

# LIMIT-BERT : Linguistics Informed Multi-Task BERT

Junru Zhou<sup>1,2,3</sup>, Zhuosheng Zhang<sup>1,2,3</sup>, Hai Zhao<sup>1,2,3\*</sup>, Shuailiang Zhang<sup>1,2,3</sup>

<sup>1</sup>Department of Computer Science and Engineering, Shanghai Jiao Tong University

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

<sup>3</sup>MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University  
{zhoujunru, zhangzs}@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn

## Abstract

In this paper, we present Linguistics Informed Multi-Task BERT (LIMIT-BERT) for learning language representations across multiple linguistics tasks by Multi-Task Learning. LIMIT-BERT includes five key linguistics tasks: Part-Of-Speech (POS) tags, constituent and dependency syntactic parsing, span and dependency semantic role labeling (SRL). Different from recent Multi-Task Deep Neural Networks (MT-DNN), our LIMIT-BERT is fully linguistics motivated and thus is capable of adopting an improved masked training objective according to syntactic and semantic constituents. Besides, LIMIT-BERT takes a semi-supervised learning strategy to offer the same large amount of linguistics task data as that for the language model training. As a result, LIMIT-BERT not only improves linguistics tasks performance, but also benefits from a regularization effect and linguistics information that leads to more general representations to help adapt to new tasks and domains. LIMIT-BERT outperforms the strong baseline Whole Word Masking BERT on both dependency and constituent syntactic/semantic parsing, GLUE benchmark, and SNLI task. Our practice on the proposed LIMIT-BERT also enables us to release a well pre-trained model for multi-purpose of natural language processing tasks once for all.

## 1 Introduction

Recently, pre-trained language models have shown greatly effective across a range of linguistics inspired natural language processing (NLP) tasks such as syntactic parsing, semantic parsing and

so on (Zhou and Zhao, 2019; Zhou et al., 2020; Ouchi et al., 2018; He et al., 2018b; Li et al., 2019), when taking the latter as downstream tasks for the former. In the meantime, introducing linguistic clues such as syntax and semantics into the pre-trained language models may furthermore enhance other downstream tasks such as various Natural Language Understanding (NLU) tasks (Zhang et al., 2020a,b). However, nearly all existing language models are usually trained on large amounts of unlabeled text data (Peters et al., 2018; Devlin et al., 2019), without explicitly exploiting linguistic knowledge. Such observations motivate us to jointly consider both types of tasks, pre-training language models, and solving linguistics inspired NLP problems. We argue such a treatment may benefit from two-fold. (1) Joint learning is a better way to let the former help the latter in a bidirectional mode, rather than in a unidirectional mode, taking the latter as downstream tasks of the former. (2) Naturally empowered by linguistic clues from joint learning, pre-trained language models will be more powerful for enhancing downstream tasks. Thus we propose Linguistics Informed Multi-Task BERT (LIMIT-BERT), making an attempt to incorporate linguistic knowledge into pre-training language representation models. The proposed LIMIT-BERT is implemented in terms of Multi-Task Learning (MTL) (Caruana, 1993) which has shown useful, by alleviating overfitting to a specific task, thus making the learned representations universal across tasks.

Since universal language representations are learning by leveraging large amounts of unlabeled data which has quite different data volume compared with linguistics tasks dataset such as Penn Treebank (PTB)<sup>1</sup> (Marcus et al., 1993).

To alleviate such data unbalance on multi-task

<sup>1</sup>PTB is an English treebank with syntactic tree annotation which only contains 50k sentences.

\* 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), Huawei-SJTU long term AI project, Cutting-edge Machine reading comprehension and language model.

learning, we apply semi-supervised learning approach that uses a pre-trained linguistics model<sup>2</sup> to annotate large amounts of unlabeled text data and to combine with gold linguistics task dataset as our final training data. For such pre-processing, it is easy to train our LIMIT-BERT on large amounts of data with many tasks concurrently by simply summing up all the concerned losses together. Moreover, since every sentence has labeled with predicted syntax and semantics, we can furthermore improve the masked training objective by fully exploiting the known syntactic or semantic constituents during the language model training process. Unlike the previous work MT-DNN (Liu et al., 2019b) which only fine-tunes BERT on GLUE tasks, our LIMIT-BERT is trained on large amounts of data in a semi-supervised way and firmly supported by explicit linguistic clues.

We verify the effectiveness and applicability of LIMIT-BERT on Propbank semantic parsing<sup>3</sup> in both span style (CoNLL-2005) (Carreras and Màrquez, 2005) and dependency style, (CoNLL-2009) (Hajič et al., 2009) and Penn Treebank (PTB) (Marcus et al., 1993) for both constituent and dependency syntactic parsing. Our empirical results show that semantics and syntax can indeed benefit the language representation model via multi-task learning and outperforms the strong baseline Whole Word Masking BERT (BERT<sub>WWM</sub>).

## 2 Tasks and Datasets

LIMIT-BERT includes five types of downstream tasks: Part-Of-Speech, constituent and dependency syntactic parsing, span and dependency semantic role labeling (SRL).

Both span (constituent) and dependency are two broadly-adopted annotation styles for either semantics or syntax, which have been well studied and discussed from both linguistic and computational perspectives (Chomsky, 1981; Li et al., 2019).

Constituency parsing aims to build a constituency-based parse tree from a sentence that represents its syntactic structure according to a phrase structure grammar. While dependency parsing identifies syntactic relations (such as an adjective modifying a noun) between word pairs in a sentence. The constituent structure

<sup>2</sup>The model may jointly predict syntax and semantics for both span and dependency annotation styles, which is from (Zhou et al., 2020) and joint learning with POS tag.

<sup>3</sup>It is also called semantic role labeling (SRL) for the semantic parsing task over the Propbank.

is better at disclosing phrasal continuity, while the dependency structure is better at indicating dependency relation among words.

Semantic role labeling (SRL) is dedicated to recognizing the predicate-argument structure of a sentence, such as **who** did **what** to **whom**, where and when, etc. For argument annotation, there are two formulizations. One is based on text spans, namely span-based SRL. The other is dependency-based SRL, which annotates the syntactic head of argument rather than the entire argument span. SRL is an important method to obtain semantic information beneficial to a wide range of NLP tasks (Zhang et al., 2019; Mihaylov and Frank, 2019).


