

Accelerating BERT Inference for Sequence Labeling via Early-Exit

Xiaonan Li, Yunfan Shao, Tianxiang Sun, Hang Yan, Xipeng Qiu*, Xuanjing Huang

Shanghai Key Laboratory of Intelligent Information Processing, Fudan University

School of Computer Science, Fudan University

{lixn20, yfshao19, txsun19, hyan19, xpqiu, xjhuang}@fudan.edu.cn

Abstract

Both performance and efficiency are crucial factors for sequence labeling tasks in many real-world scenarios. Although the pre-trained models (PTMs) have significantly improved the performance of various sequence labeling tasks, their computational cost is expensive. To alleviate this problem, we extend the recent successful early-exit mechanism to accelerate the inference of PTMs for sequence labeling tasks. However, existing early-exit mechanisms are specifically designed for sequence-level tasks, rather than sequence labeling. In this paper, we first propose **SENTEE**: a simple extension of **SENTence-level Early-Exit** for sequence labeling tasks. To further reduce computational cost, we also propose **TOKEE**: a **TOKEN-level Early-Exit** mechanism that allows partial tokens to exit early at different layers. Considering the local dependency inherent in sequence labeling, we employed a window-based criterion to decide for a token whether or not to exit. The token-level early-exit brings the gap between training and inference, so we introduce an extra self-sampling fine-tuning stage to alleviate it. The extensive experiments on three popular sequence labeling tasks show that our approach can save up to 66%~75% inference cost with minimal performance degradation. Compared with competitive compressed models such as DistilBERT, our approach can achieve better performance under the same speed-up ratios of 2×, 3×, and 4×.¹

1 Introduction

Sequence labeling plays an important role in natural language processing (NLP). Many NLP tasks can be converted to sequence labeling tasks, such as named entity recognition, part-of-speech tagging,

Chinese word segmentation and Semantic Role Labeling. These tasks are usually fundamental and highly time-demanding, therefore, apart from performance, their inference efficiency is also very important.

The past few years have witnessed the prevailing of pre-trained models (PTMs) (Qiu et al., 2020) on various sequence labeling tasks (Nguyen et al., 2020; Ke et al., 2020; Tian et al., 2020; Mengge et al., 2020). Despite their significant improvements on sequence labeling, they are notorious for enormous computational cost and slow inference speed, which hinders their utility in real-time scenarios or mobile-device scenarios.

Recently, early-exit mechanism (Liu et al., 2020; Xin et al., 2020; Schwartz et al., 2020; Zhou et al., 2020) has been introduced to accelerate inference for large-scale PTMs. In their methods, each layer of the PTM is coupled with a classifier to predict the label for a given instance. At inference stage, if the prediction is confident² enough at an earlier time, it is allowed to exit without passing through the entire model. Figure 1(a) gives an illustration of early-exit mechanism for text classification. However, most existing early-exit methods are targeted at sequence-level prediction, such as text classification, in which the prediction and its confidence score are calculated over a sequence. Therefore, these methods cannot be directly applied to sequence labeling tasks, where the prediction is token-level and the confidence score is required for each token.

In this paper, we aim to extend the early-exit mechanism to sequence labeling tasks. First, we proposed the *SENTence-level Early-Exit* (SENTEE), which is a simple extension of existing early-exit methods. SENTEE allows a sequence of tokens to exit together once the maximum uncertainty

*Corresponding author.

¹Our implementation is publicly available at <https://github.com/LeeSureman/Sequence-Labeling-Early-Exit>.

²In this paper, confident prediction indicates that the uncertainty of it is low.

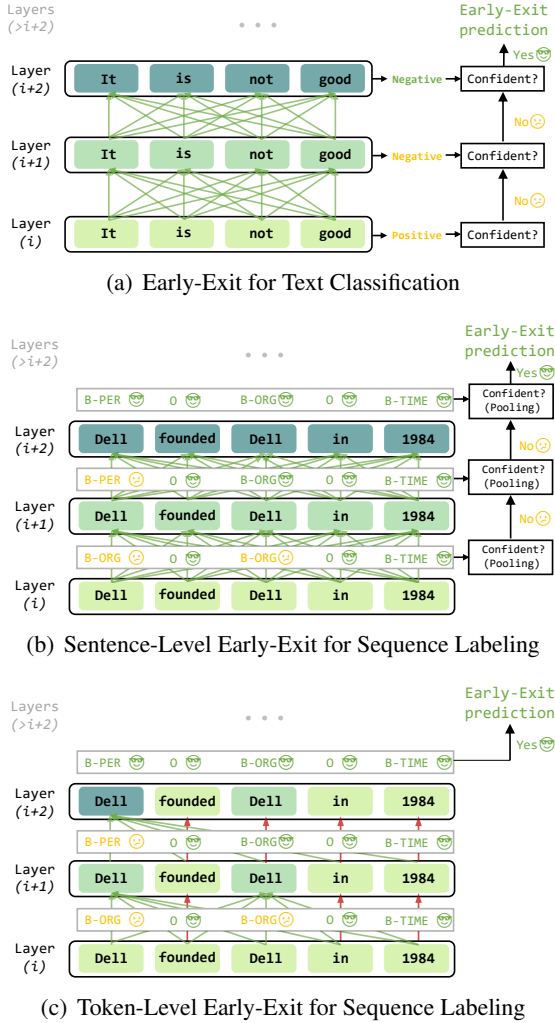


Figure 1: Early-Exit for Text Classification and Sequence Labeling, where \rightarrow represents the self-attention connection, \rightarrow represents simply copying, 😊 and 😞 represents the confident and uncertain prediction. The darker the hidden state block, the deeper layer it is from. Due to space limit, the window-based uncertainty is not reflected.

