

The Emergence of Chunking Structures with Hierarchical RNN

Zijun Wu¹, Anup Anand Deshmukh^{2,*}, Yongkang Wu³,
Jimmy Lin², and Lili Mou^{1,4}

¹University of Alberta, Department of Computing Science, Alberta Machine Intelligence Institute (Amii)

zijun4@ualberta.ca, doublepower.mou@gmail.com

²University of Waterloo, David R. Cheriton, School of Computer Science

aa2deshmukh@uwaterloo.ca, jimmylin@uwaterloo.ca

³Huawei Poisson Lab

wuyongkang7@huawei.com

⁴Canada CIFAR AI Chair

In Natural Language Processing (NLP), predicting linguistic structures, such as parsing and chunking, has mostly relied on manual annotations of syntactic structures. This article introduces an unsupervised approach to chunking, a syntactic task that involves grouping words in a non-hierarchical manner. We present a Hierarchical Recurrent Neural Network (HRNN) designed to model word-to-chunk and chunk-to-sentence compositions. Our approach involves a two-stage training process: pretraining with an unsupervised parser and finetuning on downstream NLP tasks. Experiments on multiple datasets reveal a notable improvement of unsupervised chunking performance in both pretraining and finetuning stages. Interestingly, we observe that the emergence of the chunking structure is transient during the neural model's downstream-task training. This study contributes to the advancement of unsupervised syntactic structure discovery and opens avenues for further research in linguistic theory.¹

1. Introduction

Understanding the linguistic structure of language, such as parsing and chunking, is an important research topic in natural language processing (NLP). While most previous work (Kudo and Matsumoto 2001; Zhang, Damerau, and Johnson 2002; Pradhan et al.

* Work partially done as a co-op intern at the University of Alberta/Amii.

Action Editor: Deyi Xiong. Submission received: 11 September 2023; revised version received: 17 September 2024; accepted for publication: 23 October 2024.

<https://doi.org/10.1162/coli.a.00545>

1 This paper is a substantially extended version of an EMNLP Findings paper (Deshmukh et al. 2021). The main extension is a novel approach that improves unsupervised chunking performance in a downstream NLP task, which is an interesting result by itself. Part of the text is reused with permission from our previous work (Deshmukh et al. 2021), available at <https://aclanthology.org/2021.findings-emnlp.307.pdf>, licensed under CC BY 4.0. Our code and results are available at <https://github.com/MANGA-UOFA/UCHRNN>.

2004) uses supervised machine learning methods to predict linguistic structures and can achieve high performance, they rely heavily on manual annotations of syntactic structures like Treebanks (Marcus, Santorini, and Marcinkiewicz 1993). As a result, there has been a growing interest in unsupervised linguistic structure discovery in recent years (Kim et al. 2020; Shen et al. 2018, 2019), which is important to NLP research because it sheds light on linguistic theories and can potentially benefit low-resource languages. Previous work on unsupervised syntactic structure discovery has mainly focused on unsupervised constituency parsing (Kim et al. 2020; Kim, Dyer, and Rush 2019; Shen et al. 2018, 2019) and dependency parsing (Klein and Manning 2004; Gillenwater et al. 2010; Shen et al. 2021).

In this work, we address unsupervised chunking, another meaningful task for discovering syntactic structure. Unlike parsing that induces tree structures from a sentence, chunking aims to group consecutive words of a sentence in a non-hierarchical fashion, and each chunk can be intuitively thought of as a phrase (Tjong Kim Sang and Buchholz 2000; Clark 2001). In fact, unsupervised chunking has wide applications in real-world scenarios, as understanding text fundamentally requires finding spans like noun phrases and verb phrases. It would benefit other NLP tasks, such as keyword extraction (Firoozeh et al. 2020), named entity recognition (Sato et al. 2017), open information extraction (Niklaus et al. 2018; Borgeaud et al. 2022; Izacard et al. 2022), and logical reasoning (Wu et al. 2023).

We propose a hierarchical recurrent neural network (HRNN) to accomplish the chunking task. Our HRNN is designed to explicitly model the word-to-chunk composition by a lower-level HRNN, and chunk-to-sentence composition by an upper-level HRNN. We further design a trainable chunking gate that switches between lower word-level HRNN and upper phrase-level HRNN, which is also used for chunk prediction.