BERT is typically trained on quite large unlabeled text datasets, BooksCorpus and English Wikipedia, which have 13GB plain text, while the datasets for specific linguistics tasks are less than 100MB. Thus we employ semi-supervised learning to alleviate such data unbalance on multi-task learning by using a pre-trained linguistics model to label BooksCorpus and English Wikipedia data. The pre-trained model jointly learns POS tags and the four types of structures on semantics and syntax, in which the latter is from the XLNet version of (Zhou et al., 2020), giving state-of-the-art or comparable performance for the concerned four parsing tasks. During training, we set 10% probability to use gold syntactic parsing and SRL data: Penn Treebank (PTB) (Marcus et al., 1993), span style SRL (CoNLL-2005) (Carreras and Màrquez, 2005) and dependency style SRL (CoNLL-2009) (Hajič et al., 2009).

### 2.1 Linguistics-Guided Mask Strategy

BERT applies two training objectives: Masked Language Model (LM) and Next Sentence Prediction (NSP) based on WordPiece embeddings (Wu et al., 2016) with a 30,000 token vocabulary. For Masked LM training objective, BERT uses training data generator to choose 15% of the token positions at random for mask replacement and predict the masked tokens<sup>4</sup>. Since using different masking strategy can improve model performance such as the Whole Word Masking<sup>5</sup> which masks all of the tokens corresponding to a word at once, we further improve the masking strategy by exploit-

<sup>4</sup>Actually, BERT applies three replacement strategies: (1) the [MASK] token 80% of the time (2) random token 10% of the time (3) the unchanged *i*-th token 10% of the time. This work uses the same replacement strategies.


<sup>5</sup><https://github.com/huggingface/transformers>



Span and Dependency SRL

federal paper board [MASK] paper and wood [MASK] .

(a) Semantic Phrase Masking.



Constituent Syntactic Tree

[MASK] [MASK] [MASK] sells paper and wood products .

(b) Syntactic Phrase Masking.

Figure 1: Syntactic and Semantic Phrase Masking strategy. In figure (a) the predicates *sells* and *products* are replaced by [MASK] while in figure (b) each token of constituent *federal paper board* also has been masked.

ing available linguistic clues, syntactic or semantic constituents (phrases)<sup>6</sup>, predicted by our pre-trained linguistics model as discussed in Section 2. Thus, we apply three mask strategies at random for each sentence: Syntactic Phrase Masking, Semantic Phrase Masking, and Whole Word Masking. Syntactic/Semantic Phrase Masking (SPM) means that all the tokens corresponding to a syntactic/semantic phrase are masked, as shown in Figure 1. The overall masking rate and replacement strategy remain the same as BERT, we still predict each masked WordPiece token independently. Intuitively, it makes sense that SPM is strictly more powerful than original Token Masking or Whole Word Masking, since SPM may choose and predict the meaningful words or phrases such as verb predicates or noun phrases.

### 3 LIMIT-BERT Model

#### 3.1 Overview

The architecture of the LIMIT-BERT is shown in Figure 2. Our model includes four modules: token representation, Transformer encoder, language modeling layers, task-specific layers including syn-

<sup>6</sup>Syntactic phrases indicate the constituent subtrees while semantic phrases represent as predicate or argument in span SRL.

tactic and semantic scorers and decoders. We take multi-task learning (MTL) (Caruana, 1993) sharing the parameters of token representation and Transformer encoder, while language modeling layers and the top task-specific layers have independent parameters. The training procedure is simple that we just sum up the language model loss with task-specific losses together.

#### 3.2 Token Representation

Following BERT token representation (Devlin et al., 2019), the first token is always the [CLS] token. If input  $X$  is packed by a sentence pair  $X_1; X_2$ , we separate the two sentences with a special token [SEP] (“packed by” means connect two sentences as BERT training). The Transformer encoder maps  $X$  into a sequence of input embedding vectors, one for each token, which is a sum of the corresponding word, segment, and positional embeddings.

If we apply BERT training data (BooksCorpus and English Wikipedia), we use pair sentences packed to perform next sentence prediction and only take the first sentence including [CLS] and [SEP] token for later linguistics tasks. While using gold linguistics task data (PTB, CoNLL-2005, and CoNLL-2009) with 10% probability, we only take one sentence as input that [CLS] and [SEP] are first and last tokens respectively.

Since input sequence  $X$  is based on WordPiece token, we only take the last WordPiece vector of the word in the last layer of Transformer encoder as our sole word representation for later linguistics tasks input to keep the same length of the token and label annotations.

#### 3.3 Transformer Encoder

The Transformer encoder in our model is adapted from (Vaswani et al.), which transforms the input representation vectors into a sequence of contextualized embedding vectors with shared representation across different tasks. We use the pre-trained parameters of BERT (Devlin et al., 2019) as our encoder initialization for faster convergence.

#### 3.4 Language Modeling Layers

BERT training applies masked language modeling (MLM) as a training objective which corrupts the input by replacing some tokens with a special token [MASK] and then lets the model reconstruct the original tokens. While in our LIMIT-BERT training, the linguistics specific tasks and MLM train-