of the tokens is below a threshold. Despite its effectiveness, we find it redundant for most tokens to update the representation at each layer. Thus, we proposed a *TOKEN-level Early-Exit* (TOKEE) that allows part of tokens that get confident predictions to exit earlier. Figure 1(b) and 1(c) illustrate our proposed SENTEE and TOKEE. Considering the local dependency inherent in sequence labeling tasks, we decide whether a token could exit based on the uncertainty of a window of its context instead of itself. For tokens that are already exited, we do not update their representation but just copy it to the upper layers. However, this will introduce a train-inference discrepancy. To tackle this prob-

lem, we introduce an additional fine-tuning stage that samples the token’s halting layer based on its uncertainty and copies its representation to upper layers during training. We conduct extensive experiments on three sequence labeling tasks: NER, POS tagging, and CWS. Experimental results show that our approach can save up to 66% ~75% inference cost with minimal performance degradation. Compared with competitive compressed models such as DistilBERT, our approach can achieve better performance under speed-up ratio of 2×, 3×, and 4×.

2 BERT for Sequence Labeling

Recently, PTMs (Qiu et al., 2020) have become the mainstream backbone model for various sequence labeling tasks. The typical framework consists of a backbone encoder and a task-specific decoder.

Encoder In this paper, we use BERT (Devlin et al., 2019) as our backbone encoder. The architecture of BERT consists of multiple stacked Transformer layers (Vaswani et al., 2017).

Given a sequence of tokens x_1, \dots, x_N , the hidden state of l -th transformer layer is denoted by $\mathbf{H}^{(l)} = [\mathbf{h}_1^{(l)}, \dots, \mathbf{h}_N^{(l)}]$, and $\mathbf{H}^{(0)}$ is the BERT input embedding.

Decoder Usually, we can predict the label for each token according to the hidden state of the top layer. The probability of labels is predicted by

$$\mathbf{P} = f(\mathbf{W}\mathbf{H}^{(L)}) \in \mathbb{R}^{N \times C}, \quad (1)$$

where N is the sequence length, C is the number of labels, L is the number of BERT layers, \mathbf{W} is a learnable matrix, and $f(\cdot)$ is a simple softmax classifier or conditional random field (CRF) (Lafferty et al., 2001). Since we focus on inference acceleration and PTM performs well enough on sequence labeling without CRF (Devlin et al., 2019), we do not consider using such a recurrent structure.

3 Early-Exit for Sequence Labeling

The inference speed and computational costs of PTMs are crucial bottlenecks to hinder their application in many real-world scenarios. In many tasks, the representations at an earlier layer of PTMs are usually adequate to make a correct prediction. Therefore, early-exit mechanisms (Liu et al., 2020; Xin et al., 2020; Schwartz et al., 2020; Zhou et al., 2020) are proposed to dynamically stop inference

on the backbone model and make prediction with intermediate representation.

However, these existing early-exit mechanisms are built on sentence-level prediction and unsuitable for token-level prediction in sequence labeling tasks. In this section, we propose two early-exit mechanisms to accelerate the inference for sequence labeling tasks.

3.1 Token-Level Off-Ramps

To extend early-exit to sequence labeling, we couple each layer of the PTM with token-level s that can be simply implemented as a linear classifier. Once the off-ramps are trained with the golden labels, the instance has a chance to be predicted and exit at an earlier time instead of passing through the entire model.

Given a sequence of tokens $X = x_1, \dots, x_N$, we can make predictions by the injected off-ramps at each layer. For an off-ramp at l -th layer, the label distribution of all tokens is predicted by

$$\mathbf{P}^{(l)} = f^{(l)}(X; \theta) \quad (2)$$

$$= \text{softmax}(\mathbf{W}\mathbf{H}^{(l)}), \quad (3)$$

where \mathbf{W} is a learnable matrix, $f^{(l)}$ is the token-level off-ramp at l -th layer, $\mathbf{P}^{(l)} = [\mathbf{p}_1^{(l)}, \dots, \mathbf{p}_N^{(l)}]$, $\mathbf{p}_n^{(l)} \in \mathbb{R}^C$, indicates the predicted label distribution at the l -th off-ramp for each token.

Uncertainty of the Off-Ramp With the prediction for each token at hand, we can calculate the *uncertainty* for each token as follows,

$$u_n^{(l)} = \frac{-\mathbf{p}_n^{(l)} \cdot \log \mathbf{p}_n^{(l)}}{\log C}, \quad (4)$$

where $\mathbf{p}_n^{(l)}$ is the label probability distribution for the n -th token.

3.2 Early-Exit Strategies

In the following sections, we will introduce two early-exit mechanisms for sequence labeling, at sentence-level and token-level.

3.2.1 SENTEE: Sentence-Level Early-Exit

Sentence-Level Early-Exit (SENTEE) is a simple extension for sequential labeling tasks based on existing early-exit approaches. SENTEE allows a sequence of tokens to exit together if their uncertainty is low enough. Therefore, SENTEE is to aggregate the uncertainty for each token to obtain

an overall uncertainty for the whole sequence. Here we perform a straight-forward but effective method, i.e., conduct max-pooling³ over uncertainties of all the tokens,

$$u^{(l)} = \max\{u_1^{(l)}, \dots, u_n^{(l)}\}, \quad (5)$$

where $u^{(l)}$ represents the uncertainty for the whole sentence. If $u^{(l)} < \delta$ where δ is a pre-defined threshold, we let the sentence exit at layer l . The intuition is that only when the model is confident of its prediction for the most difficult token, the whole sequence could exit.

