

Multilingual Chart-based Constituency Parse Extraction from Pre-trained Language Models

Taeuk Kim*

Dept. of Computer Science
Dept. of Artificial Intelligence
Hanyang University
kimtaeuk@hanyang.ac.kr

Bowen Li

ILCC
School of Informatics
University of Edinburgh
bowen.li@ed.ac.uk

Sang-goo Lee

Dept. of Computer Science
and Engineering
Seoul National University
sglee@europa.snu.ac.kr

Abstract

As it has been unveiled that pre-trained language models (PLMs) are to some extent capable of recognizing syntactic concepts in natural language, much effort has been made to develop a method for extracting complete (binary) parses from PLMs without training separate parsers. We improve upon this paradigm by proposing a novel chart-based method and an effective top-K ensemble technique. Moreover, we demonstrate that we can broaden the scope of application of the approach into multilingual settings. Specifically, we show that by applying our method on multilingual PLMs, it becomes possible to induce non-trivial parses for sentences from nine languages in an integrated and language-agnostic manner, attaining performance superior or comparable to that of unsupervised PCFGs. We also verify that our approach is robust to cross-lingual transfer. Finally, we provide analyses on the inner workings of our method. For instance, we discover universal attention heads which are consistently sensitive to syntactic information irrespective of the input language.

1 Introduction

Constituency parsing is a classic task in natural language processing (NLP), whose goal is to construct a phrase-structure tree for a given sentence. As parse trees have long been recognized as being integral to the meaning of sentences, there has been an enormous amount of work in the literature to develop constituency parsers (Charniak (2000); Collins (2003); Petrov et al. (2006); *inter alia*), resulting in the development of sophisticated neural supervised parsers (Kitaev and Klein, 2018; Kitaev et al., 2019; Zhou and Zhao, 2019). Although it becomes possible to obtain such neural parsers of high-quality for a few languages, there remain many other languages which, for lack of resources or attention, have yet to benefit from the progress

made in the field of constituency parsing. The main issue is that it is expensive and time-consuming to prepare adequate numbers of gold-standard trees essential for training parsers with supervision.

Considering this data scarcity problem, *unsupervised* constituency parsing methods have naturally arisen as an alternative for generating constituency trees. Work on unsupervised parsing (Shen et al. (2018b, 2019); Kim et al. (2019a,b); *inter alia*) has focused on devising linguistically-informed neural models, which are carefully designed to be more sensitive to the hierarchical nature of language structure and be able to learn this nature from raw text rather than gold-standard trees. Despite the recent progress in implementing *English* unsupervised parsers with decent performance, it has been known as not trivial to reproduce such a success in multilingual environments (Kann et al., 2019; Zhao and Titov, 2021). Moreover, as most unsupervised parsers are typically trained on monolingual text for a specific target language, it is required to prepare separate instances of them to support different languages.

On the other hand, a different line of work (Kim et al., 2020; Wu et al., 2020) has proposed a new direction of inducing syntax trees, dubbed in this paper as **Constituency Parse Extraction from Pre-trained Language Models (CPE-PLM)**, by relying on the combination of (i) simple distance metrics and (ii) the representations obtained from pre-trained language models (PLMs).¹ The core assumption underlying the methodology is that PLMs hold enough syntactic knowledge to be utilized for predicting parse trees by themselves. Although the CPE-PLM framework has demonstrated that non-trivial trees resembling gold-standard annotations can be extracted from general PLMs even without fine-

*This work was mainly conducted when TK was at SNU.

¹We use the term *pre-trained language models (PLMs)* to refer to BERT-like (Devlin et al., 2019) Transformer (Vaswani et al., 2017) models that are pre-trained with massive plain text corpora in a self-supervised fashion.

tuning on treebanks, it has been also reported that CPE-PLM’s freedom from task-specific training comes at the cost of its performance inferior to that of unsupervised parsers. Furthermore, as is the case of the foregoing unsupervised parsing literature, it has yet to be verified that CPE-PLM is also effective for languages other than English.

In this paper, we first attempt to narrow the performance gap between unsupervised parsers and CPE-PLM by introducing a novel method inspired by neural chart-based algorithms (Durrett and Klein, 2015; Stern et al., 2017; Kitaev and Klein, 2018). In contrast to the top-down CPE-PLM method (Kim et al., 2020), which focuses on detecting the boundary of two subspans in a phrase relying only on the knowledge from the two words around the boundary, our chart-based method considers all components in a phrase to judge how plausible the phrase is. Furthermore, we introduce a simple but effective ensemble technique that utilizes the pre-defined set of attention heads which are confirmed as being effective by their performance on the validation set. We show that our chart-based method outperforms or is competitive to the top-down method and that the top-K ensemble plays a key role in boosting their performance.

Second, the limitation of most previous studies for both unsupervised parsing and CPE-PLM is that they are heavily English-centric, leaving an open question whether they are universally applicable. To investigate this problem, we test CPE-PLM on several other languages. Specifically, we propose to introduce *multilingual* PLMs (Conneau and Lample, 2019; Conneau et al., 2019) into CPE-PLM to grant the framework an ability to deal with multiple languages simultaneously. We show that the CPE-PLM methods built upon multilingual PLMs are able to induce reasonable parses for sentences in nine languages in an integrated and language-agnostic manner, achieving figures superior or comparable to ones from neural PCFGs (Kim et al., 2019a; Zhao and Titov, 2021). In supplementary analyses, we provide intuitive explanations about the inner workings of our method. For instance, we confirm the existence of *universal* attention heads which seem to be responsible for capturing syntactic information irrespective of the input language.

