

Span Fine-tuning for Pre-trained Language Models

Rongzhou Bao^{1,2,3}, Zhuosheng Zhang^{1,2,3}, Hai Zhao^{1,2,3}

¹ Department of Computer Science and Engineering, Shanghai Jiao Tong University

² Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University, Shanghai, China

³ MoE Key Lab of Artificial Intelligence, Shanghai Jiao Tong University, Shanghai, China

1154904117@qq.com

zhangzs@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn

Abstract

Pre-trained language models (PrLM) have to carefully manage input units when training on a very large text with a vocabulary consisting of millions of words. Previous works have shown that incorporating span-level information over consecutive words in pre-training could further improve the performance of PrLMs. However, given that span-level clues are introduced and fixed in pre-training, previous methods are time-consuming and lack of flexibility. To alleviate the inconvenience, this paper presents a novel span fine-tuning method for PrLMs, which facilitates the span setting to be adaptively determined by specific downstream tasks during the fine-tuning phase. In detail, any sentences processed by the PrLM will be segmented into multiple spans according to a pre-sampled dictionary. Then the segmentation information will be sent through a hierarchical CNN module together with the representation outputs of the PrLM and ultimately generate a span-enhanced representation. Experiments on GLUE benchmark show that the proposed span fine-tuning method significantly enhances the PrLM, and at the same time, offer more flexibility in an efficient way. The code is available at <https://github.com/BAORONGZHOU/span-fine-tuning>.

1 Introduction

Pre-trained language models (PrLM), including ELECTRA (Clark et al., 2020), RoBERTa (Liu et al., 2019b), and BERT (Devlin et al., 2018), have demonstrated strong performance in downstream tasks (Wang et al., 2018). Leveraging a self-supervised training on large text corpora, these

models are able to provide contextualized representations in a more efficient way. For instance, BERT uses Masked Language Modeling and Next Sentence Prediction as pre-training objects and is trained on a corpus of 3.3 billion words.

In order to be adaptive for a wider range of applications, PrLMs usually generate sub-token-level representations (words or subwords) as basic linguistic units. For downstream tasks such as natural language understanding (NLU), span-level representations, e.g. phrases and name entities, are also important. Previous works manifest that by changing pre-training objectives, PrLMs' ability to capture span-level information can be strengthened to some extent. For example, base on BERT, SpanBERT (Joshi et al., 2019) focuses on masking and predicting text spans, instead of sub-token-level information for pre-training. Entity-level masking is used as a pre-training strategy by ERNIE models (Sun et al., 2019; Zhang et al., 2019a). The upper mentioned methods prove that the introduction of span-level information in pre-training to be effective for different NLU tasks.

However, the requirements for span-level information of various NLU tasks differs a lot from case to case. The methods of introducing span-level information in pre-training phase, proposed by previous works, do not fit into the requirements and cannot improve the performance for all NLU tasks. For instance, ERNIE models (Sun et al., 2019) perform remarkably well in Relation Classification, while underperforms BERT in language inference tasks, such as MNLI (Nangia et al., 2017). Therefore, it is imperative to develop a strategy to incorporate span-level information into PrLMs in a more flexible and universally adaptive way. This paper proposes a novel approach, Span Fine-tuning (SF), to leverage span-level information in fine-tuning phase and therefore formulate a task-specific strategy. Compared to existing works, our approach requires less time and computing resources, and is

*Corresponding author. This paper was partially supported by National Key Research and Development Program of China (No. 2017YFB0304100), Key Projects of National Natural Science Foundation of China (U1836222 and 61733011). This work was supported by Huawei Noah's Ask Lab

more adaptive to various NLU tasks.

In order to maximize the value and contribution of span-level information, in addition to the sub-token-level representation generated by BERT, Span Fine-tuning also applies a computationally motivated segmentation to further improve the overall experience. Although various techniques, such as dependency parsing (Zhou et al., 2019) or semantic role labeling (SRL) (Zhang et al., 2019b), have been used as auxiliary tools for sentence segmentation, these methods demand extra parsing procedure, which increase complexities in actual practice. Span Fine-tuning first leverages a pre-sampled n -gram dictionary to segment input sentences into spans. Then, the sub-token-level representations within the same span are combined to generate a span-level representation. Finally, the span-level representations are merged with sub-token-representations into a sentence-level representation. In this way, the sentence-level representation is able to contain and maximize the utilization of both sub-token-level and span-level information.