3.2.2 TOKEE: Token-Level Early-Exit

Despite the effectiveness of SENTEE (see Table 1), we find it redundant for most simple tokens to be fed into the deep layers. The simple tokens that have been correctly predicted in the shallow layer can not exit (under SENTEE) because the uncertainty of a small number of difficult tokens is still above the threshold. Thus, to further accelerate the inference for sequence labeling tasks, we propose a token-level early-exit (TOKEE) method that allows simple tokens with confident predictions to exit early.

Window-Based Uncertainty Note that a prevalent problem in sequence labeling tasks is the local dependency (or label dependency). That is, the label of a token heavily depends on the tokens around it. To that end, the calculation of the uncertainty for a given token should not only be based on itself but also its context. Motivated by this, we proposed a window-based uncertainty criterion to decide for a token whether or not to exit at the current layer. In particular, the uncertainty for the token x_n at l -th layer is defined as

$$u_n'^{(l)} = \max\{u_{n-k}^{(l)}, \dots, u_{n+k}^{(l)}\}, \quad (6)$$

where k is a pre-defined window size. Then we use $u_n'^{(l)}$ to decide whether the n th token can exit at layer l , instead of $u_n^{(l)}$. Note that window-based uncertainty is equivalent to sentence-level uncertainty when k equals to the sentence length.

Halt-and-Copy For tokens that have exited, their representation would not be updated in the upper layers, i.e., the hidden states of exited tokens are

³We also tried average-pooling, but it brings drastic performance drop. We find that the average uncertainty over the sequence is often overwhelmed by lots of easy tokens and this causes many wrong exits of difficult tokens.

directly *copied* to the upper layers.⁴ Such a *halt-and-copy* mechanism is rather intuitive in two-fold:

- **Halt.** If the uncertainty of a token is very small, there are also few chances that its prediction will be changed in the following layers. So it is redundant to keep updating its representation.
- **Copy.** If the representation of a token can be classified into a label with a high degree of confidence, then its representation already contains the label information. So we can directly copy its representation into the upper layers to help predict the labels of other tokens.

These exited tokens will not attend to other tokens at upper layers but can still be attended by other tokens thus part of the layer-specific query projections in upper layers can be omitted. By this, the computational complexity in self-attention is reduced from $\mathcal{O}(N^2d)$ to $\mathcal{O}(NMd)$, where $M \ll N$ is the number of tokens that have not exited. Besides, the computational complexity of the point-wise FFN can also be reduced from $\mathcal{O}(Nd^2)$ to $\mathcal{O}(Md^2)$.

The halt-and-copy mechanism is also similar to multi-pass sequence labeling paradigm, in which the tokens are labeled their in order of difficulty (easiest first). However, the copy mechanism results in a train-inference discrepancy. That is, a layer never processed the representation from its non-adjacent previous layers during training. To alleviate the discrepancy, we further proposed an additional fine-tuning stage, which will be discussed in Section 3.3.2.

3.3 Model Training

In this section, we describe the training process of our proposed early-exit mechanisms.

3.3.1 Fine-Tuning for SENTEE

For sentence-level early-exit, we follow prior early-exit work for text classification to jointly train the added off-ramps. For each off-ramp, the loss function is as follows,

$$\mathcal{L}_l = \sum_{n=1}^N H\left(y_n, f^{(l)}(X; \theta)_n\right), \quad (7)$$

⁴For English sequence labeling, we use the first-pooling to get the representation of the word. If a word exits, we will halt-and-copy its all wordpieces.

where H is the cross-entropy loss function, N is the sequence length. The total loss function for each sample is a weighted sum of the losses for all the off-ramps,

$$\mathcal{L}_{\text{total}} = \frac{\sum_{l=1}^L w_l \mathcal{L}_l}{\sum_{l=1}^L w_l}, \quad (8)$$

where w_l is the weight for the l -th off-ramp and L is the number of backbone layers. Following (Zhou et al., 2020), we simply set $w_l = l$. In this way, The deeper an off-ramp is, the weight of its loss is bigger, thus each off-ramp can be trained jointly in a relatively balanced way.

3.3.2 Fine-Tuning for TOKEE

Since we equip halt-and-copy in TOKEE, the common joint training off-ramps are not enough. Because the model never conducts halt-and-copy in training but does in inference. In this stage, we aim to train the model to use the hidden state from different previous layers but not only the previous adjacent layer, just like in inference.

Random Sampling A direct way is to uniformly sample halting layers of tokens. However, halting layers at the inference are not random but depends on the difficulty of each token in the sequence. So random sampling halting layers also causes the gap between training and inference.

Self-Sampling Instead, we use the fine-tuned model itself to sample the halting layers. For every sample in each training epoch, we will randomly sample a window size and threshold for it, and then we can conduct TOKEE on the trained model, under the window size and threshold, without halt-and-copy. Thus we get the exiting layer of each token, and we use it to re-forward the sample, by halting and copying each token in the corresponding layer. In this way, the exiting layer of a token can correspond to its difficulty. The deeper a token’s exiting layer is, the more difficult it is. Because we sample the exiting layer using the model itself, we think the gap between training and inference can be further shrunk. To avoid over-fitting during further training, we prevent the training loss from further reducing, similar with the flooding mechanism used by Ishida et al. (2020). We also employ the sandwich rule to stabilize this training stage (Yu and Huang, 2019). We compare self-sampling with random sampling in Section 4.4.4.

