different methods on four GLUE development sets, as shown in Figure 4. We can observe both SWM and OSC significantly improve the performance of early exiting across all tasks, especially under high speed-up ratios. This confirms the advantage of our COSEE under high acceleration scenarios, indicating the proposed SWM and OSC effectively facilitate the training of internal classifiers, particularly shallow ones.

Evaluation of Exiting Signals. Difficulty Inversion Score (DIS) is first proposed by Li et al. (2021), an evaluation metric for exiting signals. A higher value indicates a greater correlation between the exiting signal and sample difficulty, thus enabling more reliable exiting decisions. Figure 5 illustrates the DIS of exiting signals generated by different models. The results indicate that OSC explicitly enhances the correlation between the exiting signal and sample difficulty by enlarging the distribution divergence of exiting signals

Method	RTE Acc	MRPC F1/Acc/Mean	QQP F1/Acc/Mean	SST-2 Acc	QNLI Acc	MNLI Acc
BERT-base	66.4 (1.00×)	88.9/-/ (1.00×)	71.2/-/ (1.00×)	93.5 (1.00×)	90.5 (1.00×)	84.6 (1.00×)
BERT-6L [†]	63.9 (2.00×)	85.1/78.6/81.9 (2.00×)	69.7/88.3/79.0 (2.00×)	91.0 (2.00×)	86.7 (2.00×)	80.8 (2.00×)
DeeBERT [†]	64.3 (1.95×)	84.4/77.4/80.9 (2.07×)	70.4/88.8/79.6 (2.13×)	90.2 (2.00×)	85.6 (2.09×)	74.4 (1.87×)
PABEE [†]	64.0 (1.81×)	84.4/77.4/80.9 (2.01×)	70.4/88.6/79.5 (2.09×)	89.3 (1.95×)	88.0 (1.87×)	79.8 (2.07×)
BERxiT	65.7 (2.17×)	86.2/-/ (2.27×)	70.5/-/ (2.27×)	91.6 (2.86×)	89.6 (1.72×)	82.1 (2.33×)
LeeBERT	-	87.1/-/ (1.97×)	-	92.6 (1.97×)	-	83.1 (1.97×)
GPFEЕ	64.5 (2.04×)	87.0/81.8/84.4 (1.98×)	71.2/89.4/80.3 (2.18×)	92.8 (2.02×)	89.8 (1.97×)	83.3 (1.96×)
GAML-BERT	64.3 (1.96×)	87.2/-/ (1.96×)	70.9/-/ (1.96×)	92.8 (1.96×)	84.2 (1.96×)	83.3 (1.96×)
PALBERT [‡]	64.3 (1.48×)	-/80.7 (1.48×)	-/79.3 (1.48×)	91.8 (1.48×)	89.1 (1.48×)	83.0 (1.48×)
DisentangledEE	66.8 (1.25×)	-/83.8 (1.25×)	-/79.4 (1.25×)	92.9 (1.25×)	88.5 (1.25×)	83.0 (1.25×)
ConsistentEE	69.0 (1.85×)	89.0/-/ (1.59×)	-/89.0/- (1.82×)	92.9 (1.85×)	89.9 (1.72×)	83.4 (1.45×)
COSEE (ours)	68.7 (1.96×)	88.0/82.0/85.0 (2.70×)	71.4/89.4/80.4 (2.01×)	93.0 (2.14×)	90.2 (2.56×)	83.4 (1.92×)

Table 2: Performance comparison on the GLUE test set with BERT-base as the backbone. [†] denotes the results taken from GPFEЕ (Liao et al. 2021), and [‡] denotes the results taken from DisentangledEE (Ji et al. 2023). Other baseline results are from their original papers. Our COSEE uses the normalized energy score as the exiting signal. Best results are marked in bold.




Figure 6: Impact of SWM and OSC on training convergence for the SST-2 and QNLI tasks.

between easy and hard samples. Meanwhile, SWM encourages highly discriminative exiting signals by enabling each classifier to emphasize a subset of samples with certain difficulty levels. Also, it is noticeable that the improvements brought by SWM and OSC appear to be significant on shallow layers, which aligns with the observation shown in Figure 4. This is due to the constrained capability of shallow classifiers, which allows for greater potential for improvements in training than deep classifiers.