Experiments are conducted on the GLUE benchmark (Wang et al., 2018), which includes many NLU tasks, such as text classification, semantic similarity, and natural language inference. Empirical results demonstrate that Span Fine-tuning is able to further improve the performance of different PrLMs, including BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019b) and SpanBERT (Joshi et al., 2019). The result of the experiments with SpanBERT indicates that Span Fine-tuning leverages span-level information differently compared to PrLMs pre-trained with span-level information, which shows the distinguishness of our approach. It is also verified in ablation studies and analysis that Span Fine-tuning is essential for further performance improvement for PrLMs.

2 Related Work

2.1 Pre-trained language models

Learning reliable and broadly applicable word representations has been an ongoing heated focus for natural language processing community. Language modeling objectives are proved to be effective for distributed representation generation (Mnih and Hinton, 2009). By generating deep contextualized word representations, ELMo (Peters et al., 2018) advances state of the art for several NLU tasks. Leveraging Transformer (Vaswani et al., 2017),

BERT (Devlin et al., 2018) further advances the field of transfer learning. Recent PrLMs are established based on the various extensions of BERT, including using GAN-style architecture (Clark et al., 2020), applying a parameter sharing strategy (Lan et al., 2019), and increasing the efficiency of parameters (Liu et al., 2019b).

2.2 Span-level pre-training methods

Previous works manifest that the introduction of span-level information in pre-training phase can improve PrLMs' performance. In the first place, BERT leverages the prediction of single masked tokens as one of the pre-training objectives. Due to the use of WordPiece embeddings (Wu et al., 2016), BERT is able to segment sentences into sub-word level tokens, so that the masked tokens are at sub-token-level, e.g. "###ing". (Devlin et al., 2018) shows that masking the whole word, rather than only single tokens, can further enhance the performance of BERT. Later, it is proved by (Sun et al., 2019; Zhang et al., 2019a) that the masking of entities is also helpful for PrLMs. By randomly masking adjoining spans in pre-training, SpanBERT (Joshi et al., 2019) can generate better representation for given texts. AMBERT (Zhang and Li, 2020) achieves better performance than its precursors in NLU tasks by incorporating both sub-token-level and span-level tokenization in pre-training. The upper mentioned studies all focus on introducing span-level information in pre-training. To the best of our knowledge, the introduction of span-level information in fine-tuning is still a white space to explore, which makes our approach a valuable attempt.

2.3 Integration of fine-grained representation

Different formats of downstream tasks require sentence-level representations, such as natural language inference (Bowman et al., 2015; Nangia et al., 2017), semantic textual similarity (Cer et al., 2017) and sentiment classification (Socher et al., 2013). Besides directly pre-training the representation of coarser granularity (Le and Mikolov, 2014; Logeswaran and Lee, 2018), a lot of methods have been explored to obtain a task-specific sentence-level representation by integrating fine-grained token-level representations (Conneau et al., 2017). Kim (2014) shows that by applying a convolutional neural network (CNN) on top of pre-trained word vectors, we can get a sentence-level representation that is well adapted to classification