Task		NER									
Dataset/ Model	CoNLL2003		Twitter		Ontonotes 4.0		Weibo		CLUE NER		
	F1(Δ)	Speedup	F1(Δ)	Speedup	F1(Δ)	Speedup	F1(Δ)	Speedup	F1(Δ)	Speedup	
	BERT	91.20	1.00 \times	77.97	1.00 \times	79.18	1.00 \times	66.15	1.00 \times	79.11	1.00 \times
	LSTM-CRF	-0.29	-	-12.65	-	-7.37	-	-9.40	-	-7.75	-
$\sim 2X$	DistilBERT-6L	-1.42	2.00 \times	-1.79	2.00 \times	-3.76	2.00 \times	-1.95	2.00 \times	-2.05	2.00 \times
	SENTEE	+0.01	2.07 \times	-0.82	2.03 \times	-1.70	1.94 \times	-0.23	2.11 \times	-0.22	2.04 \times
	TOKEE	-0.09	2.02 \times	-0.20	2.00 \times	-0.20	2.00 \times	-0.15	2.10 \times	-0.24	1.99 \times
$\sim 3X$	DistilBERT-4L	-1.57	3.00 \times	-3.37	3.00 \times	-8.79	3.00 \times	-5.92	3.00 \times	-5.67	3.00 \times
	SENTEE	-0.75	3.02 \times	-5.54	3.03 \times	-8.90	2.95 \times	-2.51	2.87 \times	-6.17	2.98 \times
	TOKEE	-0.32	3.03 \times	-1.32	2.98 \times	-1.48	3.02 \times	-0.70	3.04 \times	-0.92	2.91 \times
$\sim 4X$	DistilBERT-3L	-2.78	4.00 \times	-5.08	4.00 \times	-12.77	4.00 \times	-9.05	4.00 \times	-8.22	4.00 \times
	SENTEE	-6.51	4.00 \times	-11.50	4.14 \times	-14.40	3.78 \times	-8.31	3.86 \times	-14.95	3.91 \times
	TOKEE	-1.44	3.94 \times	-3.96	3.86 \times	-3.81	3.81 \times	-2.04	3.92 \times	-3.16	4.19 \times

Task		POS Tagging				CWS					
Dataset/ Model	ARK Twitter		CTB5 POS		UD POS		CTB5 Seg		UD Seg		
	Acc	Speedup	F1(Δ)	Speedup	F1(Δ)	Speedup	F1(Δ)	Speedup	F1(Δ)	Speedup	
	BERT	91.51	1.00 \times	96.20	1.00 \times	95.04	1.00 \times	98.48	1.00 \times	98.04	1.00 \times
	LSTM-CRF	-0.92	-	-2.48	-	-6.30	-	-0.79	-	-3.03	-
$\sim 2X$	DistilBERT-6L	-0.17	2.00 \times	-0.27	2.00 \times	-0.19	2.00 \times	-0.12	2.00 \times	-0.15	2.00 \times
	SENTEE	-0.49	2.01 \times	-0.19	1.95 \times	-0.35	1.87 \times	-0.02	1.99 \times	-0.11	2.13 \times
	TOKEE	-0.13	1.96 \times	-0.07	2.04 \times	-0.17	1.93 \times	-0.05	2.02 \times	-0.07	2.10 \times
$\sim 3X$	DistilBERT-4L	-0.93	3.00 \times	-0.74	3.00 \times	-1.21	3.00 \times	-0.19	3.00 \times	-0.43	3.00 \times
	SENTEE	-1.98	3.03 \times	-1.63	3.12 \times	-2.47	2.95 \times	-0.11	3.03 \times	-0.47	2.99 \times
	TOKEE	-0.67	2.95 \times	-0.30	3.01 \times	-1.04	2.89 \times	-0.06	3.01 \times	-0.20	3.01 \times
$\sim 4X$	DistilBERT-3L	-1.04	4.00 \times	-1.34	4.00 \times	-4.16	4.00 \times	-0.21	4.00 \times	-1.01	4.00 \times
	SENTEE	-2.96	3.98 \times	-2.37	3.91 \times	-4.78	3.96 \times	-0.83	4.02 \times	-1.23	3.98 \times
	TOKEE	-2.17	3.89 \times	-1.29	3.98 \times	-4.03	3.89 \times	-0.14	3.85 \times	-0.53	3.97 \times

Table 1: Comparison on BERT model. The performance of LSTM-CRF is from previous paper (Tian et al., 2020; Mengge et al., 2020; Yan et al., 2019; Gui et al., 2018), and others are implemented by ourselves.

Task	NER	POS	CWS
Dataset/ Model	Ontonotes4.0 F1(Δ)/Speedup	CTB5 POS F1(Δ)/Speedup	UD Seg F1(Δ)/Speedup
RoBERTa-base			
RoBERTa	79.32 / 1.00 \times	96.29 / 1.00 \times	97.91 / 1.00 \times
RoBERTa 6L	-4.03 / 2.00 \times	-0.11 / 2.00 \times	-0.28 / 2.00 \times
SENTEE	-0.55 / 1.98 \times	-0.08 / 1.85 \times	-0.08 / 2.01 \times
TOKEE	-0.19 / 2.01 \times	-0.06 / 1.98 \times	-0.02 / 1.97 \times
RoBERTa 4L	-11.97 / 3.00 \times	-1.59 / 3.00 \times	-0.76 / 3.00 \times
SENTEE	-5.61 / 3.03 \times	-1.02 / 2.80 \times	-0.51 / 2.99 \times
TOKEE	-0.58 / 3.05 \times	-0.51 / 3.01 \times	-0.12 / 3.00 \times
ALBERT-base			
ALBERT	76.24 / 1.00 \times	95.73 / 1.00 \times	97.00 / 1.00 \times
ALBERT 6L	-5.78 / 2.00 \times	-0.42 / 2.00 \times	-0.20 / 2.00 \times
SENTEE	-1.25 / 2.01 \times	-0.23 / 1.97 \times	-0.02 / 1.98 \times
TOKEE	-0.43 / 1.97 \times	-0.08 / 2.00 \times	-0.03 / 2.01 \times
ALBERT 4L	-12.93 / 3.00 \times	-2.15 / 3.00 \times	-1.38 / 3.00 \times
SENTEE	-6.90 / 2.83 \times	-1.47 / 3.11 \times	-0.41 / 2.97 \times
TOKEE	-0.87 / 2.87 \times	-0.58 / 2.89 \times	-0.06 / 2.98 \times

Table 2: Result on RoBERTa and ALBERT.