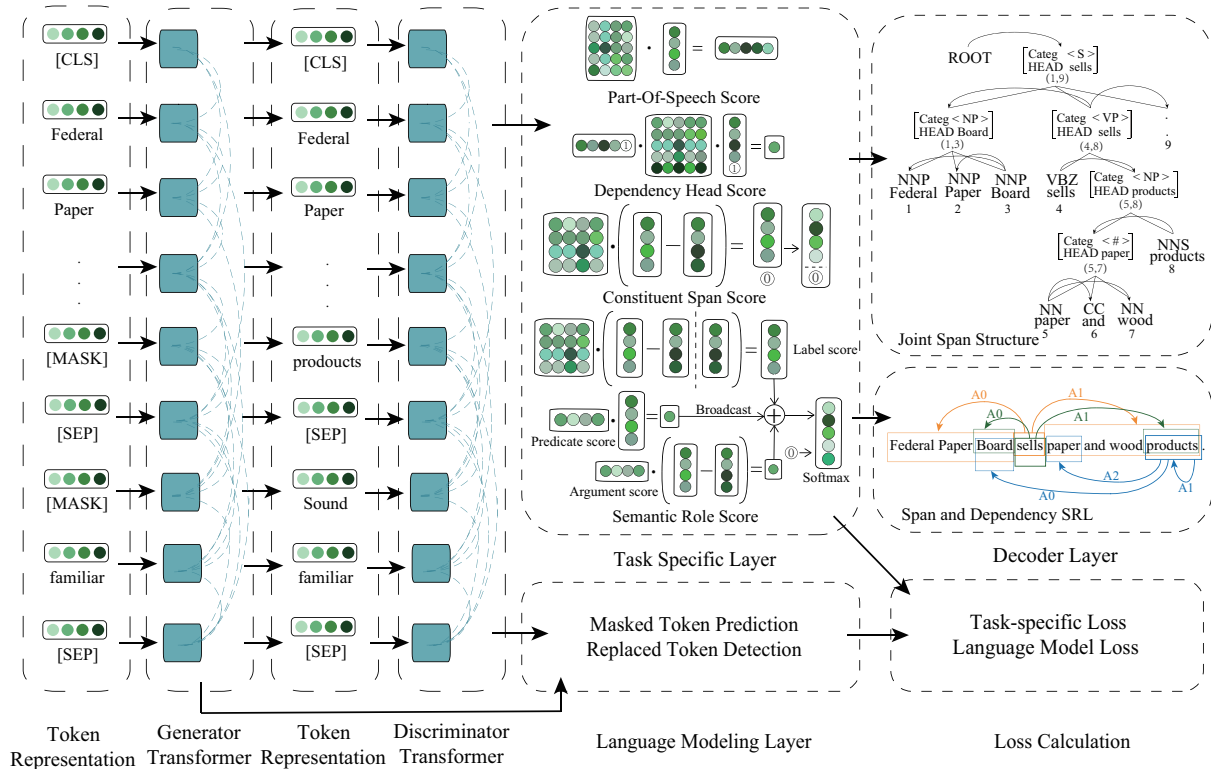


Figure 2: The framework of LIMIT-BERT.

ing take the same input; thus the [MASK] tokens raise a mismatch problem that the model sees artificial [MASK] tokens during MLM training but not when being fine-tuned and inference on linguistics tasks. Besides, due to learning bidirectional representations, MLM approaches incur a substantial computational cost increase because the network only learns from 15% of the tokens per example and needs more training time to converge.

Recently (Yang et al., 2019; Clark et al., 2020) have made attempts to alleviate such a difficulty. The latter applies a replaced token detection task in their ELECTRA model. Instead of masking the input, ELECTRA corrupts it by replacing some input tokens with plausible alternatives sampled from a small generator network, which is close to the original input without [MASK] tokens.

We adopt the ELECTRA training approach in our LIMIT-BERT, which lets the generator  $G$  and discriminator  $D$  share the same parameters and embedding as shown in Figure 2. The generator  $G$  is identical to BERT training (Devlin et al., 2019) that predicts the masked tokens and next sentence and sums token mask loss and next sentence predict loss<sup>7</sup> as  $J_G(\theta)$ . Then the discriminator  $D$  takes the

<sup>7</sup>If using gold linguistics task data, we only compute the token mask loss.

predicted tokens by generator  $G$ <sup>8</sup> and is trained to distinguish tokens that have been replaced by generator  $G$  which is a simple binary classification of each token with loss  $J_D(\theta)$ . At last, we take the output vector  $X$  of discriminator  $D$  to feed the following task-specific layers and sum the loss of  $J_G(\theta)$  and  $J_D(\theta)$  as the final language modeling loss  $J_{lm}(\theta)$ :

$$J_{lm}(\theta) = J_G(\theta) + \lambda J_D(\theta),$$

where  $\lambda$  is set to 50 as the same as ELECTRA.

### 3.5 Task-specific Layers

Firstly, we rebuild word representations from the WordPiece tokens for linguistics tasks. Then we follow (Zhou et al., 2020) to construct the task-specific layers, including scoring layer and decoder layer. The former scores three types of linguistic objectives, dependency head, syntactic constituent and semantic role. The latter is to generate the legal linguistics structures.

**Word Level Construction** Suppose that  $X$  is the output of the discriminator Transformer encoder, we pre-process the WordPiece sequence vector  $X$  for linguistics specific tasks learning which

<sup>8</sup>For the non-masked tokens, we take the original tokens as input.



are based on word level. We only take the first sentence  $X_1$  including token [CLS] and [SEP] of packed sentence pair  $(X_1; X_2)$ . Then we convert WordPiece sequence vector to word-level by simply taking the last WordPiece token vector of the word as the representation of the whole word.

### Scoring Layer

After the word-level construction, we calculate the POS tag, syntactic constituent, dependency head, and semantic role scores, following the training way as (Zhou et al., 2020) to construct syntactic constituent, dependency head, and semantic role scores objective loss which are represented as  $J_1(\theta)$ ,  $J_2(\theta)$  and  $J_3(\theta)$  respectively.

For POS tagging model training, we apply a one-layer feedforward network and minimize the negative log-likelihood of the gold POS tag  $gp_i$  of each word, which is implemented as a cross-entropy loss:

$$J_4(\theta) = -\log P_\theta(gp_i|x_i),$$

where  $x_i$  is word vector inside  $X$ .

Utilizing these specific task scores, we do a sum to obtain the linguistics task loss  $J_{lt}(\theta)$  for training:

$$J_{lt}(\theta) = J_1(\theta) + J_2(\theta) + J_3(\theta) + J_4(\theta).$$

At last, our LIMIT-BERT is trained for simply minimizing the overall loss:

$$J_{overall}(\theta) = J_{lm}(\theta) + J_{lt}(\theta).$$

**Decoder Layer** For syntactic parsing, we apply the *joint span* CKY-style algorithm to generate constituent and dependency syntactic tree simultaneously by following (Zhou and Zhao, 2019).