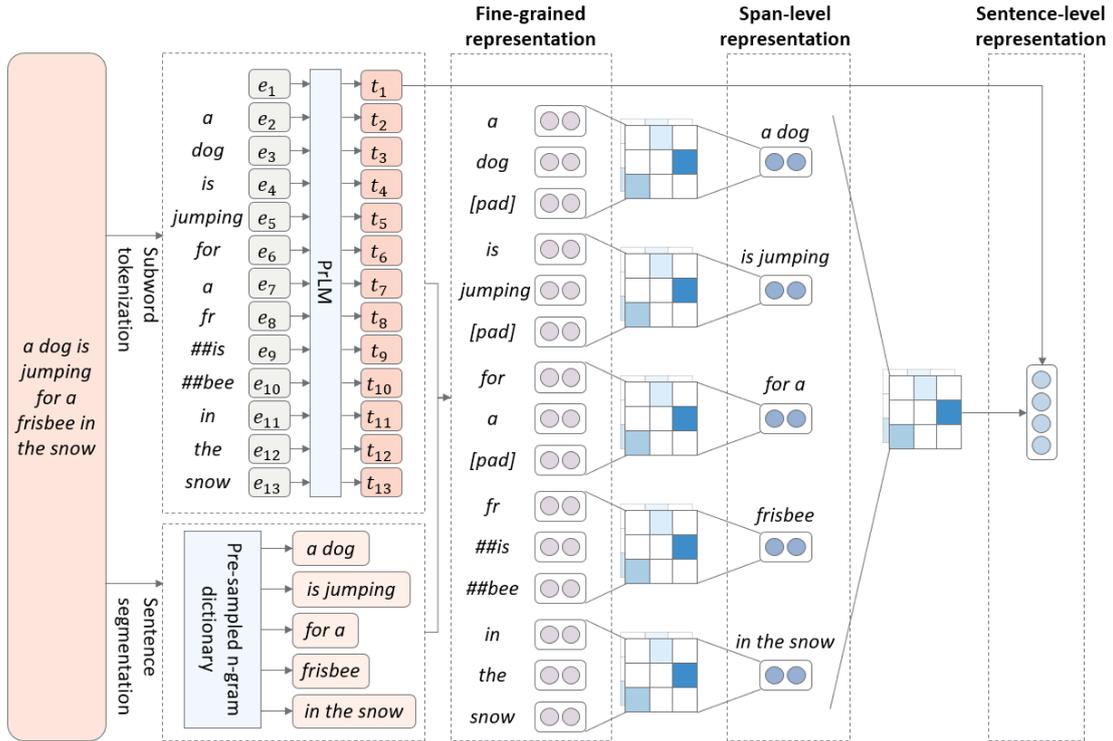


Figure 1: Overview of the framework of our proposed method

tasks. Lin et al. (2017) leverage a self-attentive module over hidden states of a BiLSTM to generate sentence-level representations. Zhang et al. (2019b) use a CNN layer to extract word-level representations from sub-word representations and combine them with word-level semantic role representations. Inspired by these methods, after a series of preliminary attempts, we choose a hierarchical CNN architecture to recombine fine-grained representations to coarse-grained ones.

3 Methodology

Figure 1 demonstrates the overall framework of Span Fine-tuning, which essentially uses BERT as a foundation and incorporates segmentation as an auxiliary tool. The figure does not exhaustively depict the details of BERT, given the model is relatively popular and ubiquitous. Further information on BERT is available in (Devlin et al., 2018). In Span Fine-tuning, an input sentence is divided into sub-word-level tokens and then sent to BERT to generate sub-token-level representations. At the same time, the input is segmented into spans based on n -gram statistics. By combining the segmentation information with sub-token-level representations generated by BERT, we divided the repre-

sentation into several spans. Then, the spans are sent through a hierarchical CNN module to obtain a span-level information enhanced representation. Finally, the sub-token-level representation of [CLS] token generated by BERT and the span-level information enhanced representation are concatenated and form a final representation, which maximized the value of both sub-token-level and span-level information for NLU tasks.

3.1 Sentence Segmentation

Semantic role labeling (SRL) (Zhang et al., 2019b) and dependency parsing (Zhou et al., 2019) have been used as auxiliary tools for segmentation by previous works. Nonetheless, these techniques demand additional parsing procedures, and therefore increase complexities for real application. In order to obtain a simpler and more convenient segmentation, based on frequency, we select meaningful n -grams appeared in wikitext-103 dataset* to form a pre-sampled dictionary.

We use the dictionary to match n -grams from the head of each input sentence. n -grams with greater lengths are prioritized, while unmatched tokens remain the same. In this way, we are able to

*PMI method has also been tried to adjust our dictionary, but the result is not competitive.

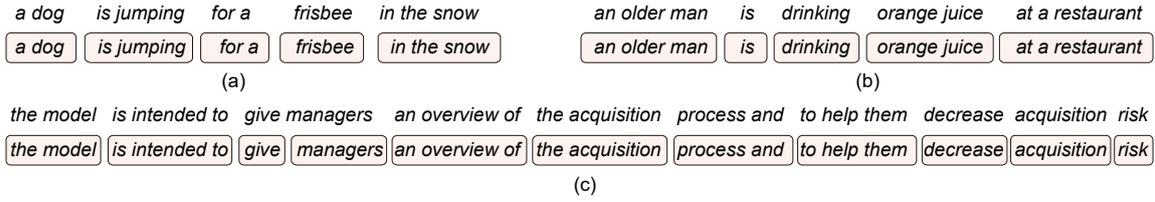


Figure 2: Segmentation Examples

obtain a specific segmentation of the input sentence. Figure 2 demonstrates some examples of sentence segmentation from the GLUE dataset.