4 Experiment

4.1 Computational Cost Measure

We use average floating-point operations (FLOPs) as the measure of computational cost, which denotes how many floating-point operations the model performs for a single sample. The FLOPs is

universal enough since it is not involved with the model running environment (CPU, GPU or TPU) and it can measure the theoretical running time of the model. In general, the lower the model’s FLOPs is, the faster the model’s inference is.

4.2 Experimental Setup

4.2.1 Dataset

To verify the effectiveness of our methods, We conduct experiments on ten English and Chinese datasets of sequence labeling, covering **NER**: CoNLL2003 (Tjong Kim Sang and De Meulder, 2003), Twitter NER (Zhang et al., 2018), Ontonotes 4.0 (Chinese) (Weischedel et al., 2011), Weibo (Peng and Dredze, 2015; He and Sun, 2017) and CLUE NER (Xu et al., 2020), **POS**: ARK Twitter (Gimpel et al., 2011; Owoputi et al., 2013), CTB5 POS (Xue et al., 2005) and UD POS (Nivre et al., 2016), **CWS**: CTB5 Seg (Xue et al., 2005) and UD Seg (Nivre et al., 2016). Besides the standard benchmark dataset like CoNLL2003 and Ontonotes 4.0, we also choose some datasets closer to real-world application to verify the actual utility of our methods, such as Twitter NER and Weibo in social media domain. We use the same dataset prepro-

cessing and split as in previous work (Huang et al., 2015; Mengge et al., 2020; Jia et al., 2020; Tian et al., 2020; Nguyen et al., 2020).

4.2.2 Baseline

We compare our methods with three baselines:

- **BiLSTM-CRF** (Huang et al., 2015; Ma and Hovy, 2016) The most widely used model in sequence labeling tasks before the pre-trained language model prevails in NLP.
- **BERT** The powerful stacked Transformer encoder model, pre-trained on large-scale corpus, which we use as the backbone of our methods.
- **DistilBERT** The most well-known distillation method of BERT. Huggingface released 6 layers DistilBERT for English (Sanh et al., 2019). For comparison, we distill {3, 4} and {3, 4, 6} layers DistilBERT for English and Chinese using the same method.

4.2.3 Hyper-Parameters

For all datasets, We use batch size=10. We perform grid search over learning rate in {5e-6, 1e-5, 2e-5}. We choose learning rate and the model based on the development set. We use the AdamW optimizer (Loshchilov and Hutter, 2019). The warmup step, weight decay is set to 0.05, 0.01, respectively.

4.3 Main Results

For English Datasets, we use the ‘BERT-base-cased’ released by Google (Devlin et al., 2019) as backbone. For Chinese Datasets, we use ‘BERT-wwm’ released by (Cui et al., 2019). The DistilBERT is distilled from the backbone BERT.

To fairly compare our methods with baselines, we turn the speedup ratio of our methods to be consistent with the corresponding static baseline. We report the average performance over 5 times under different random seeds. The overall results are shown in Table 1, where the *speedup* is based on the backbone. We can see both SENTEE and TOKEE brings little performance drop and outperforms DistilBERT in speedup ratio of 2, which has achieved similar effect like existing early-exit for text classification. Under higher speedup, 3× and 4×, SENTEE shows its weakness but TOKEE can still keep a certain performance. And under 2~4× speedup ratio, TOKEE has a lower performance drop than DistilBERT. What’s more, for datasets where BERT can show its power than LSTM-CRF,

e.g., Chinese NER, TOKEE (4×) on BERT can still outperform LSTM-CRF significantly. This indicates the potential utility of it in complicated real-world scenario.

To explore the fine-grained performance change under different speedup ratio, We visualize the speedup-performance trade-off curve on 6 datasets, in Figure 2. We observe that,

- Before the speedup ratio rises to a certain turning point, there is almost no drop on performance. After that, the performance will drop gradually. This shows our methods keep the superiority of existing early-exit methods (Xin et al., 2020).
- As the speedup rises, TOKEE will encounter the speedup turning point later than SENTEE. After both methods reach the turning point, SENTEE’s performance degradation is more drastic than TOKEE. These both indicate the higher speedup ceiling of TOKEE.
- On some datasets, such as CoNLL2003, we observe a little performance improvement under low speedup ratio, we attribute this to the potential regularization brought by early-exit, such as alleviating overthinking (Kaya et al., 2019).

To verify the versatility of our method over different PTMs, we also conduct experiments on two well-known BERT variants, RoBERTa (Liu et al., 2019)⁵ and ALBERT (Lan et al., 2020)⁶, as shown in Table 2. We can see that SENTEE and TOKEE also significantly outperform static backbone internal layer on three Representative datasets of corresponding tasks. For RoBERTa and ALBERT, we also observe the TOKEE can have a better performance than SENTEE under high speedup ratio.

4.4 Analysis

In this section, we conduct a set of detailed analysis on our methods.

4.4.1 The Effect of Window Size

We show the performance change under different k in Figure 3, keeping the speedup ratio consistent. We observe that: (1) when k is 0, in other words, not using window-based uncertainty but token-independent uncertainty, the performance is the almost lowest across different speedup ratio, because it does not consider local dependency

⁵<https://github.com/ymcui/Chinese-BERT-wwm>.

⁶<https://github.com/brightmart/albert.zh>.