For span and dependency SRL, we use a single dynamic programming decoder according to the uniform semantic role score following the non-overlapping constraints: span semantic arguments for the same predicate do not overlap (Punyakanok et al., 2008). For further details of the scoring and decoder layer, please refer to (Zhou et al., 2020).

## 4 Experiments

### 4.1 Evaluation

We use the model of (Zhou et al., 2020) with fine-tuned uncased BERT<sub>WWM</sub> (whole word masking) as the baseline<sup>9</sup>. For fairly compared to the baseline

<sup>9</sup>Our codes and the pre-trained models : <https://github.com/DoodleJZ/LIMIT-BERT>.

BERT<sub>WWM</sub>, we also extract the language modeling layer of LIMIT-BERT and use the same model of (Zhou et al., 2020) to fine-tune. We evaluate our proposed model LIMIT-BERT and baseline model BERT<sub>WWM</sub> on CoNLL-2009 shared task (Hajič et al., 2009) for dependency-style SRL, CoNLL-2005 shared task (Carreras and Màrquez, 2005) for span-style SRL both using the Propbank convention (Palmer et al., 2005), and English Penn Treebank (PTB) (Marcus et al., 1993) for constituent syntactic parsing, Stanford basic dependencies (SD) representation (de Marneffe et al., 2006) converted by the Stanford parser<sup>10</sup> for dependency syntactic parsing using the same model of (Zhou et al., 2020) to fine-tune. We follow standard data splitting and evaluate setting as (Zhou et al., 2020) and use end-to-end SRL setups of both span and dependency SRL. Since LIMIT-BERT involves all syntactic and semantic parsing tasks, it is possible to directly apply LIMIT-BERT to each task without fine-tuning and we also compare these results.

In order to evaluate the language model pre-training performance of our LIMIT-BERT, we also evaluate LIMIT-BERT on two widely-used datasets, The General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2018) which is a collection of nine NLU tasks and Stanford Natural Language Inference (SNLI) (Bowman et al., 2015) to show the superiority.

### 4.2 Implementation Details

Our implementation of LIMIT-BERT is based on the PyTorch implementation of BERT<sup>11</sup>. We use a learning rate of 3e-5 and a batch size of 32 with 1 million training steps. The optimizer and other training settings are the same as BERT (Devlin et al., 2019). For task-specific layers including syntactic and semantic scorers and decoders, we set the same hyperparameters settings as (Zhou et al., 2020). LIMIT-BERT model is trained on 32 NVIDIA GeForce GTX 1080Ti GPUs.

### 4.3 Main Results

**Syntactic Parsing Results** As shown in Table 1 and 2, LIMIT-BERT without fine-tuning obtains 95.84 F1 score of constituent parsing and 97.14% UAS and 95.44% LAS of dependency parsing. Compared with baseline BERT<sub>WWM</sub>, LIMIT-BERT

<sup>10</sup><http://nlp.stanford.edu/software/lex-parser.html>

<sup>11</sup><https://github.com/huggingface/pytorch-pretrained-BERT>. We use Whole Word Masking BERT parameters as our Transformer encoder initialization.

	UAS	LAS
Dozat and Manning (2017)	95.74	94.08
Ma et al. (2018)	95.87	94.19
Ji et al. (2019)	95.97	94.31
Fernández-González and Gómez-Rodríguez (2019)	96.04	94.43
Liu et al. (2019a)	96.09	95.03
Zhou and Zhao (2019)(BERT)	97.00	95.43
Zhou et al. (2020)(BERT)	96.90	95.32
Zhou et al. (2020)(XLNet)	97.23	95.65
Baseline (BERT <sub>WWM</sub> )	96.89	95.22
<b>Our LIMIT-BERT</b>	96.94	95.30
<b>Our LIMIT-BERT</b> <sup>†</sup>	97.14	95.44

Table 1: Dependency syntactic parsing on PTB, no finetuning result is marked by †.

	LR	LP	F1
Gaddy et al. (2018)	91.76	92.41	92.08
Kitaev and Klein (2018)(ELMo)	94.85	95.40	95.13
Kitaev et al. (2019)(BERT)	95.46	95.73	95.59
Zhou and Zhao (2019)(BERT)	95.70	95.98	95.84
Zhou et al. (2020)(BERT)	95.39	95.64	95.52
Zhou et al. (2020)(XLNet)	96.10	96.26	96.18
Baseline (BERT <sub>WWM</sub> )	95.59	95.86	95.72
<b>Our LIMIT-BERT</b>	95.67	95.92	95.80
<b>Our LIMIT-BERT</b> <sup>†</sup>	95.72	95.96	95.84

Table 2: Constituent syntactic parsing on PTB, no fine-tuning result is marked by †.

System	WSJ			Brown		
	P	R	F <sub>1</sub>	P	R	F <sub>1</sub>
<i>End-to-end Span SRL</i>						
He et al. (2018a)	81.2	83.9	82.5	69.7	71.9	70.8
He et al. (2018a)(ELMo)	84.8	87.2	86.0	73.9	78.4	76.1
Li et al. (2019)(ELMo)	85.2	87.5	86.3	74.7	78.1	76.4
Strubell et al. (2018)(ELMo)	87.13	86.67	86.90	79.02	77.49	78.25
Zhou et al. (2020)(BERT)	86.46	88.23	87.34	77.26	80.20	78.70
Zhou et al. (2020)(XLNet)	87.48	89.51	88.48	80.46	84.15	82.26
Baseline (BERT <sub>WWM</sub> )	86.48	88.59	87.52	79.4	82.68	81.01
<b>Ours LIMIT-BERT</b>	86.62	89.12	87.85	79.58	83.05	81.28
<b>Ours LIMIT-BERT</b> <sup>†</sup>	87.16	88.51	87.83	79.20	80.29	79.74
<i>End-to-end Dependency SRL</i>						
Li et al. (2019)	-	-	85.1	-	-	-
He et al. (2018b)	83.9	82.7	83.3	-	-	-
Cai et al. (2018)	84.7	85.2	85.0	-	-	72.5
Li et al. (2019)(ELMo)	84.5	86.1	85.3	74.6	73.8	74.2
Zhou et al. (2020)(BERT)	86.77	89.14	87.94	79.71	82.40	81.03
Zhou et al. (2020)(XLNet)	86.35	90.16	88.21	80.90	85.38	83.08
Baseline (BERT <sub>WWM</sub> )	85.13	89.21	87.12	79.05	83.95	81.43
<b>Ours LIMIT-BERT</b>	85.84	90.01	87.87	79.50	84.85	82.09
<b>Ours LIMIT-BERT</b> <sup>†</sup>	85.73	89.34	87.50	79.60	82.81	81.17