2 Background

In this work, we focus on a variant of unsupervised constituency parsing, which we call **Constituency**

Methodology	Unsupervised Parsing	CPE-PLM
Training data	In-domain data (e.g., raw text from PTB)	General corpora (e.g., Wikipedia)
Architecture	Task-oriented (e.g., RNN, PCFG)	Transformer
Modeling	$p(S, T)$ (T is marginalized or implicitly modeled)	$p(S)$ (T is not considered in modeling)

Table 1: Comparison between typical unsupervised parsing and constituency parse extraction from pre-trained language models (CPE-PLM).

Parse Extraction from Pre-trained Language Models (CPE-PLM) (Kim et al., 2020; Wu et al., 2020). We specify the characteristics of CPE-PLM in Table 1, comparing them with those of general unsupervised parsing methods.

Typical unsupervised parsers consist of task-oriented architectures (e.g., RNN (Kim et al., 2019b) and PCFG (Kim et al., 2019a)) which are designed to model both a sentence S and the corresponding tree T (i.e., $p(S, T)$) and are trained with in-domain plain text.² On the other hand, CPE-PLM simply employs off-the-shelf Transformer PLMs, which only model the probability of a sentence $p(S)$, as their core component and do not require additional training—the PLMs are frozen and no trainable component is augmented on top of them, meaning *parameter-free*. Instead, CPE-PLM methods take advantage of implicit syntactic knowledge residing in PLMs to reconstruct parses, by computing syntactic distances (Shen et al., 2018b) between words in a sentence using features from the PLMs. We describe their algorithmic details in Section 3. CPE-PLM’s independence from training also makes it being distinct from syntactic probes for PLMs (Hewitt and Manning, 2019; Chi et al., 2020) which demand training probing modules to investigate the latent knowledge of PLMs.

3 Method

Among various approaches that belong to CPE-PLM, we regard the top-down method proposed by Kim et al. (2020) as a starting point and aim to improve the method in several perspectives.

3.1 Top-down CPE-PLM

Kim et al. (2020) proposed a zero-shot version of top-down constituency parsing (Shen et al.,

²As it is mostly infeasible to directly model the tree T without supervision from gold annotations, unsupervised parsers usually make use of different approximation or marginalization techniques such as variational inference and sampling.

Algorithm 1 Syntactic Distance to Binary Constituency Tree (from Kim et al. (2020))

```

1:  $S = [w_1, w_2, \dots, w_n]$ : Words in a sentence of length  $n$ .
2:  $\mathbf{d} = [d_1, d_2, \dots, d_{n-1}]$ : Vector, each of whose elements  $d_i$  is the syntactic distance between  $w_i$  and  $w_{i+1}$ .
3: function D2T( $S, \mathbf{d}$ )
4:   if  $\mathbf{d} = []$  then
5:     node  $\leftarrow$  Leaf( $S[0]$ )
6:   else
7:      $i \leftarrow \arg \max_i(\mathbf{d})$ 
8:     child $_l \leftarrow$  D2T( $S_{\leq i}, \mathbf{d}_{< i}$ )
9:     child $_r \leftarrow$  D2T( $S_{> i}, \mathbf{d}_{> i}$ )
10:    node  $\leftarrow$  Node(child $_l, \text{child}_r$ )
11:   end if
12:   return node
13: end function

```

2018a), where a concept of *syntactic distance* (Shen et al., 2018b) plays a vital role. Formally, given a sequence of words in a sentence $S = [w_1, w_2, \dots, w_n]$, the corresponding syntactic distance vector $\mathbf{d} = [d_1, d_2, \dots, d_{n-1}]$ is computed as follows (each d_i is the syntactic distance between w_i and w_{i+1}):

$$d_i = f(g(w_i), g(w_{i+1})),$$

where $f(\cdot, \cdot)$ and $g(\cdot)$ are a distance measure function and representation extractor function. For g , the authors suggest utilizing G^v and G^d . Given l as the number of layers in a PLM and a as the number of attention heads per layer, G^v refers to a set of functions $\{g_j^v | j = 1, \dots, l\}$, each of which outputs the hidden representation of a given word on the j^{th} layer of the PLM. Similarly, G^d is defined as $\{g_{(j,k)}^d | j = 1, \dots, l, k = 1, \dots, a + 1\}$ ³, each of whose elements computes the attention distribution of an input word by using the k^{th} attention head on the j^{th} layer of the PLM. For the function f , there also exist two options, i.e., $F^v = \{\text{COS}, \text{L1}, \text{L2}\}$ and $F^d = \{\text{JSD}, \text{HEL}\}$, where COS, L1, L2, JSD, and HEL correspond to the Cosine, L1, and L2, Jensen-Shannon, and Hellinger distance respectively. Note that F^v is only compatible with G^v while F^d is only with G^d .

Finally, given the input sentence S and syntactic distance vector \mathbf{d} , Algorithm 1 is adopted to induce a complete (binary) constituency parse tree, recursively splitting S in a top-down manner.

3.2 Chart-based CPE-PLM

Although the top-down method has shown its effectiveness in extracting non-trivial phrase structures

³Given a attention heads on the j^{th} layer, Kim et al. (2020) also consider the $(a + 1)^{\text{th}}$ head that corresponds to the average of all attention distributions on the j^{th} layer.