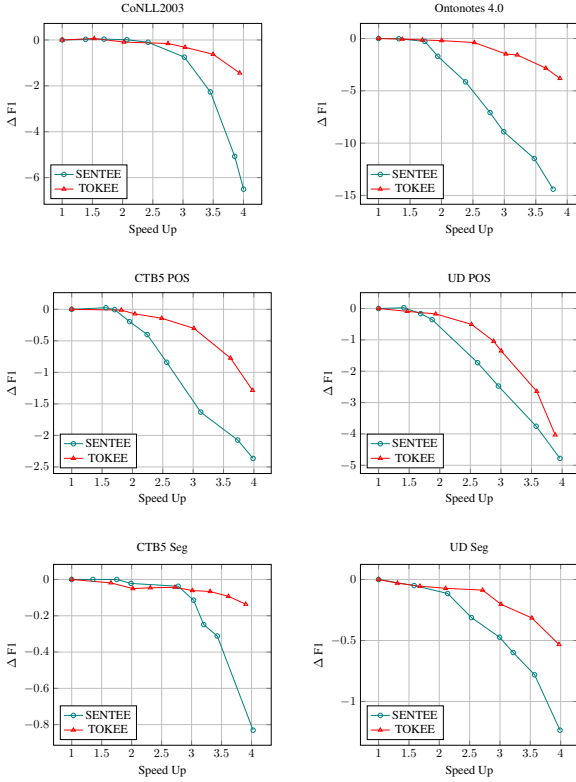


Figure 2: The performance-speedup trade-off curve on six datasets of SENTE and TOKEE.

at all. This shows the necessity of the window-based uncertainty. (2) When k is relatively large, it will bring significant performance drop under high speedup ratio ($3\times$ and $4\times$), like SENTE. (3) It is necessary to choose an appropriate k under high speedup ratio, where the effect of different k has a high variance.

4.4.2 Accuracy V.S. Uncertainty

Liu et al. (2020) verified ‘the lower the uncertainty, the higher the accuracy’ on text classification. Here, we’d like to verify our window-based uncertainty

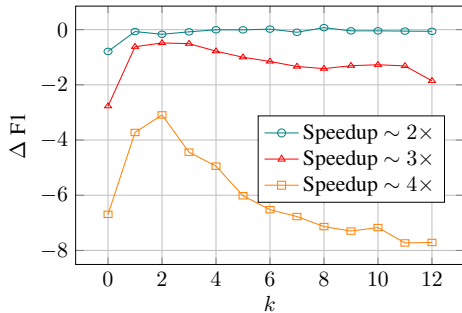


Figure 3: The performance change over different window size under the same speedup ratio.

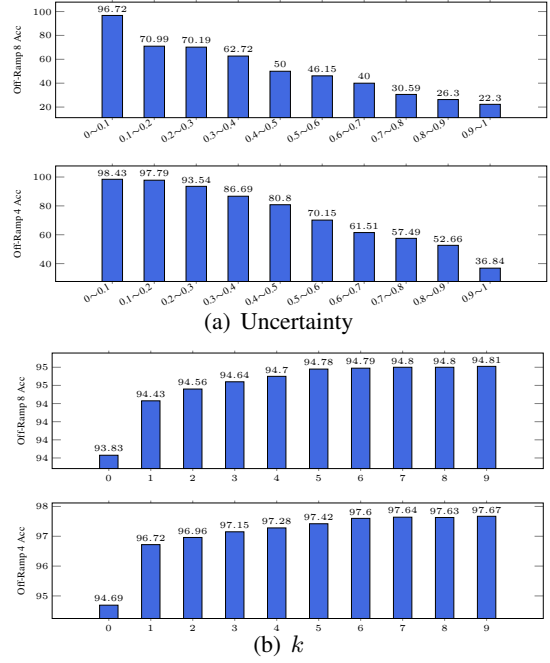


Figure 4: The relation of off-ramps’ accuracy between uncertainty and window size. Two off-ramps of shallow and deep layers are analyzed on CoNLL2003.

on sequence labeling. In detail, we verify the entire window-based uncertainty and its specific hyperparameter, k , on CoNLL2003, shown in Figure 4. For the uncertainty, we intercept the 4th and 8th off-ramps and calculate their accuracy in each uncertainty interval, when $k=2$. The result shown in Figure 4(a) indicates that ‘the lower the window-based uncertainty, the higher the accuracy’, similar as in text classification. For k , we set a certain threshold = 0.3, and calculate accuracy of tokens whose window-based uncertainty is small than the threshold under different k , shown in Figure 4(b). The result shows that, as k increases: (1) The accuracy of screened tokens is higher. This shows that the wider of a token’s low-uncertainty neighborhood, the more accurate the token’s prediction is. This also verifies the validity of window-based uncertainty strategy. (2) The accuracy improvement slows down. This shows the low relevance of distant tokens’ uncertainty and explains why large k performs not well under high speedup ratio: it does not help improving more accurate exiting but slowing down exiting.

4.4.3 Influence of Sequence Length

Transformer-based PTMs, e.g. BERT, face a challenge in processing long text, due to the $O(N^2d)$ computational complexity brought by self-

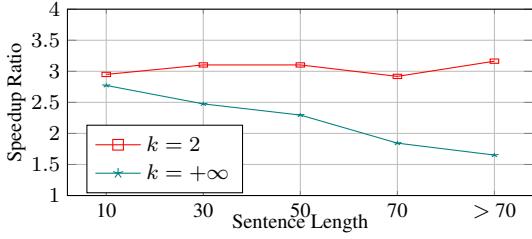


Figure 5: Influence of sentence length.