Table 3: Span SRL and dependency SRL results on on CoNLL-2005 and CoNLL-2009 test sets in end-to-end mode, no finetuning result is marked by †.

outperforms the baseline model both of fine-tuning or not. Particularly, LIMIT-BERT without fine-tuning exceeds more than 0.2 in UAS of dependency and 0.1 F1 of constituent syntactic parsing which are considerable improvements on such strong baselines.

**Semantic Parsing Results** Table 3 shows results on CoNLL-2005, CoNLL-2009 in-domain (WSJ) and out-domain (Brown) test sets and compares our LIMIT-BERT with previous published

Model	Dev	Test
DRCN (Kim et al., 2018)	-	90.1
SJRC (Zhang et al., 2019)	-	91.3
MT-DNN (Liu et al., 2019b)	92.2	91.6
SemBERT (Zhang et al., 2020a)	-	91.9
Baseline (BERT <sub>WWM</sub> )	91.7	91.4
<b>LIMIT-BERT</b>	92.3	91.7

Table 4: Leaderboards of SNLI dataset. Both our LIMIT-BERT and BERT<sub>WWM</sub> are single models.

state-of-the-art models in end-to-end mode. The upper part of the table presents results from span SRL while the lower part shows results of dependency SRL. Compared with baseline, LIMIT-BERT with fine-tuning outperforms BERT<sub>WWM</sub> on all four SRL datasets, exceeding more than 0.3 in F1 of in-domain span SRL and 0.7 F1 of dependency SRL, which demonstrate that LIMIT-BERT can furthermore improve SRL performance even over strong baselines.

The results of syntactic and semantic parsing empirically illustrate that incorporating linguistic knowledge into pre-trained language model by multi-task and semi-supervised learning can significantly enhance downstream tasks.

**SNLI Results** Table 4 includes the best results reported in the leaderboards<sup>12</sup> of SNLI. We see that LIMIT-BERT outperforms the strong baseline model BERT<sub>WWM</sub> in 0.3 F1 score on the SNLI benchmark.

**GLUE Results** We fine-tuned LIMIT-BERT for each GLUE task on task-specific data. The dev results in Table 5 show that LIMIT-BERT outperforms the strong baseline model and achieves remarkable results compared to other state-of-the-art models in literature.

## 4.4 Discussions

**Ablation Study** LIMIT-BERT contains three key components: Multi-Task learning, ELECTRA training approach, and Syntactic/Semantic Phrase Masking (SPM). To evaluate the contribution of each component in LIMIT-BERT, we remove each component from the model for training and then fine-tune on downstream NLU tasks and linguistics tasks for evaluation. In consideration of computational cost, we apply BERT<sub>base</sub> as the start of training and only use one-tenth of the BERT training corpus. We employ the same training setting

<sup>12</sup><https://nlp.stanford.edu/projects/snli/>

Model	CoLA (mc)	SST-2 (acc)	MRPC (F1/acc)	STS-B (pc/sc)	QQP (acc/F1)	MNLI m/mm(acc)	QNLI (acc)	RTE (acc)	Score -
<i>Dev set results for Comparison</i>									
BERT	60.6	93.2	-/88.0	-/90.0	91.3/-	-/86.6	92.3	70.4	84.0
MT-DNN	63.5	94.3	91.0/87.5	90.7/90.6	91.9/89.2	87.1/86.7	92.9	83.4	-
ELECTRA	69.3	96.0	-/90.6	-/92.1	92.4/-	-/90.5	94.5	86.8	89.0
Baseline (BERT <sub>WWM</sub> )	63.6	93.6	90.8/87.0	90.5/90.2	91.7/88.8	87.4/87.2	93.9	77.3	85.6
<b>LIMIT-BERT</b>	64.0	94.0	94.0/91.7	91.5/91.3	91.6/88.6	87.4/87.3	93.5	85.2	87.3
<i>Test set results for models with standard single-task finetuning</i>									
BiLSTM+ELMo+Attn	36.0	90.4	84.9/77.9	75.1/73.3	84.7/64.8	76.4/76.1	-	56.8	70.5
BERT	60.5	94.9	89.3/85.4	87.6/86.5	89.3/72.1	86.7/85.9	92.7	70.1	80.5
MT-DNN	62.5	95.6	91.1/88.2	89.5/88.8	89.6/72.7	86.7/86.0	93.1	81.4	82.7
SemBERT	62.3	94.6	91.2/88.3	87.8/86.7	89.8/72.8	87.6/86.3	94.6	84.5	82.9
ELECTRA	71.7	97.1	93.1/90.7	92.9/92.5	90.8/75.6	91.3/90.8	95.8	89.8	89.35
<b>LIMIT-BERT</b>	62.5	94.5	90.9/88.0	90.3/89.7	89.5/71.9	87.1/86.2	94.0	83.0	83.3

Table 5: Comparison of GLUE dev and test sets. Our model is in boldface. MT-DNN dev results are from (Liu et al., 2019b) and other dev results are from (Clark et al., 2020).

System	GLUE	SNLI	SNLI
	Dev	Dev	Test
<b>LIMIT-BERT</b>	82.6	<b>90.6</b>	<b>91.0</b>
w/o Multi-Task	<b>82.9</b>	90.5	90.7
w/o ELECTRA	81.3	90.4	90.5
w/o SPM	80.2	90.6	90.8

Table 6: Ablation study of LIMIT-BERT (base) on GLUE and SNLI.