Algorithm 2 Chart to Syntactic Distance

```

1:  $n$ : Length of an input sentence  $S$ .
2:  $C \in \mathbb{R}^{n \times n}$ : Chart matrix whose elements are  $s_{span}(i, j)$ .
3:  $P \in \mathbb{R}^{n \times n}$ : Matrix, whose  $(i, j)^{\text{th}}$  element is the split point of the span  $(i, j)$  of the sentence  $S$ .
4:  $s$ : Start position, initialized as 1.
5:  $e$ : End position, initialized as  $n$ .
6: function C2D( $C, P, s, e$ )
7:   if  $s = e$  then
8:     return [] (empty vector)
9:   else
10:     $v \leftarrow C[s][e]$ 
11:     $p \leftarrow P[s][e]$ 
12:    return [C2D( $C, P, s, p$ );  $v$ ; C2D( $C, P, p+1, e$ )]
13:    ( $[\cdot; \cdot]$ : vector concatenation)
14:   end if
15: end function

```

from PLMs, there still remains much room for improvement, considering that this method by nature operates in a greedy fashion rather than taking account of the probabilities of all possible subtrees. In other words, the top-down CPE-PLM method only relies on the information obtained from the representations of two words to estimate the likelihood of the space between the two words becoming a target to be split. To overcome this limitation, we propose a novel approach based on chart parsing which executes an exact inference to find the most probable parse while effectively considering all possibilities with dynamic programming.

Following the previous work on chart parsing (Stern et al., 2017; Kitaev and Klein, 2018), we assign a real-valued score $s_{tree}(T)$ for each tree candidate T , which decomposes as $s_{tree}(T) = \sum_{(i,j) \in T} s_{span}(i, j)$, where $s_{span}(i, j)$ is a score (or cost) for a constituent that is located between positions i and j ($1 \leq i \leq j \leq n$) in a sentence. Specifically, $s_{span}(i, j)$ is defined as follows:

$$s_{span}(i, j) = \begin{cases} s_{comp}(i, j) + \min_{i \leq k < j} s_{split}(i, k, j) & \text{if } i < j \\ 0 & \text{if } i = j \end{cases}$$

where $s_{split}(i, k, j) = s_{span}(i, k) + s_{span}(k + 1, j)$. In other words, $s_{comp}(i, j)$ measures the validity or compositionality of the span (i, j) itself while $s_{split}(i, k, j)$ indicates how plausible it is to divide the span (i, j) into two subspans (i, k) and $(k + 1, j)$. We choose the most probable k that brings us the minimum cost of $s_{split}(i, k, j)$. Note that each constituent is by definition evaluated with the scores of its children in addition to its own score. Once $s_{comp}(\cdot, \cdot)$ is properly defined, it is straightforward to compute every $s_{span}(i, j)$ by utilizing the CKY algorithm (Cocke, 1969; Kasami, 1966; Younger, 1967). In the following subsections, we

formulate two variants of $s_{comp}(\cdot, \cdot)$ in detail.

Finally, our parser outputs \hat{T} , the tree that requires the lowest score (cost) to build, as a prediction for the parse tree of the input sentence: $\hat{T} = \arg \min_T s_{tree}(T)$.

3.2.1 Pair Score Function

The methodology introduced in the previous section abstracted over the choice of $s_{comp}(\cdot, \cdot)$; in what follows we propose two candidates for it.

First, we propose a pair score function $s_p(\cdot, \cdot)$ which is defined as follows:

$$s_p(i, j) := \binom{j-i+1}{2}^{-1} \sum_{(w_x, w_y) \in \text{pair}(i, j)} f(g(w_x), g(w_y)),$$

where $\text{pair}(i, j)$ returns a set consisting of all combinations of two words from a span (i, j) —e.g., $\text{pair}(1, 3) = \{(w_1, w_2), (w_1, w_3), (w_2, w_3)\}$. The intuition behind this formulation is that every pair of words in a constituent should have similar attention distributions so that the pair’s embeddings become similar to each other in the subsequent layers of PLMs. For $f(\cdot, \cdot)$ and $g(\cdot)$, we again take advantage of F^d and G^d specified in Section 3.1.⁴

3.2.2 Characteristic Score Function

We also propose another candidate for $s_{comp}(\cdot, \cdot)$, namely a characteristic score function $s_c(\cdot, \cdot)$. Instead of measuring the similarities of all pairs of attention distributions, we pre-define c as the characteristic vector of a given constituent and evaluate the cost of each word in the constituent with regard to this value. Although c can be realized in many ways, for simplicity, we here use the average of all the attention distributions of words in a constituent. As a consequence, $s_c(i, j)$ is formalized as follows:

$$s_c(i, j) := \frac{1}{j-i+1} \sum_{i \leq x \leq j} f(g(w_x), c),$$

where $c = \frac{1}{j-i+1} \sum_{i \leq y \leq j} g(w_y)$.

3.3 Top-K Ensemble for CPE-PLM

The part remaining ambiguous so far in clarifying CPE-PLM algorithms is about how to properly select the distance measure function f and representation extractor function g from the set of candidates, i.e., F^d and G^d . Basically, we can consider a typical case where we acquire the best combination of

⁴This implies that we make only use of the attention distributions of PLMs, as it is verified by the previous work (Kim et al., 2020) and our preliminary experiments that attention distributions offer more useful signals in this setting.

f and g using the validation set and apply it to the test set. In addition, we introduce one more option, called *top-K ensemble*, that enables us to integrate the knowledge from several attention heads.