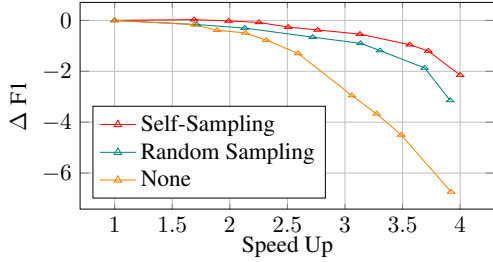


Figure 6: Comparison between different training way.

attention. Since the TOKEE reduces the layer-wise computational complexity from $O(N^2d + Nd^2)$ to $O(NMd + Md^2)$ and SENTEE does not, we’d like to explore their effect over different sentence length. We compare the highest speedup ratio of TOKEE and SENTEE when performance drop < 1 on Ontonotes 4.0, shown in Figure 5. We observe that TOKEE has a stable computational cost saving as the sentence length increases, but SENTEE’s speedup ratio will gradually reduce. For this, we give an intuitive explanation. In general, a longer sentence has more tokens, it is more difficult for the model to give them all confident prediction at the same layer. This comparison reveals the potential of TOKEE on accelerating long text inference.

4.4.4 Effects of Self-Sampling Fine-Tuning

To verify the effect of self-sampling fine-tuning in Section 3.3.2, we compare it with random sampling and no extra fine-tuning on CoNLL2003. The performance-speedup trade-off curve of TOKEE is shown in Figure 6, which shows self-sampling is always better than random sampling for TOKEE. As speedup ratio rises, this trend is more significant. This shows the self-sampling can help more in reducing the gap of training and inference. As for no extra fine-tuning, it will deteriorate drastically at high speedup ratio. But it can roughly keep a certain capability at low speedup ratio, which we attribute to the residual-connection of PTM and similar results were reported by Veit et al. (2016).

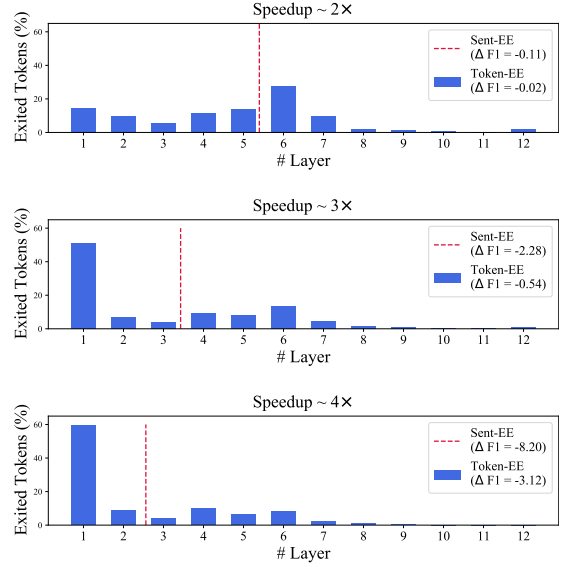


Figure 7: TOKEE layer distribution in a sentence.

4.4.5 Layer Distribution of Early-Exit

In TOKEE, by halt-and-copy mechanism, each token goes through a different number of PTM layers according to the difficulty. We show the average distribution of a sentence’s tokens exiting layers under different speedup ratio on CoNLL2003, in Figure 7. We also draw the average exiting layer number of SENTEE under the same speedup ratio. We observe that as speedup ratio rises, more tokens will exit at the earlier layer but a bit of tokens can still go through the deeper layer even when $4\times$, meanwhile, the SENTEE’s average exiting layer number reduces to 2.5, where the PTM’s encoding power is severely cut down. This gives an intuitive explanation of why TOKEE is more effective than SENTEE under high speedup ratio: although both SENTEE and TOKEE can dynamically adjust computational cost on the sample-level, TOKEE can adjust do it in a more fine-grained way.

5 Related Work

PTMs are powerful but have high computational cost. To accelerate them, many attempts have been made. A kind of methods is to reduce its size, such as distillation (Sanh et al., 2019; Jiao et al., 2020), structural pruning (Michel et al., 2019; Fan et al., 2020) and quantization (Shen et al., 2020).

Another kind of methods is early-exit, which dynamically adjusts the encoding layer number of different samples (Liu et al., 2020; Xin et al., 2020; Schwartz et al., 2020; Zhou et al., 2020; Li et al., 2020). While they introduced early-exit mecha-

nism in simple classification tasks, our methods are proposed for the more complicated scenario: sequence labeling, where it has not only one prediction probability and it's necessary to consider the dependency of token exitings. [Elbayad et al. \(2020\)](#) proposed Depth-Adaptive Transformer to accelerate machine translation. However, their early-exit mechanism is designed for auto-regressive sequence generation, in which the exit of tokens must be in left-to-right order. Therefore, it is unsuitable for language understanding tasks. Different from their method, our early-exit mechanism can consider the exit of all tokens simultaneously.

6 Conclusion and Future Work

In this work, we propose two early-exit mechanisms for sequence labeling: SENTEE and TOKEE. The former is a simple extension of sequence-level early-exit while the latter is specially designed for sequence labeling, which can conduct more fine-grained computational cost allocation. We equip TOKEE with window-based uncertainty and self-sampling finetuning to make it more robust and faster. The detailed analysis verifies their effectiveness. SENTEE and TOKEE can achieve $2\times$ and $3\sim 4\times$ speedup with minimal performance drop.

For future work, we wish to explore: (1) leveraging the exited token's label information to help the exiting of remained tokens; (2) introducing CRF or other global decoding methods into early-exit for sequence labeling.

Acknowledgments

We thank anonymous reviewers for their detailed reviews and great suggestions. This work was supported by the National Key Research and Development Program of China (No. 2020AAA0106700), National Natural Science Foundation of China (No. 62022027) and Major Scientific Research Project of Zhejiang Lab (No. 2019KD0AD01).