3.2 Sentence Encoder Architecture

An input sentence $X = \{x_1, \dots, x_n\}$ is given with a length n . The sentence is firstly divided into sub-word tokens (with a special token [CLS] at the beginning) and converted to sub-token-level representations $E = \{e_1, \dots, e_m\}$ (usually m is larger than n) according to embeddings proposed by (Wu et al., 2016). Then, the transformer encoder (BERT) captures the contextual information for each token by self-attention and generates a sequence of sub-token-level contextual embeddings $T = \{t_1, \dots, t_m\}$, in which t_1 is the contextual representation of special token [CLS]. Based on the segmentation generated by the n -gram statistics, the sub-token-level contextual representations are combined into several spans $\{C_1, \dots, C_r\}$, with r as a hyperparameter indicating the max number of spans for all processed sentences. Each C_i contains several contextual sub-token-level representations extracted from T denoted as $\{t_1^i, t_2^i, \dots, t_l^i\}$. l is another hyperparameter representing the max number of tokens for all the spans. A CNN-Maxpooling module is applied to each C_i to get a span-level representation c_i :

$$\begin{aligned} c_j^i &= \text{ReLU}(W_1 [t_j^i, t_{j+1}^i, \dots, t_{j+k-1}^i] + b_1), \\ c_i &= \text{MaxPooling}(c_1^i, \dots, c_r^i), \end{aligned} \quad (1)$$

where W_1 and b_1 are trainable parameters and k is the kernel size. Based on the span-level representations $\{c_1, \dots, c_r\}$, another CNN-Maxpooling module is applied to obtain a sentence-level representation s with enhanced span-level information:

$$\begin{aligned} s'_i &= \text{ReLU}(W_2 [c_i, c_{i+1}, \dots, c_{i+k-1}] + b_2), \\ s &= \text{MaxPooling}(s'_1, \dots, s'_r), \end{aligned} \quad (2)$$

Finally, we concatenate s with the contextual sub-token-level representation t_1 of special token [CLS] provided by BERT, and generate a sentence-level representation s^* that maximizes the value of both sub-token-level and span-level information for NLU tasks: $s^* = s \diamond t_1$.

3.3 Tasks and Datasets

To evaluate Span Fine-tuning, experiments are conducted on nine NLU benchmark datasets, including text classification, natural language inference, semantic similarity. Eight of which are available from the GLUE benchmark (Wang et al., 2018). And the rest one is SNLI (Bowman et al., 2015), a widely accepted natural language inference dataset.

3.4 Pre-trained Language Model

We leverage the PyTorch implementation of BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019b) and SpanBERT (Joshi et al., 2019) based on HuggingFace’s codebase[†] (Wolf et al., 2019) as our PrLMs and baselines.

4 Experiments

4.1 Set Up

We select all the n -grams with $n \leq 5$, which occurs more than ten times in the wikitext-103 dataset, to form a dictionary. The pre-sampled dictionary, containing more than 400 thousand n -grams, is used to segment input sentences. During segmentation, two hyperparameters are involved: r representing the largest number of spans in a sentence, and l indicating the largest number of tokens included in a span. In order to maintain different dimensions of features for each input sentence, padding and tail are employed. We set r equals to 16, and based on NLU tasks, choose l in $\{64, 128\}$.

The fine-tuning procedure is as the same as BERT’s. Adam is used as the optimizer. The initial learning rate is in $\{1e-5, 2e-5, 3e-5\}$, the warm-up rate is 0.1, and the L2 weight decay is 0.01. The