System	SEM <sub>span</sub>	SEM <sub>dep</sub>	SYN <sub>con</sub>	SYN <sub>dep</sub>	
	F <sub>1</sub>	F <sub>1</sub>	F <sub>1</sub>	UAS	LAS
<b>LIMIT-BERT</b>	<b>86.25</b>	<b>85.74</b>	95.34	96.59	<b>94.71</b>
w/o Multi-Task	85.60	85.24	95.16	96.38	94.38
w/o ELECTRA	86.20	85.71	95.32	96.59	94.70
w/o SPM	86.21	85.72	<b>95.44</b>	<b>96.64</b>	94.70

Table 7: Ablation study of LIMIT-BERT (base) on linguistics tasks.

for each ablation model: 200k training steps, 1e-5 learning rate and 32 batch size. After language model training, we extract the layers of BERT<sub>base</sub> and fine-tune on downstream tasks for evaluation.

The ablation study is conducted on NLU tasks and linguistics tasks shown in Table 6 and 7 respectively. For NLU tasks, GLUE and SNLI results both decrease in w/o ELECTRA and w/o SPM setting showing the effectiveness of the ELECTRA training approach and SPM for NLU tasks. For linguistics tasks, w/o Multi-Task setting hurts the performance obviously from LIMIT-BERT both of semantic and syntactic parsing, which shows the effectiveness of multi-tasks learning for linguistics tasks. Besides, the ELECTRA training approach

System	SEM <sub>span</sub>	SEM <sub>dep</sub>	SYN <sub>con</sub>	SYN <sub>dep</sub>	
	F <sub>1</sub>	F <sub>1</sub>	F <sub>1</sub>	UAS	LAS
Baseline	86.55	86.10	95.52	96.54	94.71
LIMIT-BERT	<b>87.13</b>	<b>86.77</b>	95.55	96.54	94.74
LIMIT-BERT <sup>†</sup>	87.04	86.38	<b>95.72</b>	<b>96.82</b>	<b>94.82</b>

Table 8: Fine-tuning effect analysis on English dev sets, no finetuning result is marked by †.

and SPM also can improve performance when fine-tuning on linguistics tasks.

Comparing the results in Tables 6 and 7, ELECTRA training approach and SPM are more effective for NLU tasks while multi-tasks learning can improve the linguistics tasks performance significantly. The possible explanation is that multi-tasks learning enables LIMIT-BERT to ‘remember’ the linguistics information and thus lead to better performance in downstream linguistics tasks.

**Fine-tuning Effect** We examine the fine-tuning effect of LIMIT-BERT on linguistics tasks. The results in Table 8 show that LIMIT-BERT with or without finetuning still outperforms BERT<sub>WWM</sub> baseline consistently among all tasks. In such a case, fine-tuning is necessary to boost the semantic parsing performance while no-fine-tuning performs better on syntactic parsing. As shown in Table 8, the accuracy improves 0.1 F1 and 0.4 F1 of span SRL and dependency SRL after fine-tuning respectively but no-fine-tuning performs better nearly 0.2 F1 of syntactic parsing. The possible explanation is that no-fine-tuning LIMIT-BERT use semi-supervised training data which contains much more long sentence samples and benefits syntactic parsing more.

Model	Test
Yasunaga et al. (2018)	97.59
Akbik et al. (2018)	97.85
Bohnet et al. (2018)	97.96
LIMIT-BERT	97.71

Table 9: POS tagging accuracy on the WSJ test set, with other top-performing systems.




Figure 3: The performance of baseline model and LIMIT-BERT while varying sentence length of four linguistics tasks on the English dev set.

**Part-Of-Speech Performance** Table 9 lists the results of POS tagging on WSJ test set showing that our LIMIT-BERT achieves competitive results compared with other state-of-the-art models. Note that we only apply simple one-layer decoder without those complicated ones such as conditional random field (CRF) (Lafferty et al., 2001) as the POS tagging task is not the main concern of our model.

**Sentences Length Performance** The performance of baseline model and LIMIT-BERT while varying the sentence length of four linguistics tasks on the English dev set is shown in Figure 3. The statistics show that our LIMIT-BERT outperforms the baseline model of over all sentence lengths. For different sentence lengths, LIMIT-BERT outperforms much better than baseline model on long sentence (larger than 50) of both syntactic and semantic parsing. The possible explanation is that LIMIT-

BERT uses semi-supervised training data, which contains much more long sentence samples and benefits parsing performance on long sentences.

## 5 Related Work

**Linguistics Inspired NLP Tasks** With the impressive success of deep neural networks in various NLP tasks (Chen and Manning, 2014; Dozat and Manning, 2017; Ma et al., 2018; Strubell et al., 2018; Luo and Zhao, 2020; Li et al., 2020; He et al., 2019; Luo et al., 2020; Zhang et al., 2018a; Li et al., 2018a; Zhang et al., 2018b), syntactic parsing and semantic role labeling have been well developed with neural network and achieve very high performance (Chen and Manning, 2014; Dozat and Manning, 2017; Ma et al., 2018; Kitaev and Klein, 2018; Zhou and Zhao, 2019). Semantic role labeling is deeply related to syntactic structure and a number of works try to incorporate syntactic information in semantic role labeling models by different methods such as concatenation of lexicalized embedding (He et al., 2018b), usage of syntactic GCN (Li et al., 2018b) and multi-task learning (Strubell et al., 2018; Zhou et al., 2020). Besides semantic role labeling and syntactic parsing are two key tasks of semantics and syntax so that they are included into our linguistics tasks for multi-task learning.