We propose a two-stage training framework for HRNN without using linguistic annotations, namely, pretraining with an unsupervised parser and finetuning with a downstream NLP task. In the first stage of pretraining, we adopt the recent advances of unsupervised parsing (Kim et al. 2020; Kim, Dyer, and Rush 2019), and propose a maximal left-branching heuristic to induce sensible (albeit noisy and imperfect) chunk labels from an unsupervised parser. In the second stage, we finetune HRNN chunking by feeding its upper-level HRNN representations to a downstream-task network. Our intuition is that HRNN serves as the bottleneck, and the downstream-task finetuning can force the HRNN to learn more meaningful latent structures.

We conducted experiments on the CoNLL-2000 dataset (Tjong Kim Sang and Buchholz 2000) and the English Web Treebank (Bies et al. 2012). Results show that the unsupervised parser-pretrained HRNN significantly improves the best-performing unsupervised baseline, with a considerable margin of up to 6 percentage points in terms of the phrase F1 score. We then finetuned the pretrained HRNN model with three downstream text generation tasks: summarization, paraphrasing, and translation. Compared with our pretrained model, it achieves a chunking performance improvement of up to 6 percentage points in downstream datasets, and 2 percentage points on CoNLL-2000. The results suggest that our method not only bridges the gap between supervised and unsupervised chunking methods but also shows the generalizability across different downstream tasks and datasets.

In our experiment, we also find an intriguing linguistic phenomenon during the finetuning stage: The neural model's emergence of linguistic structure is transient. That is, although the downstream performance consistently improves, the chunking performance improves significantly but only in the first few thousand steps. Then, it starts to decrease and eventually drops to the initial level of the finetuning stage or even

lower, suggesting that linguistic structures are a convenient vehicle for a downstream task when the training begins and the model capacity is small. However, the neural network tends to discard such linguistic structures to achieve higher performance in the downstream task as the training proceeds.

In summary, our main contributions are as follows:

- We address the task of unsupervised chunking, and propose a HRNN that can explicitly model the word-to-chunk and chunk-to-sentence compositions.
- We propose a two-stage training framework for HRNN, largely outperforming previous unsupervised chunking methods.
- We observe that the neural model’s emergence of chunking structure is transient, which may inspire further study of linguistic theory.

2. Approach

In this section, we first introduce the proposed HRNN model. Then, we discuss the process of pretraining HRNN with the induced labels from a state-of-the-art unsupervised parser. Finally, we propose to finetune the model with a downstream task to improve chunking performance.

2.1 Hierarchical RNN

We model the chunking task in the BI schema (Ramshaw and Marcus 1995), where “B” denotes the beginning of a chunk, and “I” denotes the inside of a chunk. We observe that encoder-only HRNNs or Transformers may not be suitable for the chunking task because they lack autoregressiveness, which means that the prediction at a time step is unaware of previously predicted chunks. Feeding predicted chunk labels like a sequence-to-sequence model is also inadequate because a BI label only provides one-bit information and does not offer meaningful autoregressive information.

To this end, we design a HRNN to model the autoregressiveness of predicted chunks. Our HRNN contains a lower-level HRNN operating at the word level and an upper-level HRNN operating at the chunk level. We design a gating mechanism that switches between the two HRNNs in a soft manner, also serving as the chunk label. The bottom-left part of Figure 1 shows the structure of HRNN.

Given a sentence with words $x^{(1)}, \dots, x^{(n)}$, we first apply a pretrained Transformer encoder (Devlin et al. 2019) to obtain the contextual representations of the words that help our model to understand the previous and future context of the sentence. We denote the contextual embeddings of these words by $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$. For a step t , we predict a switching gate $m^{(t)} \in (0, 1)$ as the chunking decision.²

$$m_{\text{logit}}^{(t)} = W[\underline{\mathbf{h}}^{(t-1)}; \bar{\mathbf{h}}^{(t-1)}; \mathbf{x}^{(t)}] \quad (1)$$

$$m^{(t)} = \sigma(m_{\text{logit}}^{(t)}) \quad (2)$$

² Here, $m^{(t)} = 1$ corresponds to “B,” i.e., a new chunk, and $m^{(t)} = 0$ corresponds to “I,” i.e., inside of a chunk.