[†]<https://github.com/huggingface>

Method	CoLA (mc)	SST-2 (acc)	MNLI m/mm(acc)	QNLI (acc)	RTE (acc)	MRPC (F1)	QQP (F1)	STS-B (pc)	Avg	Gain
<i>In literature</i>										
BERT _{BASE}	52.1	93.5	84.6/83.4	-	66.4	88.9	71.2	87.1	78.3	
BERT _{LARGE}	60.5	94.9	86.7/85.9	92.7	70.1	89.3	72.1	87.6	80.5	
BERT-1seq [‡]	63.5	94.8	88.0/87.4	93.0	72.1	91.2	72.1	89.0	83.5	1.0
SpanBERT	64.3	94.8	88.1/87.7	94.3	79.0	90.9	71.9	89.9	84.5	
<i>Our implementation</i>										
BERT _{BASE}	51.4	92.1	84.4/83.5	90.3	67.1	88.3	71.3	85.1	79.3	
BERT _{BASE} + SF	55.1	93.6	84.8/84.3	90.6	69.6	88.7	71.9	86.5	80.4	1.1
BERT _{WWM}	61.1	93.6	87.1/86.5	93.9	77.3	90.0	71.9	88.1	83.3	1.1
BERT _{WWM} + SF	62.9	94.1	87.6/87.0	94.3	81.4	91.1	72.4	89.1	84.4	

Table 1: Test sets performance on GLUE benchmark. All the results are obtained from (Liu et al., 2019a), (Radford et al., 2018). For a simple demonstration, problematic WNLI set are excluded, and we do not show the accuracy of the datasets have F1 scores. *mc* and *pc* denote the Matthews correlation and Pearson correlation respectively.

batch size is set in {16, 32, 48}. The maximum number of epochs is set in {2,3,4,5} based on NLU tasks. Input sentences are divided into subtokens and converted to WordPiece embeddings, with a maximum length in {128, 256}. The output size of the CNN layer is the same as the hidden size of PrLM, and the kernel size is set to 3.

4.2 Results with BERT as PrLM

Two released BERT (Devlin et al., 2018), BERT Large Whole Word Masking and BERT Base, are first used as pre-trained encoder and baselines for Span Fine-tuning. Compared with BERT Large, BERT Large Whole Word Masking reach a better performance, since it uses whole-word masking in pre-training phase. Therefore, we select BERT Large Whole Word Masking as a stronger baseline. The results indicate that Span Fine-tuning can maximize the contribution of span-level information, even when compared to a stronger baseline.

Table 1 exhibits the results on the GLUE datasets, showing that Span Fine-tuning can significantly improve the performance of PrLMs. Since our approach leverages BERT as a foundation, and undergoes the the same evaluation procedure, it is evident that the performance gain is fully contributed by the newly introduced Span Fine-tuning.

In order to test the statistical significance of the results, we follow the procedure of (Zhang et al., 2020). We use the McNemars test, this test is designed for paired nominal observations, and it is appropriate for binary classification tasks. The p-value is defined as the probability of obtaining a result equal to or more extreme than what was ob-

served under the null hypothesis. The smaller the p-value, the higher the significance. A commonly used level of reliability of the result is 95%, written as $p = 0.05$. As shown in table 2, compared with the baseline, for all the binary classification tasks of GLUE benchmark, our method pass the significance test.

	CoLA	SST-2	QNLI	RTE	MRPC	QQP
p-value	0.005	0.012	0.023	0.009	0.008	0.031

Table 2: Results of McNemars tests for binary classification tasks of GLUE benchmark, tests are conducted based on the results of best run of BERT_{BASE} and BERT_{BASE} + SF.

Span Fine-tuning can reach the same performance improvement as previous methods. As illustrated in Table 1, on average, SpanBERT can improve the result by one percentage point over the baseline (BERT-1seq), while Span Fine-tuning is able to achieve an improvement of 1.1 percentage points over our baseline. However, as showed in Table 3, Span Fine-tuning requires considerably less time and computing resources compared to the large-scale pre-training for span-level information incorporation. When the Span Fine-tuning is adopted, the extra parameters are only 3 percent of the total parameters of the adopted PrLMs for every downstream task, and introduce little extra overhead.

Besides, Span Fine-tuning is more flexible and adaptive compared to previous methods. Table

[‡]The baseline of SpanBert, a BERT pre-trained without next sentence prediction object.

Method	Time	Resource
Pre-train	32 days	32 Volta V100
Span Fine-tune	12 hours max	2 Titan RTX

Table 3: The comparison between incorporation of span-level information in pre-training and Span Fine-tuning .

1 shows that Span Fine-tuning is able to achieve stronger results on all NLU tasks compared to the baseline, whereas the results of SpanBERT in certain tasks, such as Quora Question Pairs and Microsoft Research Paraphrase Corpus, are worse than its baseline. Since for spanBERT, the utilization of span-level information is fixed for every downstream task. Whereas in our method, an extra module designed to incorporate span-level information is trained during the fine-tuning, which can be more dynamically adapted to different downstream tasks.