Training Curves. To further explore the convergence speed of our COSEE during training, we plot the model’s training curves across different training methods on SST-2 and QNLI tasks, as shown in Figure 6. The results indicate that the proposed SWM effectively accelerates the model’s convergence during training. We attribute this to SWM’s ability to reduce data complexity during training by enabling each classifier to emphasize a different subset of samples. Additionally, we can observe that incorporating the OSC objective can slightly impact the model’s convergence speed. Nevertheless, our COSEE framework still maintains an advantage over the vanilla training method.

5 In-depth Analysis

5.1 Visualization of Sample Exiting Layers

To examine the consistency between training and testing under our COSEE framework, we visualize the exiting layer distribution in training and development sets at various

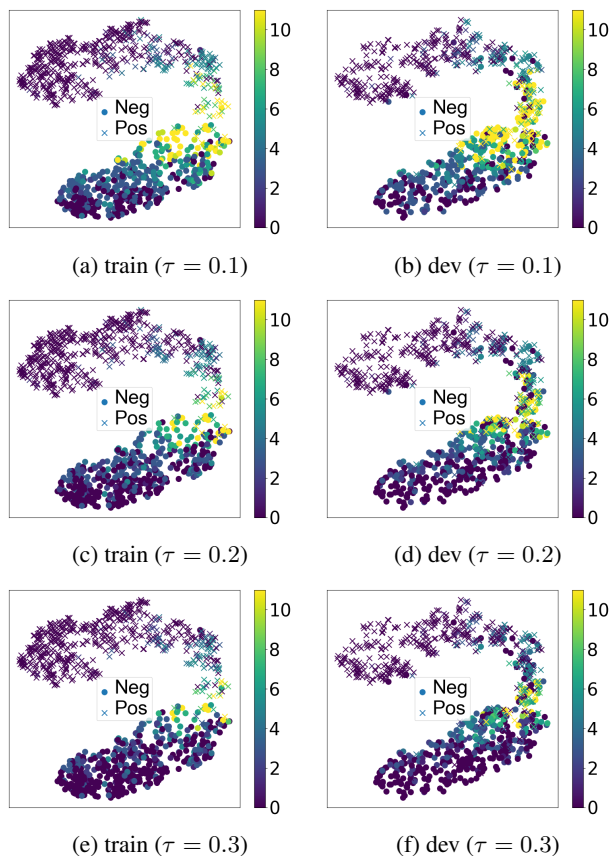


Figure 7: Exiting layer distribution on the training and development sets of SST-2 task under different thresholds τ . Neg and Pos denote negative and positive samples, respectively. The results exhibit consistency between training and development sets across different thresholds.

thresholds, respectively. Figure 7 shows the visualization results for the SST-2 task. At different thresholds, training and development sets demonstrate consistent exiting layer distributions, which verifies the interpretability of our design. Additionally, we can observe that samples near the classi-




Figure 8: Impact of SWM and OSC on the trade-off between performance and efficiency for COSEE with entropy. We include COSEE with energy in comparison.

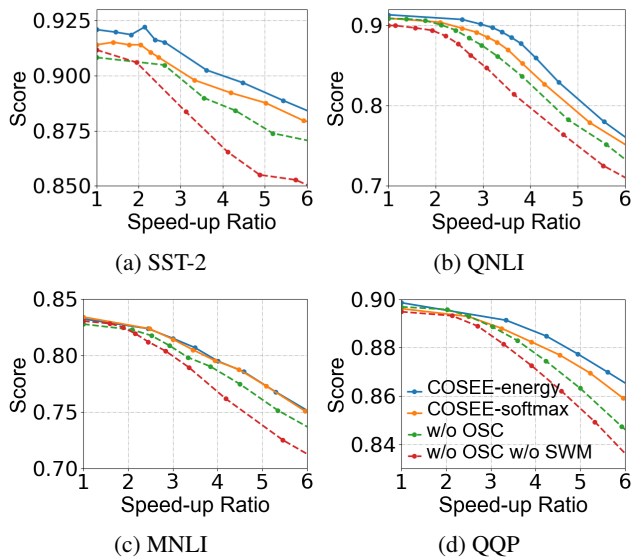


Figure 9: Impact of SWM and OSC on the trade-off between performance and efficiency for COSEE with softmax score.

fication boundary (hard samples) tend to exit at deep classifiers while samples far from the classification boundary (easy samples) tend to exit at shallow classifiers. Furthermore, increasing the threshold will cause a decrease in exiting layers, thus achieving a higher speed-up ratio. These observations align with our intuitive understanding.