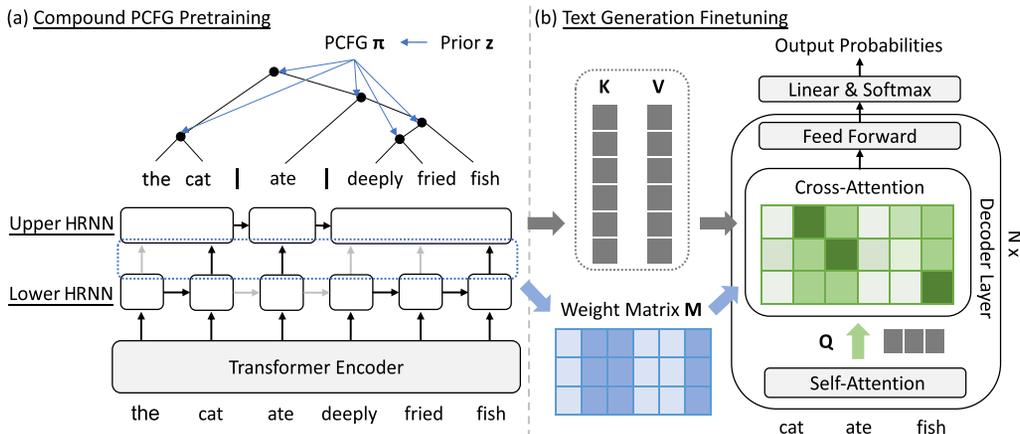


Figure 1 An overview of our chunking induction method. (a) Pretraining HRNN using the chunk labels induced from the Compound PCFG parser. (b) HRNN is finetuned with text generation, specifically, a summarization task in this example. A weight matrix is created from the switching gate’s values and then inserted into the Transformer’s encoder-decoder attention modules.

where $\underline{h}^{(t-1)}$ is the hidden state of the lower HRNN and $\bar{h}^{(t-1)}$ is that of the upper HRNN. The semicolon represents vector concatenation, and σ represents the sigmoid function.

The switching gate is also used to control the information flow between the lower-level word HRNN and the upper-level chunk HRNN. In this way, it provides meaningful autoregressive information that enables HRNN to be aware of previously detected chunks.

Suppose our model predicts the t th word as the beginning of a chunk, which essentially “cuts” the sequence into two parts at this step. The lower HRNN and upper HRNN are updated by

$$\underline{h}_{\text{cut}}^{(t)} = f(\underline{x}^{(t)}, \underline{h}^{(\text{sos})}) \tag{3}$$

$$\bar{h}_{\text{cut}}^{(t)} = \bar{f}(\underline{h}^{(t-1)}; \bar{h}^{(t-1)}) \tag{4}$$

where f and \bar{f} are the transition functions of the lower and upper HRNNs, respectively, both utilizing the hyperbolic tangent (\tanh) as their activation function.

In Equation (3), the lower HRNN ignores its previous hidden state and starts anew from a learnable initial state $\underline{h}^{(\text{sos})}$, because a new chunk is predicted at this step. On the other hand, in Equation (4), the upper HRNN picks the newly formed chunk $\underline{h}^{(t-1)}$ captured by the lower HRNN, and fuses it with the previous upper HRNN’s state $\bar{h}^{(t-1)}$.

Suppose our model predicts that the t th word is not the beginning of a chunk, i.e., “no cut” occurs at this step. The HRNNs are then updated by

$$\underline{h}_{\text{nocut}}^{(t)} = f(\underline{x}^{(t)}, \underline{h}^{(t-1)}) \tag{5}$$

$$\bar{h}_{\text{nocut}}^{(t)} = \bar{h}^{(t-1)} \tag{6}$$

In this scenario, the lower HRNN updates its hidden state with the input $\mathbf{x}^{(t)}$ in the same manner as a typical HRNN, whereas the upper HRNN remains idle because no chunk is formed.