Table 5 indicates that Span Fine-tuning also enhances the result of PrMLs on the SNLI benchmark. The improvement achieved by Span Fine-tuning is similar to published state-of-the-art accomplished by SemBERT. However, compared to SemBERT, Span Fine-tuning saves a lot more time and computing resources. Span Fine-tuning merely leverages a pre-sampled dictionary to facilitate segmentation, whereas SemBERT leverages a pre-trained semantic role labeller, which brings extra complexities to the whole segmentation process.

Furthermore, Span Fine-tuning is different from SemBERT in terms of motivation, method and contribution factors. The motivation of SemBERT is to enhance PrLMs by incorporating explicit contextual semantics, whereas the motivation of our work is to let PrLMs leverage span-level information in fine-tuning. When it comes to method, SemBERT concatenate the original representations given by BERT with representations of semantic role labels, in comparison, our work directly leverages a segmentation given by a pre-sampled dictionary to generate span-enhanced representation and requires no pre-trained semantic role labeler. The gain of SemBERT comes from semantic role labels while the gain of our work comes from the specific segmentation, which is very different.

It’s worth noticing that semantic role labeler can also generate segmentation. However, semantic role labeler will generate multiple segmentation for sentence which has various predicate-argument

structures. Furthermore, such segmentation is sometimes coarse-grained (with span more than ten words), which is unpractical for our work.

4.3 Results with Stronger PrLMs

In addition to BERT, we also apply Span Fine-tuning to stronger PrLMs, such as RoBERTa (Liu et al., 2019b) and SpanBERT (Joshi et al., 2019), which optimize BERT by enhancing pre-training procedure and predicting text spans rather than single tokens respectively.

Table 4 shows that Span Fine-tuning can strengthen both RoBERTa and SpanBERT. RoBERTa is a already very strong baseline, we remarkably improve the performance of RoBERTa on RTE by four percentage points. SpanBERT already incorporated span-level information during the pre-training, but the results still support that Span Fine-tuning utilizes the span-level formation and improves the performance of PrLMs in a different dimension.

5 Analysis

5.1 Ablation Study

In order to determine the key factors in Span Fine-tuning, a series of studies are conducted on the dev sets of eight NLU tasks. BERT_{BASE} is chosen as the PrLM for the ablation studies. As shown in Table 6, three sets of ablation studies are performed. For experiment BERT_{BASE} + CNN, only a hierarchical CNN structure is applied in to evaluate whether the improvement comes from the extra parameters. To illustrate, we firstly apply two layers of CNN over the token-level representations given by BERT. Then, a max pooling operation is applied to get the sentence-level representation. Finally, the sentence-level representation and the 'CLS' representation of BERT is concatenated and sent to the classifier. In this way, the parameters of BERT_{BASE} + CNN are the same as in our method. For experiment BERT_{BASE} + CNN + Random SF, random sentence segmentation is applied to the experiment to test if the proposed segmentation method of Span Fine-tuning really functions in span-level information incorporation. For experiment BERT_{BASE} + CNN + NLTK SF, we conduct the experiments using a pre-trained chunker from Natural Language Toolkit to see whether the proposed segmentation method of Span Fine-tuning can achieve further improvements.

Method	CoLA (mc)	SST-2 (acc)	MNLI m/mm(acc)	QNLI (acc)	RTE (acc)	MRPC (F1)	QQP (acc)	STS-B (pc)	Avg. -
SpanBERT _{LARGE}	64.3	94.8	88.1/87.7	94.3	79.0	90.9	89.5	89.9	86.5
SpanBERT _{LARGE} + SF	65.9	95.1	88.4/88.1	94.3	83.3	92.1	90.9	90.1	87.6
RoBERTa _{LARGE}	68.0	96.4	90.2/ 90.2	94.7	86.6	90.9	92.2	92.4	89.0
RoBERTa _{LARGE} + SF	68.9	96.1	90.3/90.2	94.3	90.6	92.8	92.2	92.4	89.8

Table 4: Results on test sets of GLUE benchmark with stronger baseline, we average results from three different random seeds.