5.2 Generality of the COSEE Framework

In this subsection, we explore the generality of our method on various exiting signals and backbones.

Figure 8 and Figure 9 present the experimental results of

Method	Speed-up	QQP	SST-2	QNLI	MNLI	AVG
ALBERT-base [†]	1.00×	79.6	93.3	92.0	85.2	87.5
PABEE [†]	1.95×	79.8	92.4	90.9	84.2	86.8
PALBERT	1.21×	79.1	91.4	90.9	83.2	86.2
DisentangledEE	1.26×	79.3	92.2	91.0	83.5	86.5
COSEE-energy	2.12×	79.6	92.9	91.8	84.8	87.3

Table 3: Test results of different early exiting methods with ALBERT-base as the backbone. The speed-up ratio is averaged across 4 tasks. We report the mean of accuracy and F1-score for QQP, and accuracy for other tasks. [†] denotes results taken from GPFEE (Liao et al. 2021). Other baseline results are taken from DisentangledEE (Ji et al. 2023).

Model	#Params	
	$C = 2$	$C = 3$
BERT-base	109.48M	109.48M
COSEE	+16.92K	+25.38K

Table 4: Parameter volume comparison. C is the number of classes. COSEE introduces negligible extra parameters.

our COSEE framework with entropy and softmax scores, respectively. The results demonstrate the generality of our framework across various exiting signals. Notably, COSEE with energy outperforms COSEE with entropy or softmax scores on most tasks (SST-2, QNLI, and QQP). We attribute this to the superiority of energy scores in distinguishing easy and hard samples compared to entropy and softmax scores, as theoretically demonstrated by Akbari, Banitalebi-Dehkordi, and Zhang (2022). Therefore, we primarily implement COSEE with energy scores in this paper.

Table 3 presents the performance comparison with backbone ALBERT-base. We observe that our COSEE with energy outperforms competitive baseline methods on most tasks, demonstrating its generality across different PLMs.

5.3 Storage Costs Analysis

Table 4 compares the parameter volume of our COSEE model with that of the original BERT-base. Our COSEE model introduces less than 0.03% additional parameters due to incorporating internal classifiers. Notably, the proposed SWM is parameter-free, yet it effectively generates appropriate loss weights for each sample to enhance the training of multi-exit networks.

6 Conclusion

In this paper, we point out that the performance bottleneck of existing early exiting methods primarily lies in the challenge of ensuring consistency between training and testing while enabling flexible adjustments of the speed-up ratio. To remedy this, we propose COSEE, which mimics the test-time early exiting process under various acceleration scenarios based on calibrated exiting signals and then produces the sample-wise loss weights at all classifiers according to the sample’s exiting layer. Our framework is both simple and intuitive. Extensive experiments on the GLUE benchmark demonstrate the superiority and generality of our framework across various exiting signals and backbones.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62376198), the National Key Research and Development Program of China (No. 2022YFB3104700), and the Shanghai Baiyulan Pujiang Project (No. 08002360429).