The “cut” and “no cut” cases can be unified by

$$\bar{\mathbf{h}}^{(t)} = m^{(t)}\bar{\mathbf{h}}_{\text{cut}}^{(t)} + (1 - m^{(t)})\bar{\mathbf{h}}_{\text{nocut}}^{(t)} \quad (7)$$

$$\underline{\mathbf{h}}^{(t)} = m^{(t)}\underline{\mathbf{h}}_{\text{cut}}^{(t)} + (1 - m^{(t)})\underline{\mathbf{h}}_{\text{nocut}}^{(t)} \quad (8)$$

In our work, we keep $m^{(t)}$ as a real number, which is essentially a gating mechanism that fuses “cut” and “no cut” in a soft manner. This is because chunking, by its nature, may be ambiguous, and our soft gating mechanism can better handle such ambiguity.

Our proposed HRNN supports end-to-end learning, which can be weakly supervised by downstream tasks. However, a randomly initialized HRNN may have difficulty modeling the desired chunking structure and suffer from the cold-start problem. We thus propose to pretrain HRNN by heuristically induced chunking labels from an unsupervised parser, which will be discussed in Section 2.2.

2.2 Pretraining HRNN by Unsupervised Parsing

We propose to induce chunk labels from an unsupervised parser. The intuition is that the chunking structure can be considered a flattened parse tree, thus sharing some commonalities with the parsing structure. Our approach is able to take advantage of recent advances in unsupervised parsing (Kim et al. 2020; Kim, Dyer, and Rush 2019).

Specifically, we adopt a state-of-the-art unsupervised parser, Compound Probabilistic Context-Free Grammar (PCFG; Kim, Dyer, and Rush 2019). A PCFG is a tuple $\mathcal{G} = (\mathcal{S}, \mathcal{N}, \mathcal{P}, \Sigma, \mathcal{R})$, where \mathcal{S} is a start symbol; \mathcal{N} , \mathcal{P} , and Σ are finite sets of nonterminal, preterminal, and terminal symbols, respectively. \mathcal{R} is a finite set of rules taking one of the following forms:

$$S \rightarrow A \quad A \in \mathcal{N} \quad (9)$$

$$A \rightarrow BC \quad B, C \in \mathcal{N} \cup \mathcal{P} \quad (10)$$

$$T \rightarrow w \quad T \in \mathcal{P}, w \in \Sigma \quad (11)$$

where $S \rightarrow A$ is the start of a sentence and $T \rightarrow w$ indicates the generation of a word. $A \rightarrow BC$ models the bifurcations of a binary constituency tree, where a constituent node is not explicitly associated with a type (e.g., noun phrase) in our setting.

In addition, the model maintains a continuous random vector at the sentence level, which serves as the prior of PCFG. To train Compound PCFG, the maximum likelihood of the text is utilized, and a Viterbi-like algorithm marginalizes the PCFG. Amortized variational inference is employed to handle the continuous distribution from the sentence-level vector. We refer readers to Kim, Dyer, and Rush (2019) for details.

We aim to induce chunk labels from the state-of-the-art unsupervised parser, Compound PCFG. Given a sentence, we obtain its parse tree by applying the Viterbi-like CYK algorithm (Younger 1967; Jurafsky and Martin 2000) to Compound PCFG. Then, we propose a simple yet effective heuristic that extracts maximal left-branching subtrees as chunks. It is known that the English language has a strong bias toward right-branching structures (Williams, Drozdov, and Bowman 2018; Li, Mou, and Keller

2019). We observe, on the other hand, that a left-branching structure often indicates words that are closely related. Here, a left-branching subtree means that the words are grouped in the form of $((\dots((x_i x_{i+1}) x_{i+2}) \dots) x_{i+n-1})$. A left-branching subtree for words $x_i \dots x_{i+n-1}$ is maximal if neither $x_{i-1} x_i \dots x_{i+n-1}$ nor $x_i \dots x_{i+n-1} x_{i+n}$ is left-branching. We extract all maximal left-branching subtrees as chunks. Our heuristic can provide unambiguous chunk labels for any sentence with any parse tree, as demonstrated by the following theorem.