Specifically, we first pick an arbitrary candidate for f , dubbed \hat{f} .⁵ Then, we compute the parsing performance of every possible combination of \hat{f} and $g \in G^d$ on the validation set and sort G^d according to the performance of its elements. After that, we simply choose the top K elements from the sorted G^d and allow all of them (G_{topK}^d) to participate in parsing instead of just leveraging the best single one. At test time, given an input sentence, we predict K separate trees using every element from G_{topK}^d , and then convert the trees into corresponding syntactic distance vectors (Algorithm 2). Finally, we compute the average of the syntactic distance vectors and translate this averaged vector into the final tree prediction (Algorithm 1).

By introducing the top-K ensemble technique, it becomes possible to obtain a more accurate tree prediction while seamlessly combining diverse syntactic signals provided by different attention heads.

4 Experiments

4.1 General Configurations

To evaluate, we prepare the PTB (Marcus et al., 1993) dataset for English and the SPMRL (Seddah et al., 2013) dataset for the eight other languages: Basque, French, German, Hebrew, Hungarian, Korean, Polish, and Swedish. We use the standard split of each dataset, and the datasets are pre-processed following Kim et al. (2019a) and Zhao and Titov (2021)—removing punctuation marks. We leverage the *unlabelled sentence-level F1* (percentage) score as a primary metric to evaluate the extent to which tree predictions resemble corresponding gold-standard trees. The hyperparameter K , which determines the number of attention heads engaging in the top-K ensemble, is decided by grid search on some reasonable candidates ($\{5, 10, 20, 30\}$). We empirically found that $K=20$ is versatile across different settings. From now on, we employ the abbreviations TD, CP, and CC to refer to the *top-down* (Section 3.1; Kim et al. (2020)), *chart-pair* (Section 3.2.1; our approach with s_p), and *chart-characteristic* method (Section 3.2.2; our method with s_c) respectively.

⁵In practice, we test every element $f \in F^d$ exhaustively to select the best one as F^d consists of only two elements.




Figure 1: Performance of CPE-PLM methods on PTB. Chart-based (CP and CC) approaches show superior figures in most cases compared to TD. The top-K ensemble also provides orthogonal improvements. Best viewed in color.

Language	French			German			Korean			Swedish		
Method	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC
Naïve baselines												
Random		16.2			13.9			22.2			16.4	
Left-branching		5.7			10.0			18.5			8.4	
Right-branching		26.4			14.7			19.2			30.4	
PCFGs												
N-PCFG [†]		42.2			37.8			25.7			14.5	
C-PCFG [†]		40.5			37.3			27.7			23.7	
Mono. PLMs	41.4	42.4	42.8	38.4	39.6	39.7	51.1	47.3	47.4	35.6	38.4	38.9

Table 2: CPE-PLM’s performance on French, German, Korean, and Swedish. The best score for each language is in **bold**. †: results from Zhao and Titov (2021).

4.2 Experiments on Monolingual Settings

We first assess CPE-PLM on the PTB dataset. We apply our methods to three different categories of English PLMs—BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and XLNet (Yang et al., 2019).⁶ We also test the effect of the top-K ensemble by combining it with each of the CPE-PLM method. In Figure 1, we confirm that our chart-based methods mostly outperform the top-down approach, showing an improvement of up to nearly five points (RoBERTa-large: 33.8 → 38.6). Moreover, we reveal that the top-K ensemble provides significant improvements on parsing performance in an orthogonal manner, regardless of the accompanying method. This result implies that the cues that can contribute to inducing parse trees may be distributed across different parts of PLMs, rather than concentrated on a specific attention head. We attain the 46.4 F1 score by combining XLNet, the CP method, and the top-K ensemble, which is six points higher than the best performance (40.1) reported in the previous work (Kim et al., 2020).

Next, we evaluate CPE-PLM methods on French, German, Korean, and Swedish, for which language-

⁶We prepare two variants for each PLM: (i) *X-base* consists of 12 layers, 12 attention heads, and 768 hidden dimensions. (ii) *X-large* has 24 layers, 16 heads, and 1024 dimensions.

specific BERT variants are available.⁷ As baselines, we prepare three naïve methods—random and left/right-branching trees—in addition to N(eural)-PCFG and C(ompound)-PCFG (Kim et al., 2019a), which are representative unsupervised parsers.⁸ For CPE-PLM, we consider both the top-down and chart methods, all of which are combined with the top-K ensemble. In Table 2, the CPE-PLM methods demonstrate performance comparable (French / German) or superior (Korean / Swedish) to that of the strong baselines. In particular, CPE-PLM shows much better performance in Korean and Swedish, where PCFGs failed to obtain meaningful results. We conjecture this discrepancy in part comes from the availability of subword-level features, to which PLMs have access while PCFGs do not, considering that the SPMRL dataset is originally constructed for testing morphologically-rich languages. Meanwhile, CP and CC outperform TD on 3 out of 4 languages, albeit the gaps are relatively small compared to the English case. This outcome leads to two implications: (i) the effectiveness of a CPE-PLM method depends on the language where it is applied, and (ii) our top-K ensemble is broadly helpful for all the parsing methods, reducing the gap between their performance.

4.3 Experiments on Multilingual Settings

Theoretically, multilingual PLMs have a potential to be a core asset for CPE-PLM, given that they are able to deal with sentences from over a hundred languages simultaneously. However, it has not yet been investigated whether they can play a role as expected. To shed light on this issue, we conduct experiments with the CPE-PLM

⁷The PLMs used per language are listed in Appendix.