Method	Dev	Test
BERT _{WWM}	92.0	91.4
BERT _{WWM} + SF	92.3	91.7
SemBERT _{WWM}	92.2	91.9

Table 5: Accuracy on dev and test sets of SNLI. SemBERT_{WWM} (Zhang et al., 2019b) is the published SoTA on SNLI.

method	Avg Score
BERT _{BASE}	82.6
BERT _{BASE} + CNN	82.5
BERT _{BASE} + Random SF [§]	83.0
BERT _{BASE} + NLTK SF [¶]	83.7
BERT _{BASE} + SF	84.2

Table 6: Ablation studied on dev sets of GLUE benchmark, we average results from three different random seeds.

The results of the experiment BERT_{BASE} + CNN suggest that the improvement is unlikely to come from the extra parameters, since it reduce the overall performance by 0.1 percent. The experiment BERT_{BASE} + Random SF and BERT_{BASE} + NLTK SF indicate that the segmentation generated by a pre-train chunker or even random segmentation can also achieve enhancement under the Span Fine-tuning structure. However, a pre-trained chunker demands additional part-of-speech parsing process, while our segmentation method only relies on a pre-sampled dictionary and saves a lot more time, and at the same time, achieves greater improvement. Our Span Fine-tuning is able to remarkably enhance the result on all NLU tasks, raising average score by 1.6 percentage points. Overall, the

[§]Random SF represents Span Fine-tuning with randomly segmented sentences.

[¶]NLTK SF represents Span Fine-tuning with segmentation generated by an NLTK pre-trained chunker.

result of experiments indicate that the performance improvement is primarily a result of our unique segmentation method.

5.2 Encoder Architecture

(Conneau et al., 2017) mentions that the influence of sentence encoder architectures on PrLM performance varies a lot from case to case. (Toshniwal et al., 2020) also suggests that different span representations can affect NLPs tasks greatly.

Method	Dev	Test
CNN-Max	90.9	90.9
CNN-CNN	91.3	91.1
Attention -Max	90.7	90.5
Attention-Attention	90.8	90.8

Table 7: Accuracy on dev and test sets of SNLI. SemBERT_{WWM} (Zhang et al., 2019b) is the published SOTA on SNLI.

To evaluate the effectiveness of our encoder architecture, we replace the component of the encoding layer and the overall structure respectively. For the component of the encoding layer, CNN (Kim, 2014) and the Self-attentive module (Lin et al., 2017) are compared. For the overall structure, two structures are considered: a single layer structure with the max-pooling operation and a hierarchical structure.

By matching every component of the encoding layer with the overall structure, four different encoder architectures are generated: CNN-Maxpooling, CNN-CNN, Attention-Maxpooling, Attention-Attention. Experiments are conducted on SNLI dev and test sets. Table 7 suggests that the hierarchical CNN (CNN-CNN) is most suitable encoder architecture for us.

^{||}Attention indicate the Self-attentive module (Lin et al., 2017).

5.3 Size of n -gram Dictionary

Since our segmentation method is based on a pre-sampled dictionary, the size of dictionaries will have a large impact on segmentation results. Figure 3 depicts how the average number of spans in the sentences changed along with dictionary size in CoLA and MRPC datasets. At the origin, where no segmentation is applied, every token is considered as a span. The number of spans drops significantly, as the dictionary size grows and more n -grams are matched and grouped together.

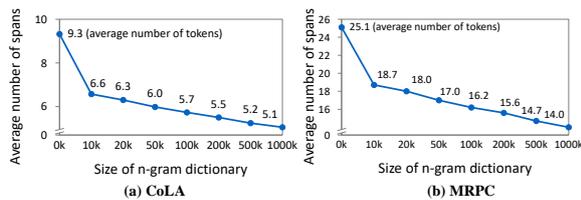


Figure 3: Influence of dictionary size on the average number of spans in the sentences

To evaluate the influence of dictionary size on PrLM performance, experiments on the dev sets of two NLU tasks are implemented: CoLA and MRPC. To concentrate on the impact of segmentation and reduce the impacts from sub-token-level representations provided by PrLM, the concatenation process is not applied to this experiment. Rather, the span-level information enhanced representations are directly sent to a dense layer to generate prediction. As demonstrated in figure 4, the incorporation of pre-sampled n -gram dictionary generates a stronger performance compared to random segmentation. Moreover, dictionaries of medium sizes (20k to 200k) commonly result in better performance. Such trend can be explained by intuition, give dictionaries of small sizes are likely to omit meaningful n -grams, whereas the ones of large sizes tend to over-combine meaningless n -grams.