Theorem 1

Given any binary parse tree, every word will belong to one and only one chunk by the maximal left-branching heuristic.

Proof. [Existence] A single word itself is a left-branching subtree, which belongs to some maximal left-branching subtree.

[Uniqueness] We will show that two different maximal left-branching subtrees s_1 and s_2 cannot overlap. Assume by way of contradiction that there exists a word x_i in both s_1 and s_2 . Then, s_1 must be a substructure of s_2 or vice versa; otherwise, the paths, $\text{root}-s_1-x_i$ and $\text{root}-s_2-x_i$, violate the acyclic nature of a tree. But s_1 being a subtree of s_2 (or vice versa) contradicts the maximality of s_1 and s_2 . \square

The theorem shows that our maximal left-branching heuristic can unambiguously assign chunk labels based on a binary parse tree of any sentence.

We induce the labels by our heuristic and pretrain HRNN on those labels. While the heuristic is simple and may not be perfect, it still achieves satisfactory chunking results. Let T be the length of a sentence. Our pretraining objective is to minimize the following loss function for a sample:

$$\mathcal{L}_{\text{pretrain}} = - \sum_{t=1}^T [\tilde{y}_t \cdot \log m^{(t)} + (1 - \tilde{y}_t) \cdot \log(1 - m^{(t)})] \quad (12)$$

where \tilde{y}_t is the induced chunk label for the t th step, and $m^{(t)}$ is the chunking decision predicted by HRNN. As a machine learning model, the HRNN is able to smooth out the noise from the heuristic and yield more meaningful chunks, shown in Section 3.3.

2.3 Finetuning Hierarchical RNN with Downstream Tasks

We propose to finetune the HRNN in a weakly supervised manner to learn better chunking structures from a downstream task. Our intuition is that accomplishing a downstream NLP task requires understanding meaningful semantic units/chunks of a sentence (Wu et al. 2023), and that optimizing the downstream task may benefit the chunk prediction.

Specifically, we consider text generation tasks, namely, summarization (Rush, Chopra, and Weston 2015; Liu, Huang, and Mou 2022; Liu 2019), paraphrasing (Liu et al. 2020a; Li et al. 2019), and translation (Bojar et al. 2014; Liu et al. 2020b). Although the HRNN may be finetuned with classification tasks as well, we found in our preliminary experiments that classification tasks yield marginal improvement, probably because the classification training signal is too sparse (one label per sample) in comparison to generation tasks. Therefore, we only consider finetuning HRNN with text generation tasks in this work.

We propose a decoder network, whose attention is based on the upper-level HRNN (chunk representations). This is accomplished by performing conventional attention over the HRNN’s states but reweighing it by whether a chunk is predicted at the step. Since the HRNN’s gate is differentiable, we can train it end to end with the downstream tasks.

We first argue that the chunk representation is indeed given by the upper-level HRNN where a “cut” is predicted. Suppose we have a chunk $x^{(i)} \dots x^{(i+k)}$, that is, HRNN predicts two consecutive “cuts” (end of a chunk) at the time steps $i - 1$ and $i + k$. According to our design (Section 2.1), the lower-level HRNN initializes its hidden state with $\underline{h}^{(\text{sos})}$ when $t = i - 1$, and updates its hidden state $\underline{h}^{(t)}$ with the input words $x^{(i)} \dots x^{(i+k)}$. In the meantime, the upper-level HRNN is idle for time steps i to $i + k - 1$, but picks up the $\underline{h}^{(i+k)}$ at step $i + k$. This follows a conventional HRNN, demonstrating that the chunk representation is given by the upper-level HRNN at a “cut” step.