⁸We rely on Zhao and Titov (2021) to report PCFG models’ performance. Note that the authors assume they have access to parses in the PTB dev set to select the best hyperparameters following Kim et al. (2019a). This condition is exactly the same for our English models and cross-lingual transfer settings in Table 3. Meanwhile, we additionally employ the SPMRL validation set in Table 2 and the Multi-ling. part of Table 3, which can give undesirable extra gains to CPE-PLM.

Language	English			Basque			French			German			Hebrew			Hungarian			Korean			Polish			Swedish		
Method	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC
PCFGs																											
N-PCFG [†]	50.8			30.2			42.2			37.8			41.0			37.9			25.7			31.7			14.5		
C-PCFG [†]	<u>55.7</u>			27.9			40.5			37.3			39.2			38.3			27.7			32.4			23.7		
Mono-ling.	42.9	46.4	45.0				41.4	42.4	42.8	38.4	39.6	39.7							51.1	47.3	47.4				35.6	38.4	38.9
Multi-ling.																											
M-BERT	40.3	45.0	44.8	40.0	40.9	41.6	42.9	44.6	45.6	39.3	40.6	40.3	42.3	42.5	42.0	38.2	39.1	40.4	52.1	50.9	49.8	41.6	43.0	42.9	37.3	37.4	39.3
XML	44.2	47.7	46.2	43.3	44.1	44.7	43.7	46.0	46.0	39.2	41.3	40.4	43.9	44.2	44.3	40.8	42.3	42.2	43.0	41.6	41.0	44.3	44.9	44.4	39.0	40.1	40.7
XML-R	45.5	46.7	47.0	43.7	43.8	45.1	45.8	44.2	45.5	41.4	42.2	41.6	45.0	43.2	45.3	42.4	44.0	43.4	55.9	55.7	54.3	43.1	43.7	44.6	39.5	40.6	41.5
XML-R-L	41.7	44.6	45.1	44.3	44.1	45.2	39.5	42.4	42.9	38.9	41.0	40.7	43.7	44.1	46.3	39.8	41.5	41.3	51.8	52.6	51.8	41.7	43.5	44.5	36.2	38.6	39.4
Cross-ling.																											
M-BERT				39.8	39.8	41.1	42.2	44.6	45.5	37.7	40.3	39.3	39.7	42.8	42.8	36.2	39.4	38.0	48.9	47.0	45.7	39.8	41.9	42.3	36.0	39.1	38.7
(+, -)				(-0.2)	(-1.1)	(-0.5)	(-0.7)	(0.0)	(-0.1)	(-1.6)	(-0.3)	(-1.0)	(-2.6)	(+0.3)	(+0.8)	(-2.0)	(+0.3)	(-2.4)	(-3.2)	(-3.9)	(-4.1)	(-1.8)	(-1.1)	(-0.6)	(-1.3)	(+1.7)	(-0.6)
XML				40.6	41.2	42.1	44.2	46.3	46.1	38.9	41.5	40.3	42.4	45.8	43.9	38.0	42.2	40.7	40.1	39.5	38.4	42.2	44.5	44.4	38.2	40.9	40.9
(+, -)				(-2.7)	(-2.9)	(-2.6)	(+0.5)	(+0.3)	(+0.1)	(-0.3)	(+0.2)	(-0.1)	(-1.5)	(+1.6)	(-0.4)	(-2.8)	(-0.1)	(-1.5)	(-2.9)	(-2.1)	(-2.6)	(-2.1)	(-0.4)	(0.0)	(-0.8)	(+0.8)	(+0.2)
XML-R				43.4	42.1	43.7	45.4	45.1	46.2	41.5	42.2	41.5	45.5	45.2	46.3	41.3	43.4	41.9	52.6	49.6	48.9	44.3	45.4	44.8	40.4	41.0	41.4
(+, -)				(-0.3)	(-1.7)	(-1.4)	(-0.4)	(+0.9)	(+0.7)	(+0.1)	(0.0)	(-0.1)	(+0.5)	(+2.0)	(+1.0)	(-1.1)	(-0.6)	(-1.5)	(-3.3)	(-6.1)	(-5.4)	(+1.2)	(+1.7)	(+0.2)	(+0.9)	(+0.4)	(-0.1)
XML-R-L				43.9	42.6	43.6	39.4	42.3	43.2	38.6	40.6	40.6	42.8	44.7	45.4	38.6	39.9	40.7	51.6	51.3	50.5	42.6	44.9	45.1	37.1	39.6	40.0
(+, -)				(-0.4)	(-1.5)	(-1.6)	(-0.1)	(-0.1)	(+0.3)	(-0.3)	(-0.4)	(-0.1)	(-0.9)	(+0.6)	(-0.9)	(-1.2)	(-1.6)	(-0.6)	(-0.2)	(-1.3)	(-1.3)	(+0.9)	(+1.4)	(+0.6)	(+0.9)	(+1.0)	(+0.6)

Table 3: Performance of CPE-PLM on 9 languages. **Mono-ling.**: CPE-PLM’s performance in monolingual settings. **Multi-ling.**: the results when combined with multilingual PLMs. **Cross-ling.**: the performance when relying on cross-lingual transfer, in addition to the relative losses or gains (+, -) compared to the original results. The best score per PLM is in **bold** while the best for each language is underlined. †: results from Zhao and Titov (2021).

methods built upon multilingual PLMs. We apply four multilingual PLMs to nine languages in total. We use a multilingual version of the BERT-base model (M-BERT, Devlin et al. (2019)), the XLM model trained on 100 languages (XLM, Conneau and Lample (2019)), XLM-R, and XLM-R-L(arge) (Conneau et al., 2019). For baselines, we only consider PCFGs as we verified in Section 4.2 that they can subsume naïve baselines. We also list CPE-PLM’s performance in monolingual settings for reference. Again, we utilize the TD, CP, and CC methods combined with the top-K ensemble.