References

- Akbari, M.; Banitalebi-Dehkordi, A.; and Zhang, Y. 2022. E-LANG: Energy-Based Joint Inferencing of Super and Swift Language Models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5229–5244. Dublin, Ireland: Association for Computational Linguistics.
- Balagansky, N.; and Gavrilov, D. 2022. Palbert: Teaching albert to ponder. *Advances in Neural Information Processing Systems*, 35: 14002–14012.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT (1)*, 4171–4186. Association for Computational Linguistics.
- Gao, X.; Zhu, W.; Gao, J.; and Yin, C. 2023. F-PABEE: flexible-patience-based early exiting for single-label and multi-label text classification tasks. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- He, J.; Zhang, Q.; Ding, W.; Miao, D.; Zhao, J.; Hu, L.; and Cao, L. 2024. DE³-BERT: Distance-Enhanced Early Exiting for BERT based on Prototypical Networks. *CoRR*, abs/2402.05948.
- Ji, Y.; Wang, J.; Li, J.; Chen, Q.; Chen, W.; and Zhang, M. 2023. Early exit with disentangled representation and equiangular tight frame. In *Findings of the Association for Computational Linguistics: ACL 2023*, 14128–14142.
- Kaya, Y.; Hong, S.; and Dumitras, T. 2019. Shallow-Deep Networks: Understanding and Mitigating Network Overthinking. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, 3301–3310. PMLR.
- Lan, Z.; Chen, M.; Goodman, S.; Gimpel, K.; Sharma, P.; and Soricut, R. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *ICLR*. OpenReview.net.
- Li, L.; Lin, Y.; Chen, D.; Ren, S.; Li, P.; Zhou, J.; and Sun, X. 2021. CascadeBERT: Accelerating Inference of Pre-trained Language Models via Calibrated Complete Models Cascade. In *EMNLP (Findings)*, 475–486. Association for Computational Linguistics.
- Liao, K.; Zhang, Y.; Ren, X.; Su, Q.; Sun, X.; and He, B. 2021. A Global Past-Future Early Exit Method for Accelerating Inference of Pre-trained Language Models. In *NAACL-HLT, 2013–2023*. Association for Computational Linguistics.
- Liu, W.; Zhou, P.; Wang, Z.; Zhao, Z.; Deng, H.; and Ju, Q. 2020. FastBERT: a Self-distilling BERT with Adaptive Inference Time. In *ACL*, 6035–6044. Association for Computational Linguistics.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR*, abs/1907.11692.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR (Poster)*. OpenReview.net.
- Mangrulkar, S.; MS, A.; and Sembium, V. 2022. BE3R: BERT based Early-Exit Using Expert Routing. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3504–3512.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Schwartz, R.; Stanovsky, G.; Swayamdipta, S.; Dodge, J.; and Smith, N. A. 2020. The Right Tool for the Job: Matching Model and Instance Complexities. In *ACL*, 6640–6651. Association for Computational Linguistics.
- Sun, T.; Liu, X.; Zhu, W.; Geng, Z.; Wu, L.; He, Y.; Ni, Y.; Xie, G.; Huang, X.; and Qiu, X. 2022. A Simple Hash-Based Early Exiting Approach For Language Understanding and Generation. In *ACL (Findings)*, 2409–2421. Association for Computational Linguistics.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *ICLR (Poster)*. OpenReview.net.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Scao, T. L.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. M. 2020. Transformers: State-of-the-Art Natural Language Processing. In *EMNLP (Demos)*, 38–45. Association for Computational Linguistics.
- Xin, J.; Tang, R.; Lee, J.; Yu, Y.; and Lin, J. 2020. DeeBERT: Dynamic Early Exiting for Accelerating BERT Inference. In *ACL*, 2246–2251. Association for Computational Linguistics.
- Xin, J.; Tang, R.; Yu, Y.; and Lin, J. 2021. BERxiT: Early Exiting for BERT with Better Fine-Tuning and Extension to Regression. In *EACL*, 91–104. Association for Computational Linguistics.
- Zeng, Z.; Hong, Y.; Dai, H.; Zhuang, H.; and Chen, C. 2024. ConsistentEE: A Consistent and Hardness-Guided Early Exiting Method for Accelerating Language Models Inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 19506–19514.
- Zhang, J.; Tan, M.; Dai, P.; and Zhu, W. 2023. Leco: Improving early exiting via learned exits and comparison-based exiting mechanism. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, 298–309.
- Zhang, Z.; Zhu, W.; Zhang, J.; Wang, P.; Jin, R.; and Chung, T. 2022. PCEE-BERT: Accelerating BERT Inference via Patient and Confident Early Exiting. In *NAACL-HLT (Findings)*, 327–338. Association for Computational Linguistics.

Zhou, W.; Xu, C.; Ge, T.; McAuley, J. J.; Xu, K.; and Wei, F. 2020. BERT Loses Patience: Fast and Robust Inference with Early Exit. In *NeurIPS*.

Zhu, W. 2021. LeeBERT: Learned early exit for BERT with cross-level optimization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2968–2980.

Zhu, W.; Wang, P.; Ni, Y.; Xie, G.; and Wang, X. 2023. BADGE: speeding up BERT inference after deployment via block-wise bypasses and divergence-based early exiting. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, 500–509.

Zhu, W.; Wang, X.; Ni, Y.; and Xie, G. 2021. GAML-BERT: improving BERT early exiting by gradient aligned mutual learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3033–3044.