Now consider the scenarios with “soft” chunks, since we keep the chunking decision $m^{(t)}$ as a real number in our HRNN. When the “cut” strength is high ($m^{(t)}$ close to 1), the step is more likely to be a chunk, and we would like to attend to this step more from the decoder. Conversely, if the “cut” strength is low, it should be less attended. This can be achieved by applying conventional encoder–decoder attention to every step, but reweighing it by the “cut” strength. For each attention head, we form a query vector $\mathbf{q}^{(j)}$ for the j th decoder step and a key vector $\mathbf{k}^{(i)}$ for the i th encoder step. Here, the key vector $\mathbf{k}^{(i)}$ is given by the upper-level HRNN’s state $\bar{\mathbf{h}}^{(i)}$. We reweigh the attention by first computing an unnormalized measure:

$$\tilde{\alpha}(\mathbf{q}^{(j)}, \mathbf{k}^{(i)}, m_{\text{logit}}^{(i)}) = \exp\left(\frac{(\mathbf{q}^{(j)})^\top \mathbf{k}^{(i)}}{\sqrt{d}} + \gamma m_{\text{logit}}^{(i)}\right) \tag{13}$$

where γ is a coefficient hyperparameter. Then, we normalize it to attention probabilities as

$$\alpha^{(j,i)} = \frac{\tilde{\alpha}(\mathbf{q}^{(j)}, \mathbf{k}^{(i)}, m_{\text{logit}}^{(i)})}{\sum_{i'} \tilde{\alpha}(\mathbf{q}^{(j)}, \mathbf{k}^{(i')}, m_{\text{logit}}^{(i')})} \tag{14}$$

Empirically, we find that an end-to-end reweighted attention head tends to ignore $m_{\text{logit}}^{(i)}$ and make it close to 0. Consequently, the gating probability $m^{(i)}$ will be close to 0.5 unanimously, and the HRNN will be in a situation of neither “cut” nor “no cut.” We alleviate this by proposing an auxiliary training objective that pushes $m^{(i)}$ to either 0 or 1. We introduce a chunking strength hyperparameter $\kappa\%$, which is the desired chunk–token ratio. We then define the loss at the i th encoder step as

$$L^{(i)} = \begin{cases} (m^{(i)} - 1)^2, & \text{if } m^{(i)} \text{ is among the top-}\kappa\% \text{ gate values,} \\ (m^{(i)} - 0)^2, & \text{otherwise.} \end{cases} \tag{15}$$

A detailed analysis of this hyperparameter is presented in Section 3.4.4, and the value of $\kappa\%$ is decided by validation performance.

The overall auxiliary loss is $L_{\text{aux}} = \sum_i L^{(i)}$, and is combined with the end-to-end downstream task by $L_{\text{task}} + \eta L_{\text{aux}}$, where η is a hyperparameter balancing the two

Table 1

Details on the dataset and implementation. Symbol η represents the balance between two training objectives, while γ is the hyperparameter for the reweight operation, as defined in Equation 13. The term top- κ denotes the chunking strength heuristics.

Task	Dataset	#Train	#Dev	#Test	Batch size	Learning rate	η	γ	top- κ %
Chunking	CoNLL-2000	8K	1K	2K	32	5e-5	-	-	-
	English Web Treebank	-	1.9K	1K	-	-	-	-	-
Downstream	Gigaword	3.8M	2K	2K	32	5e-5	0.1	0.1	50%
	MNLI-Gen	131K	3.5K	3.5K	32	4e-5	0.1	0.1	60%
	WMT-14 (En-De)	4.5M	3K	2.7K	32	1e-5	0.1	0.1	50%

training objectives. In this way, we are able to learn a meaningful chunking gate in the downstream NLP task.

3. Experiments

3.1 Datasets and Metrics

Table 1 summarizes the datasets used in our study. We adopted a widely used chunking dataset, CoNLL-2000 (Tjong Kim Sang and Buchholz 2000), to evaluate the general chunking performance of a model. CoNLL-2000 contains 8K training, 1K validation, and 2K test samples. Each sample is labeled with the BIO schema, where “B” indicates the beginning of a chunk, “I” indicates inside a chunk, and “O” indicates outside a chunk (mainly punctuation). We used the BI schema and ignore the “O” tokens.

In this work, however, we addressed chunking in the unsupervised setting, and thus we did not use the groundtruth chunk labels in the training set, but only the training sentences. Groundtruth chunk labels are used for validation³ and test purposes.