In the **Multi-ling.** section of Table 3, we report CPE-PLM’s performance with multilingual PLMs when the best attention heads are separately selected for each language, relying on the validation sets of respective languages. We observe that the CPE-PLM framework works pretty well across languages when it is built upon multilingual PLMs, outperforming PCFGs except for English. Surprisingly, we discover that for every language we consider, there exists at least one multilingual PLM that outperforms its monolingual counterpart. For instance, we achieve the F1 score of 47.7 in English with the XLM model, which is higher than all the scores we achieved for English in monolingual settings. In conclusion, we confirm that multilingual PLMs can serve as a core component for an integrated CPE-PLM framework that processes different languages simultaneously. Regarding the effect of parsing strategies, we identify that CP and

CC generally outperform TD, and that the only exception occurs in Korean. We assume this is related to the linguistic properties of target languages, but we leave a thorough analysis on this as future work.

Next, we evaluate CPE-PLM in a harsher condition where the validation set is given only for English. Concretely, we attempt to parse sentences in eight other languages with the CPE-PLM methods optimized for English (i.e., the attention heads are chosen based on the PTB validation set), performing *zero-shot cross-lingual transfer* from English to others. Note that this constraint facilitates revealing the true value of CPE-PLM by answering the following research question: *given no access to parsers or gold-standard trees in target languages at all, can we induce non-trivial parse trees by solely relying on the knowledge residing in PLMs?*

In the **Cross-ling.** section, we present the performance of cross-lingual transfer and relative performance losses or gains (+, -) compared against the language-specific optimization (**Multi-ling.**). To our surprise, we reveal that the cross-lingual transfer leads to negligible losses or even small gains in most cases. This is also in line with the reports from related work (Pires et al. (2019); Cao et al. (2020); i.a.) that multilingual PLMs are effective in cross/multi-lingual NLP tasks. Our finding implies that there exist *universal* attention heads that are sensitive to the phrase structures of sentences irrespective of the input language. We seek to analyze this phenomenon in detail in the following section.

Language	English			Basque			French			German			Hebrew			Hungarian			Korean			Polish			Swedish		
	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC	TD	CP	CC
Pre-training data																											
Tokens (M) [†]	300.8			2.0			56.8			66.6			31.6			58.4			54.2			44.6			12.1		
Size (GiB) [†]	55608			270			9780			10297			3399			7807			5644			6490			77		
Val. & Test data																											
Size (Validation)	1700			948			1235			5000			500			1051			2066			821			494		
Size (Test)	2416			946			2540			4999			716			1009			2287			822			666		
XML-R	45.5	46.7	47.0	43.7	43.8	45.1	45.8	44.2	45.5	41.4	42.2	41.6	45.0	43.2	45.3	42.4	44.0	43.4	55.9	55.7	54.3	43.1	43.7	44.6	39.5	40.6	41.5
Cross-lingual				43.4	42.1	43.7	45.4	45.1	46.2	41.5	42.2	41.5	45.5	45.2	46.3	41.3	43.4	41.9	52.6	49.6	48.9	44.3	45.4	44.8	40.4	41.0	41.4
(+, -)				(-0.3)	(-1.7)	(-1.4)	(-0.4)	(+0.9)	(+0.7)	(+0.1)	(0.0)	(-0.1)	(+0.5)	(+2.0)	(+1.0)	(-1.1)	(-0.6)	(-1.5)	(-3.3)	(-6.1)	(-5.4)	(+1.2)	(+1.7)	(+0.2)	(+0.9)	(+0.4)	(-0.1)

Table 4: Factor correlation analysis. The first section describes the statistics of the data utilized for training XLM-R. The second section displays the characteristics of the validation and test sets. †: from [Conneau et al. \(2019\)](#).

5 Analysis

We present several analyses that enrich our understanding about CPE-PLM. We employ XLM-R as a backbone and the CC method as a parsing scheme.

5.1 Factor Correlation Analysis

First, we attend to two factors that may affect CPE-PLM’s performance: (i) the amount of the data consumed to train a PLM, and (ii) the number of sentences in the validation and test sets. In Table 4, we do not notice a clear relationship between the amount of pre-training data and performance. We conjecture this result is rooted in the sampling technique exploited when pre-training XLM-R. Specifically, the technique readjusts the probability of sampling a sentence from each language, increasing the number of tokens sampled from low-resource languages while mitigating the bias towards high-resource languages ([Conneau et al., 2019](#)). On the other hand, we discover that the languages for which the size of the validation sets are relatively small (i.e., Hebrew, Polish, and Swedish) tend to benefit from cross-lingual transfer, implying that the insufficient number of examples in the validation set might cause some noise or lead to the suboptimal in the selection of attention heads.