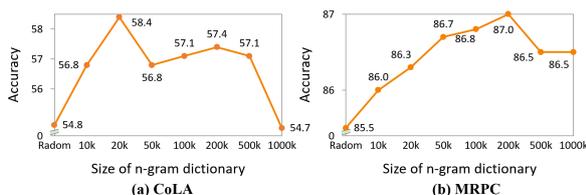


Figure 4: The influence of the size of n -gram dictionary

5.4 Span Fine-tuning for Token-Level Tasks

The upper mentioned experiments are conducted on the GLUE benchmark, whose tasks are all at the sentence level. Nevertheless, token-level representations are needed in many other NLU task, such as name-entity recognition (NER). Our approach can be applied to token-level tasks with simple modification of encoder architecture (e.g. removing the pooling layer of CNN module). Table 8 shows the results of our approach on the CoNLL-2003 Named Entity Recognition (NER) task (Tjong Kim Sang and De Meulder, 2003) with BERT as our PrLM.

	BERT _{BASE}	BERT _{BASE} +SF	BERT _{LARGE}	BERT _{LARGE} +SF
Dev	91.7	92.1	92.3	92.5
Test	95.7	96.2	96.5	96.8

Table 8: F1 on dev and test sets of named entity recognition from CoNLL-2003, we average results from three different random seeds.

6 Conclusion

This paper proposes Span Fine-tuning that maximize the advantages of flexible span-level information in fine-tuning with sub-token-level representations generated by PrLMs. Leveraging a reasonable segmentation provided by a pre-sampled n -gram dictionary, Span Fine-tuning can further enhance the performance of PrLMs on various downstream tasks. Compared with previous span pre-training methods, our Span Fine-tuning remains competitive for the following reasons:

Task-adaptive For methods that incorporate span-level information in pre-training, the utilization of span-level information is unlikely easily adjusted for every downstream task as span pre-training has been fixed after tremendous computational cost. In our method, the extra module designed to incorporate span-level information is trained during the fine-tuning, resulting in a more dynamically adaptation to different downstream tasks.

Flexible to PrLMs Our approach can be generally applied to various PrLMs including RoBERTa and SpanBERT.

Novelty Our approach can further improve the performance of PrLMs pre-trained with span-level information (e.g. SpanBERT). Such result implies that we our method utilizes the span-level

information in a different manner comparing with PrLMs pre-trained with span-level information, which makes our method distinguished comparing with previous works.

7 Acknowledgement

This paper was partially funded by Chinese National Key Laboratory of Science and Technology on Information System Security, we thank them for their generous support.

References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *ICLR*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2019. Spanbert: Improving pre-training by representing and predicting spans. *arXiv preprint arXiv:1907.10529*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196.
- Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lajanugen Logeswaran and Honglak Lee. 2018. An efficient framework for learning sentence representations. *arXiv preprint arXiv:1803.02893*.
- Andriy Mnih and Geoffrey E Hinton. 2009. A scalable hierarchical distributed language model. In *Advances in neural information processing systems*.
- Nikita Nangia, Adina Williams, Angeliki Lazaridou, and Samuel R Bowman. 2017. The repeval 2017 shared task: Multi-genre natural language inference with sentence representations. In *RepEval*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL-HLT*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *Technical report*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. ERNIE: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *NAACL*.
- Shubham Toshniwal, Haoyue Shi, Bowen Shi, Lingyu Gao, Karen Livescu, and Kevin Gimpel. 2020. A cross-task analysis of text span representations. In *Proceedings of the 5th Workshop on Representation Learning for NLP*.
- 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 *NIPS*.

- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *2018 EMNLP Workshop BlackboxNLP*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Xinsong Zhang and Hang Li. 2020. Ambert: A pre-trained language model with multi-grained tokenization. *arXiv preprint arXiv:2008.11869*.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019a. ERNIE: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*.
- Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2019b. Semantics-aware bert for language understanding. *arXiv preprint arXiv:1909.02209*.
- Zhuosheng Zhang, Junjie Yang, and Hai Zhao. 2020. Retrospective reader for machine reading comprehension.
- Junru Zhou, Zhuosheng Zhang, and Hai Zhao. 2019. Limit-bert: Linguistic informed multi-task bert. *arXiv preprint arXiv:1910.14296*.