Further, we used the English Web Treebank (Bies et al. 2012) to evaluate the general chunking performance on the online review domain, following Deshmukh et al. (2021). The dataset does not contain linguistically annotated chunk labels, and therefore we resorted to the NLTK toolkit (Bird and Loper 2004) to obtain pseudo-groundtruth chunk labels.

We finetuned the HRNN model with several downstream text generation tasks, whose datasets are introduced as follows.

- Gigaword (Rush, Chopra, and Weston 2015): A widely used article-headline dataset for text summarization with about 3.8M training, 20K validation, and 2K test samples. We randomly picked 2K samples (same size as the test set) for more efficient validation.
- MNLI-Gen: MNLI (Williams, Nangia, and Bowman 2018) is a massive dataset for the natural language inference task. We used the entailment subset, generating the entailed hypothesis based on the premise. Since the test set of MNLI is private, we adopted the matched section of the

³ We will discuss the use of labeled validation dataset in Section 3.4.6.

original validation set as our test set, while the mismatched section served as our validation set. The resulting entailment generation dataset contains 131K training, 3.5K validation, and another 3.5K test samples.

- WMT-14 (Bojar et al. 2014): A multilingual dataset used for machine translation. We used the English-to-German subset, having nearly 4.5M training, 3K validation, and 2.7K test samples.

In addition to the general chunking performance on the CoNLL-2000 and English Web Treebank corpora, we are interested in in-domain chunking performance after finetuning with a downstream task, whose test set, unfortunately, does not have human annotation of chunking labels. Therefore, we also used NLTK (Bird and Loper 2004) to produce pseudo-groundtruth labels for the source sentences of the downstream datasets.

Regarding the evaluation metrics, we adopted the standard phrase F1 score and tagging accuracy. In our implementation, they are realized by the CoNLL-2000 evaluation script (Tjong Kim Sang and Buchholz 2000).

3.2 Implementation Details

We utilized an encoder-only BERT model and two encoder–decoder models, BART and mBART, in our experiments. For the BERT model (Devlin et al. 2019), we adopted its base version that has 12 Transformer layers with 768 hidden dimensions, and its total parameter number is 110 million. For the BART model (Lewis et al. 2020), we also used its base version, which has 6 encoder layers and another 6 decoder layers. It has 768 dimensions and 139 million parameters in total. The mBART model (Liu et al. 2020b) has 610 million parameters, both the encoder and decoder having 12 layers and 1,024 dimensions. We only used the encoder module of BART or mBART for the chunking pretraining, which is then combined with the respective pretrained decoder for finetuning on downstream tasks. Our HRNN has the same dimension as the underlying Transformer models.

We used the Adam optimizer (Kingma and Ba 2015) for both pretraining and finetuning. The optimizer hyperparameters were set as $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We applied a learning rate warm-up for the first 10 percent of total steps, followed by a linear decay. The detailed hyperparameters for both stages can also be found in Table 1.

The pretraining phase generally concludes within three epochs, whereas the finetuning process requires less than one. Although one-epoch finetuning may appear unreasonable at first glance, we observed that Transformers are very powerful models, which may accomplish a downstream task while bypassing syntactic information. To let meaningful syntax emerge in the downstream task, we have to consider early stopping and only very few iterations are needed. Details will be presented in Section 3.4.4.

3.3 Main Results

Table 2 presents the main results of general chunking performance. In the experiment, we followed Deshmukh et al. (2021) and trained all models on CoNLL-2000 but used the English Web Treebank for further evaluation.

For comparison of inducing chunks from unsupervised parsing, we include Compound PCFG (discussed in Section 2.2) and another recent unsupervised parser based

Table 2

Main results on the linguistic annotated data.