5.2 Visualization of Attention Heads

We revealed in Section 4.3 that CPE-PLM’s performance for most languages does not suffer much from cross-lingual transfer, suggesting that there would exist significant overlaps among the sets of the attention heads selected for respective languages. To verify our hypothesis, in Figure 2, we visualize the language-specific sets of the top 20 heads existing in XLM-R. We observe that the heads effective for CPE-PLM are distributed over the middle-to-upper (6-12) layers of XLM-R, implying that phrase-level information is pervasive




Figure 2: Visualization of the sets of the top 20 attention heads (in XLM-R) for 9 languages. Each cell is filled with the color assigned for a language if the corresponding head is responsible for parsing the language.

in the upper layers rather than the lower ones. In addition, we discover that most of the heads detected as sensitive to syntax respond to multiple languages simultaneously and that there exist a few heads proven to be important for dealing with all the nine languages we consider. Our finding of the existence of such *universal* attention heads explains why CPE-PLM is robust to cross-lingual transfer in multilingual settings, in addition to providing a partial clue on why multilingual PLMs are excel at cross-lingual transfer as reported in previous work ([Pires et al. \(2019\)](#); [Cao et al. \(2020\)](#); *inter alia*).

5.3 Recall Scores on Noun and Verb Phrases

To assess CPE-PLM’s performance in a more fine-grained manner and probe the extent to which it detects the core components of sentences, we present its recall scores on gold-standard noun and verb phrases. We only target the languages whose

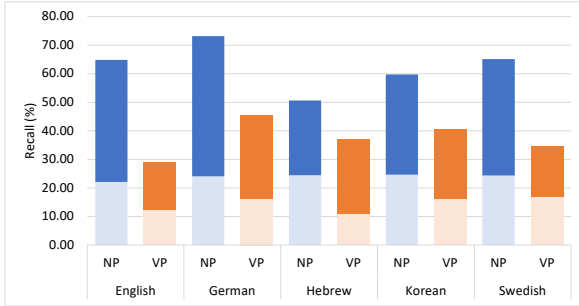


Figure 3: Recall scores on gold-standard NPs and VPs. The light bars indicate the random baseline’s performance while the dark ones show that of the CC method.

gold-standard trees contain proper tags in their test sets. In Figure 3, we confirm that compared to the random baseline, CPE-PLM has a decent ability to identify noun phrases, succeeding in retrieving more than half of NPs for every language. On the contrary, CPE-PLM seems relatively weak in recognizing VPs which are generally longer and more complex than NPs. This implies that CPE-PLM might struggle with grasping the whole structure of sentences (e.g., VPs), although it successfully perceives small phrasal components (e.g., NPs).

6 Limitations and Future Work

We here mention a few limitations of our approach and propose avenues for future work. First, analogous to several unsupervised parsers (Shi et al., 2020) including PCFGs (Zhao and Titov, 2021), the current form of our method relies on a few gold-standard annotations from the validation set to determine the best hyperparameters (i.e., the best choice for attention head selection). This dependency makes it hard to say that our approach is entirely unsupervised, although it steps aside from a typical way of learning parsers with supervision. A next, promising yet challenging, step will be therefore to develop a remedy that enables our method to be free from the annotations, similar to Li et al. (2020). Note that our cross-lingual transfer experiments also shed some light on how to relieve such dependency.

While we have shown that CPE-PLM can be superior or comparable to PCFGs and that it can function as an effective tool for analyzing PLMs, its performance still falls short of expectations in terms of whether it can practically substitute standard parsers, similar to the case of unsupervised parsers. To improve its performance, we have a plan as future work to design an ensemble method that

gathers information from heterogeneous PLMs.

Finally, chart parsing algorithm, whose time complexity is $O(n^3)$, is inherently much more expensive than other efficient parsing strategies such as top-down parsing. Therefore, when we need to decide which parsing algorithm to employ, we should keep in mind the trade-off between accuracy and efficiency.

7 Related work

Pre-trained language models (PLMs) now lie at the heart of many studies in the literature. Following the trend, much effort has been made to develop English PLMs (Devlin et al. (2019); Liu et al. (2019); Radford et al. (2019); Yang et al. (2019), *inter alia*), to construct non-English PLMs (Martin et al. (2019); *i.a.*), and to train multilingual variants (Conneau and Lample, 2019; Conneau et al., 2019). We have explored the potential use of these PLMs as parsers. The trees induced by our method can also be leveraged as a tool for probing PLMs, similar to recent work that attempt to explore the knowledge in PLMs (Clark et al., 2019; Jawahar et al., 2019). In particular, Chi et al. (2020) extended Hewitt and Manning (2019) to multilingual settings, analogous to our work. Still, it is different from ours in that it requires explicit supervision and devotes itself to dependency grammar.

In this study, we have named a line of work (Rosa and Mareček, 2019; Wu et al., 2020; Kim et al., 2020) that extracts trees directly from PLMs as CPE-PLM. We have extended its application to multilingual scenarios as well as improving its performance. Notably, Mareček and Rosa (2019) developed an approach similar to ours, but they focused on Transformers trained for machine translation rather than language models. Work on neural unsupervised parsing (Shen et al. (2018b, 2019); Kim et al. (2019a); Shi et al. (2020), *inter alia*) also seeks to generate parse trees without supervision from gold-standard trees. It is worth noting that some work such as Kann et al. (2019) and Zhao and Titov (2021) attempt to evaluate unsupervised parsers in multilingual settings, akin to our work. The difference between ours and theirs is that our method does not require training a parser for each language, instead relying on off-the-shelf PLMs.