In addition, both span and dependency are popularly adopted annotation styles for both semantics and syntax and some work on jointly learning of semantic and syntactic (Henderson et al., 2013; Lluís et al., 2013; Swayamdipta et al., 2016). Researchers are interested in two styles of SRL models that may benefit from each other rather than their separated development, which has been roughly discussed in (Johansson and Nugues, 2008). On the other hand, researchers have discussed how to encode lexical dependencies in phrase structures, like lexicalized tree adjoining grammar (LTAG) (Schabes et al., 1988), Combinatory Categorical Grammar (CCG) (Steedman, 2000) and head-driven phrase structure grammar (HPSG) (Pollard and Sag, 1994) which is a constraint-based highly lexicalized non-derivational generative grammar framework. To absorb both strengths of span and dependency structure, we apply both span (constituent) and dependency representations of semantic role labeling and syntactic parsing. Thus, it is a natural idea to study the relationship between constituent and dependency structures, and



the joint learning of constituent and dependency syntactic parsing (Klein and Manning, 2004; Charniak and Johnson, 2005; Farkas et al., 2011; Green and Žabokrtský, 2012; Ren et al., 2013; Xu et al., 2014; Yoshikawa et al., 2017).

**Pre-trained Language Modeling** Recently, deep contextual language model has been shown effective for learning universal language representations by leveraging large amounts of unlabeled data, achieving various state-of-the-art results in a series of NLU benchmarks. Some prominent examples are Embedding from Language models (ELMo) (Peters et al., 2018), Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019) and Generalized Autoregressive Pretraining (XLNet) (Yang et al., 2019).

Many latest works make attempts to modify the language model based on BERT such as ELECTRA (Clark et al., 2020) and MT-DNN (Liu et al., 2019b). ELECTRA focuses on the [MASK] tokens mismatch problem and thus combines the idea of Generative Adversarial Networks GANs (Goodfellow et al.). MT-DNN applies multi-task learning to language model pre-training and achieves new state-of-the-art results on GLUE benchmark. Besides, (Gururangan et al., 2020) finds that multiphase adaptive pretraining offers large gains in task performance which is similar with our semi-supervised learning strategy.

## 6 Conclusions

In this work, we present LIMIT-BERT which applies multi-task learning with multiple linguistic tasks by semi-supervised learning. We use five key syntax and semantics tasks : Part-Of-Speech (POS) tags, constituent and dependency syntactic parsing, span and dependency semantic role labeling (SRL). and further improve the masking strategy of BERT training by effectively exploiting the available syntactic and semantic clues for language model training. The experiments show that LIMIT-BERT outperforms the strong baseline BERT<sub>wwm</sub> on four benchmark parsing treebanks and two NLU tasks. The results of GLUE and SNLI empirically illustrate that incorporating linguistic knowledge into pre-training language BERT by multi-task and semi-supervised learning can also enhance downstream tasks. There are many future areas to explore to improve LIMIT-BERT, including a deeper understanding of model structure sharing in MTL, a more effective training method that leverages relat-

edness among multiple tasks, for both fine-tuning and pre-training, and ways of incorporating the linguistic structure of text in a more explicit and controllable manner.

## References

- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. [Contextual string embeddings for sequence labeling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Bernd Bohnet, Ryan McDonald, Gonçalo Simões, Daniel Andor, Emily Pitler, and Joshua Maynez. 2018. [Morphosyntactic tagging with a meta-BiLSTM model over context sensitive token encodings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2642–2652, Melbourne, Australia. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. [A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware?](#) In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2753–2765, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Xavier Carreras and Lluís Màrquez. 2005. [Introduction to the CoNLL-2005 shared task: Semantic role labeling](#). In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 152–164, Ann Arbor, Michigan. Association for Computational Linguistics.
- Richard A Caruana. 1993. [Multitask Learning: A Knowledge-Based Source of Inductive Bias](#). In *Machine Learning Proceedings*.
- Eugene Charniak and Mark Johnson. 2005. [Coarse-to-Fine n-Best Parsing and MaxEnt Discriminative Reranking](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Danqi Chen and Christopher Manning. 2014. [A Fast and Accurate Dependency Parser using Neural Networks](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- N. Chomsky. 1981. *Lectures on Government and Binding*. Mouton de Gruyter.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [{ELECTRA}: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- 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 (NAACL:HLT)*.
- Timothy Dozat and Christopher D. Manning. 2017. Deep Biaffine Attention for Neural Dependency Parsing. In *International Conference on Learning Representations 2017 (ICLR)*.
- Richàrd Farkas, Bernd Bohnet, and Helmut Schmid. 2011. Features for Phrase-Structure Reranking from Dependency Parses. In *Proceedings of the 12th International Conference on Parsing Technologies*.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2019. [Left-to-Right Dependency Parsing with Pointer Networks](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*.
- David Gaddy, Mitchell Stern, and Dan Klein. 2018. What’s Going On in Neural Constituency Parsers? An Analysis. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL: HLT)*.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems (NIPS)*.
- Nathan Green and Zdeněk Žabokrtský. 2012. Hybrid Combination of Constituency and Dependency Trees into an Ensemble Dependency Parser. In *Proceedings of the Workshop on Innovative Hybrid Approaches to the Processing of Textual Data*.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t Stop Pretraining: Adapt Language Models to Domains and Tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jan Hajič, Massimiliano Ciaramita, Richard Johansson, Daisuke Kawahara, Maria Antònia Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Štěpánek, Pavel Straňák, Mihai Surdeanu, Nianwen Xue, and Yi Zhang. 2009. [The CoNLL-2009 shared task: Syntactic and semantic dependencies in multiple languages](#). In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009): Shared Task*, pages 1–18, Boulder, Colorado. Association for Computational Linguistics.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018a. [Jointly predicting predicates and arguments in neural semantic role labeling](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 364–369, Melbourne, Australia. Association for Computational Linguistics.
- Shexia He, Zuchao Li, and Hai Zhao. 2019. [Syntax-aware Multilingual Semantic Role Labeling](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018b. [Syntax for semantic role labeling, to be, or not to be](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2061–2071, Melbourne, Australia. Association for Computational Linguistics.
- James Henderson, Paola Merlo, Ivan Titov, and Gabriele Musillo. 2013. [Multilingual joint parsing of syntactic and semantic dependencies with a latent variable model](#). *Computational Linguistics*, 39(4):949–998.
- Tao Ji, Yuanbin Wu, and Man Lan. 2019. [Graph-based dependency parsing with graph neural networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2475–2485, Florence, Italy. Association for Computational Linguistics.
- Richard Johansson and Pierre Nugues. 2008. [Dependency-based semantic role labeling of PropBank](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 69–78, Honolulu, Hawaii. Association for Computational Linguistics.
- Seonhoon Kim, Jin-Hyuk Hong, Inho Kang, and Nojun Kwak. 2018. Semantic sentence matching with densely-connected recurrent and co-attentive information. *arXiv preprint arXiv:1805.11360*.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. [Multilingual Constituency Parsing with Self-Attention and Pre-Training](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Nikita Kitaev and Dan Klein. 2018. Constituency Parsing with a Self-Attentive Encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.