References

Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association*

for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Maha Elbayad, Jiatao Gu, Edouard Grave, and Michael Auli. 2020. [Depth-adaptive transformer](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Angela Fan, Edouard Grave, and Armand Joulin. 2020. [Reducing transformer depth on demand with structured dropout](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. [Part-of-speech tagging for twitter: Annotation, features, and experiments](#). In *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA - Short Papers*, pages 42–47. The Association for Computer Linguistics.

Tao Gui, Qi Zhang, Jingjing Gong, Minlong Peng, Di Liang, Keyu Ding, and Xuanjing Huang. 2018. [Transferring from formal newswire domain with hypernet for Twitter POS tagging](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2540–2549, Brussels, Belgium. Association for Computational Linguistics.

Hangfeng He and Xu Sun. 2017. [F-score driven max margin neural network for named entity recognition in chinese social media](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers*, pages 713–718. Association for Computational Linguistics.

Zhiheng Huang, Wei Xu, and Kai Yu. 2015. [Bidirectional LSTM-CRF models for sequence tagging](#). *CoRR*, abs/1508.01991.

Takashi Ishida, Ikko Yamane, Tomoya Sakai, Gang Niu, and Masashi Sugiyama. 2020. [Do we need zero training loss after achieving zero training error?](#) In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4604–4614. PMLR.

Chen Jia, Yuefeng Shi, Qinrong Yang, and Yue Zhang. 2020. [Entity enhanced BERT pre-training for Chinese NER](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6384–6396, Online. Association for Computational Linguistics.

- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. [Tinybert: Distilling BERT for natural language understanding](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 4163–4174. Association for Computational Linguistics.
- Yigitcan Kaya, Sanghyun Hong, and Tudor Dumitras. 2019. [Shallow-deep networks: Understanding and mitigating network overthinking](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3301–3310. PMLR.
- Zhen Ke, Liang Shi, Erli Meng, Bin Wang, Xipeng Qiu, and Xuanjing Huang. 2020. [Unified multi-criteria chinese word segmentation with bert](#).
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Lei Li, Yankai Lin, Shuhuai Ren, Deli Chen, Xuancheng Ren, Peng Li, Jie Zhou, and Xu Sun. 2020. [Accelerating pre-trained language models via calibrated cascade](#). *CoRR*, abs/2012.14682.
- Weijie Liu, Peng Zhou, Zhiruo Wang, Zhe Zhao, Haotang Deng, and Qi Ju. 2020. [FastBERT: a self-distilling BERT with adaptive inference time](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6035–6044, Online. Association for Computational Linguistics.
- 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](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Xue Mengge, Bowen Yu, Zhenyu Zhang, Tingwen Liu, Yue Zhang, and Bin Wang. 2020. [Coarse-to-Fine Pre-training for Named Entity Recognition](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6345–6354, Online. Association for Computational Linguistics.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 14014–14024.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D. Manning, Ryan T. McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. [Universal dependencies v1: A multilingual treebank collection](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*. European Language Resources Association (ELRA).
- Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A. Smith. 2013. [Improved part-of-speech tagging for online conversational text with word clusters](#). In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*, pages 380–390. The Association for Computational Linguistics.
- Nanyun Peng and Mark Dredze. 2015. [Named entity recognition for Chinese social media with jointly trained embeddings](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 548–554, Lisbon, Portugal. Association for Computational Linguistics.
- Xipeng Qiu, TianXiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. [Pre-trained models for natural language processing: A survey](#). *SCIENCE CHINA Technological Sciences*, 63(10):1872–1897.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of bert: smaller, faster, cheaper and lighter](#). In *NeurIPS EMC2 Workshop*.
- Roy Schwartz, Gabriel Stanovsky, Swabha Swayamdipta, Jesse Dodge, and Noah A. Smith.

2020. [The right tool for the job: Matching model and instance complexities](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6640–6651, Online. Association for Computational Linguistics.
- 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](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8815–8821.
- Yuanhe Tian, Yan Song, Xiang Ao, Fei Xia, Xiaojun Quan, Tong Zhang, and Yonggang Wang. 2020. [Joint Chinese word segmentation and part-of-speech tagging via two-way attentions of auto-analyzed knowledge](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8286–8296, Online. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Andreas Veit, Michael J. Wilber, and Serge J. Belongie. 2016. [Residual networks behave like ensembles of relatively shallow networks](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 550–558.
- Ralph Weischedel, Sameer Pradhan, Lance Ramshaw, Martha Palmer, Nianwen Xue, Mitchell Marcus, Ann Taylor, Craig Greenberg, Eduard Hovy, Robert Belvin, et al. 2011. *Ontonotes release 4.0. LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium*.
- Ji Xin, Raphael Tang, Jaejun Lee, Yaoliang Yu, and Jimmy Lin. 2020. [DeeBERT: Dynamic early exiting for accelerating BERT inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2246–2251, Online. Association for Computational Linguistics.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. [CLUE: A Chinese language understanding evaluation benchmark](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Naiwen Xue, Fei Xia, Fu-Dong Chiou, and Martha Palmer. 2005. [The penn chinese treebank: Phrase structure annotation of a large corpus](#). *Nat. Lang. Eng.*, 11(2):207–238.
- Hang Yan, Bocao Deng, Xiaonan Li, and Xipeng Qiu. 2019. [TENER: adapting transformer encoder for named entity recognition](#). *CoRR*, abs/1911.04474.
- Jiahui Yu and Thomas S. Huang. 2019. [Universally slimmable networks and improved training techniques](#). In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 1803–1811. IEEE.
- Qi Zhang, Jinlan Fu, Xiaoyu Liu, and Xuanjing Huang. 2018. [Adaptive co-attention network for named entity recognition in tweets](#).
- Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian J. McAuley, Ke Xu, and Furu Wei. 2020. [BERT loses patience: Fast and robust inference with early exit](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.