8 Conclusion

In this work, we study **Constituency Parse Extraction from Pre-trained Language Models (CPE-PLM)**, a novel

paradigm of inducing parses directly from PLMs. We introduce a chart-based method and top-K ensemble for improving performance, and extend the range of application of the paradigm to different languages by applying multilingual PLMs. We hope our work can function as the foundation for future research on (i) unsupervised constituency parsing for under-represented languages and (ii) probing the inner workings of multilingual PLMs.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2020-0-01373, Artificial Intelligence Graduate School Program (Hanyang University) and No.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)).

References

- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. Multilingual alignment of contextual word representations. In *ICLR*.
- Eugene Charniak. 2000. A maximum-entropy-inspired parser. In *NAACL*.
- Ethan A. Chi, John Hewitt, and Christopher D. Manning. 2020. Finding universal grammatical relations in multilingual BERT. In *ACL*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What does BERT look at? an analysis of BERT’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- John Cocke. 1969. *Programming languages and their compilers: Preliminary notes*. New York University.
- Michael Collins. 2003. Head-driven statistical models for natural language parsing. *Computational Linguistics*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *NeurIPS*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Greg Durrett and Dan Klein. 2015. Neural CRF parsing. In *ACL-IJCNLP*.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. In *NAACL*.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does BERT learn about the structure of language? In *ACL*.
- Katharina Kann, Anhad Mohanane, Samuel R. Bowman, and Kyunghyun Cho. 2019. Neural unsupervised parsing beyond English. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*.
- Tadao Kasami. 1966. An efficient recognition and syntax-analysis algorithm for context-free languages. *Coordinated Science Laboratory Report no. R-257*.
- Taeuk Kim, Jihun Choi, Daniel Edmiston, and Sangwoo Lee. 2020. Are pre-trained language models aware of phrases? simple but strong baselines for grammar induction. In *ICLR*.
- Yoon Kim, Chris Dyer, and Alexander Rush. 2019a. Compound probabilistic context-free grammars for grammar induction. In *ACL*.
- Yoon Kim, Alexander Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019b. Unsupervised recurrent neural network grammars. In *NAACL*.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. Multilingual constituency parsing with self-attention and pre-training. In *ACL*.
- Nikita Kitaev and Dan Klein. 2018. Constituency parsing with a self-attentive encoder. In *ACL*.
- Bowen Li, Taeuk Kim, Reinald Kim Amplayo, and Frank Keller. 2020. Heads-up! unsupervised constituency parsing via self-attention heads. In *ACL-IJCNLP*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*.

- David Mareček and Rudolf Rosa. 2019. From balustrades to pierre vinken: Looking for syntax in transformer self-attentions. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villamonte de la Clergerie, Djamé Seddah, and Benoît Sagot. 2019. Camembert: a tasty french language model. *arXiv preprint arXiv:1911.03894*.
- Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. Learning accurate, compact, and interpretable tree annotation. In *COLING-ACL*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual BERT? In *ACL*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Rudolf Rosa and David Mareček. 2019. Inducing syntactic trees from bert representations. *arXiv preprint arXiv:1906.11511*.
- Djamé Seddah, Reut Tsarfaty, Sandra Kübler, Marie Candito, Jinho D. Choi, Richárd Farkas, Jennifer Foster, Iakes Goenaga, Koldo Gojenola Galletebeitia, Yoav Goldberg, Spence Green, Nizar Habash, Marco Kuhlmann, Wolfgang Maier, Joakim Nivre, Adam Przepiórkowski, Ryan Roth, Wolfgang Seeker, Yannick Versley, Veronika Vincze, Marcin Woliński, Alina Wróblewska, and Eric Villamonte de la Clergerie. 2013. Overview of the SPMRL 2013 shared task: A cross-framework evaluation of parsing morphologically rich languages. In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*.
- Yikang Shen, Zhouhan Lin, Athul Paul Jacob, Alessandro Sordoni, Aaron Courville, and Yoshua Bengio. 2018a. Straight to the tree: Constituency parsing with neural syntactic distance. In *ACL*.
- Yikang Shen, Zhouhan Lin, Chin wei Huang, and Aaron Courville. 2018b. Neural language modeling by jointly learning syntax and lexicon. In *ICLR*.
- Yikang Shen, Shawn Tan, Alessandro Sordoni, and Aaron Courville. 2019. Ordered neurons: Integrating tree structures into recurrent neural networks. In *ICLR*.
- Haoyue Shi, Karen Livescu, and Kevin Gimpel. 2020. On the role of supervision in unsupervised constituency parsing. In *EMNLP*.
- Mitchell Stern, Jacob Andreas, and Dan Klein. 2017. A minimal span-based neural constituency parser. In *ACL*.
- 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 *NeurIPS*.
- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020. Perturbed masking: Parameter-free probing for analyzing and interpreting BERT. In *ACL*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *NeurIPS*.
- Daniel H Younger. 1967. Recognition and parsing of context-free languages in time n³. *Information and control*, pages 189–208.
- Yanpeng Zhao and Ivan Titov. 2021. An empirical study of compound pcfgs. *arXiv preprint arXiv:2103.02298*.
- Junru Zhou and Hai Zhao. 2019. Head-driven phrase structure grammar parsing on Penn treebank. In *ACL*.

A Appendix

A.1 List of Non-English Monolingual PLMs

The PLMs we utilize per language are listed as follows. **German:** bert-base-german (<https://deepset.ai/german-bert>). **French:** camembert (Martin et al., 2019). **Swedish:** bert-base-swedish (https://github.com/huggingface/transformers/tree/master/model_cards/KB/bert-base-swedish-cased). **Korean:** KoBERT-base (<https://github.com/SKTBrain/KoBERT>).