- Dan Klein and Christopher Manning. 2004. [Corpus-Based Induction of Syntactic Structure: Models of Dependency and Constituency](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- 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)*, pages 282–289.
- Zuchao Li, Jiaxun Cai, Shexia He, and Hai Zhao. 2018a. [Seq2seq dependency parsing](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3203–3214, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Zuchao Li, Shexia He, Jiaxun Cai, Zhuosheng Zhang, Hai Zhao, Gongshen Liu, Linlin Li, and Luo Si. 2018b. [A unified syntax-aware framework for semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2401–2411, Brussels, Belgium. Association for Computational Linguistics.
- Zuchao Li, Shexia He, Hai Zhao, Yiqing Zhang, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2019. [Dependency or span, end-to-end uniform semantic role labeling](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*.
- Zuchao Li, Rui Wang, Kehai Chen, Masso Utiyama, Eiichiro Sumita, Zhuosheng Zhang, and Hai Zhao. 2020. [Data-dependent gaussian prior objective for language generation](#). In *International Conference on Learning Representations*.
- Linlin Liu, Xiang Lin, Shafiq Joty, Simeng Han, and Lidong Bing. 2019a. [Hierarchical Pointer Net Parsing](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019b. [Multi-task deep neural networks for natural language understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- Xavier Lluís, Xavier Carreras, and Lluís Màrquez. 2013. [Joint arc-factored parsing of syntactic and semantic dependencies](#). *Transactions of the Association for Computational Linguistics*, 1:219–230.
- Ying Luo and Hai Zhao. 2020. [Bipartite Flat-Graph Network for Nested Named Entity Recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ying Luo, Hai Zhao, and Junlang Zhan. 2020. [Named Entity Recognition Only from Word Embeddings](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Xuezhe Ma, Zecong Hu, Jingzhou Liu, Nanyun Peng, Graham Neubig, and Eduard Hovy. 2018. [Stack-Pointer Networks for Dependency Parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a Large Annotated Corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2).
- Marie-Catherine de Marneffe, Bill MacCartney, and Christopher D. Manning. 2006. [Generating Typed Dependency Parses from Phrase Structure Parses](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*.
- Todor Mihaylov and Anette Frank. 2019. [Discourse-Aware Semantic Self-Attention for Narrative Reading Comprehension](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Hiroki Ouchi, Hiroyuki Shindo, and Yuji Matsumoto. 2018. [A span selection model for semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1630–1642, Brussels, Belgium. Association for Computational Linguistics.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The proposition bank: An annotated corpus of semantic roles](#). *Computational Linguistics*, 31(1):71–106.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep Contextualized Word Representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL: HLT)*.
- Carl Pollard and Ivan A Sag. 1994. *Head-Driven Phrase Structure Grammar*. University of Chicago Press.
- Vasin Punyakanok, Dan Roth, and Wen-tau Yih. 2008. [The importance of syntactic parsing and inference in semantic role labeling](#). *Computational Linguistics*, 34(2).
- Xiaona Ren, Xiao Chen, and Chunyu Kit. 2013. [Combine Constituent and Dependency Parsing via Reranking](#). In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*.



- Yves Schabes, Anne Abeille, and Aravind K. Joshi. 1988. Parsing strategies with 'lexicalized' grammars: Application to tree adjoining grammars. In *Coling Budapest 1988 Volume 2: International Conference on Computational Linguistics (COLING)*.
- Mark Steedman. 2000. *The Syntactic Process*. MIT Press, Cambridge, MA, USA.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038, Brussels, Belgium. Association for Computational Linguistics.
- Swabha Swayamdipta, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. Greedy, joint syntactic-semantic parsing with stack LSTMs. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 187–197, Berlin, Germany. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems (NIPS)*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *CoRR*, abs/1609.08144.
- Wenduan Xu, Stephen Clark, and Yue Zhang. 2014. Shift-Reduce CCG Parsing with a Dependency Model. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Thirty-Third Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- Michihiro Yasunaga, Jungo Kasai, and Dragomir Radev. 2018. Robust multilingual part-of-speech tagging via adversarial training. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 976–986, New Orleans, Louisiana. Association for Computational Linguistics.
- Masashi Yoshikawa, Hiroshi Noji, and Yuji Matsumoto. 2017. A\* CCG Parsing with a Supertag and Dependency Factored Model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Zhisong Zhang, Rui Wang, Masao Utiyama, Eiichiro Sumita, and Hai Zhao. 2018a. Exploring recombination for efficient decoding of neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4785–4790, Brussels, Belgium. Association for Computational Linguistics.
- Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. 2018b. Modeling multi-turn conversation with deep utterance aggregation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3740–3752, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Zhuosheng Zhang, Yuwei Wu, Zuchao Li, and Hai Zhao. 2019. Explicit contextual semantics for text comprehension. In *The 33rd Pacific Asia Conference on Language, Information and Computation (PACLIC)*.
- Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020a. Semantics-aware BERT for Language Understanding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhuosheng Zhang, Yuwei Wu, Junru Zhou, Sufeng Duan, Hai Zhao, and Rui Wang. 2020b. SG-Net: Syntax-Guided Machine Reading Comprehension. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*.
- Junru Zhou, Zuchao Li, and Hai Zhao. 2020. Parsing All: Syntax and Semantics, Dependencies and Spans. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Junru Zhou and Hai Zhao. 2019. Head-driven phrase structure grammar parsing on Penn treebank. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2396–2408, Florence, Italy. Association for Computational Linguistics.