#	Method	CoNLL-2000		English Web Treebank	
		Phrase F1	Tag Acc.	Phrase F1	Tag Acc.
Supervised Methods					
1	NLTK-tagger-chunker	83.71	89.51	–	–
2	Supervised HMM	87.68	93.99	98.62	99.44
Unsupervised Methods					
3	PMI Chunker	35.64	64.50	32.28	65.34
4	Baum–Welch HMM	25.04	58.93	24.17	58.02
5	LM Chunker	42.05	68.74	31.23	62.55
6	Compound PCFG Chunker	62.89	81.64	58.17	79.33
7	Llama2-7B Prompting	48.12	73.28	43.79	64.34
Our Pretrain Methods					
8	HRNN w/ BERT	68.12	83.90	64.32	83.25
9	HRNN w/ BART	68.57	84.04	69.71	83.08
10	HRNN w/ mBART	68.65	84.42	67.53	82.18
Our Finetune Methods					
11	HRNN w/ BART (Gigaword)	70.83	85.16	72.05	84.16
12	HRNN w/ BART (MNLI-Gen)	69.68	84.72	70.47	83.12
13	HRNN w/ mBART (WMT-14)	69.77	84.69	69.57	83.06

on the features of a pretrained language model proposed by Kim et al. (2020). The latter, called an LM Chunker, thresholds the BERT (Devlin et al. 2019) similarity of consecutive words for chunking. We observe from lines 5 and 6 in Table 2 that the LM-based unsupervised chunker is worse than the Compound PCFG in phrase F1 (42.05 vs. 62.89 on CoNLL-2000; 31.23 vs. 58.17 on English Web Treebank) and tag accuracy (68.74 vs. 81.64 on CoNLL-2000; 62.55 vs. 79.33 on English Web Treebank). Therefore, we use the chunking labels induced by Compound PCFG to pretrain our HRNN model in all experiments.

We include several traditional unsupervised chunking methods as baselines: a pointwise mutual information (PMI) chunker (Van de Cruys 2011) cuts two consecutive words if their PMI score is below a threshold; Baum–Welch HMM (Baggenstoss 2001) applies expectation–maximization to a hidden Markov model (HMM, Rabiner 1989). These methods perform significantly worse than recent advances in unsupervised syntactic structure discovery.

However, we notice that both Compound PCFG and LM chunkers (lines 4–5 in Table 3) show a significant decrease in performance when transferring to other datasets. For example, when the two chunkers are transferred from CoNLL-2000 to Gigaword, their tag accuracy is dropped by 25 percent and 37 percent, respectively. By contrast, the PMI chunker has a 7 percent drop, while the Baum–Welch HMM even has a 1 percent increase (lines 3–4). The results indicate that traditional unsupervised chunking methods are more stable, although it achieves lower performance on CoNLL-2000. Overall, unsupervised chunkers may suffer from the problem of poor transferability in addition to less satisfactory chunking performance in general.

We experimented with a baseline by prompting a large language model (LLM). Specifically, we prompted Llama2-7B (Touvron et al. 2023) to generate the chunked

Table 3

Main results of the chunking performance on the downstream datasets. The numbers with [†] indicate that the model is tested on the same datasets as the finetuning tasks.

#	Method	Gigaword		MNLI-Gen		WMT-14	
		Phrase F1	Tag Acc.	Phrase F1	Tag Acc.	Phrase F1	Tag Acc.
Supervised Methods							
1	Supervised HMM	87.54	94.37	86.81	93.40	86.49	93.50
Unsupervised Methods							
2	PMI Chunker	37.44	65.18	44.37	67.14	45.70	67.20
3	Baum-Welch HMM	38.35	54.75	29.87	32.06	34.87	45.00
4	LM Chunker	29.31	51.34	25.85	49.57	32.84	50.72
5	Compound PCFG Chunker	25.69	51.70	32.96	55.73	52.16	73.73
6	Llama2-7B Prompting	34.38	65.72	38.59	60.91	43.31	66.56
Our Pretrain Methods							
7	HRNN w/ BERT	64.20	82.61	66.38	81.60	67.27	82.70
8	HRNN w/ BART	64.60	82.24	68.81	82.04	72.09	84.87
9	HRNN w/ mBART	67.08	83.49	67.25	81.35	71.24	84.55
Our Finetune Methods							
10	HRNN w/ BART (Gigaword)	70.66[†]	85.15[†]	73.96	84.93	75.87	86.66
11	HRNN w/ BART (MNLI-Gen)	67.15	83.36	70.79 [†]	83.11 [†]	73.26	85.43
12	HRNN w/ mBART (WMT-14)	67.91	83.58	69.93	82.75	71.85 [†]	84.69 [†]

sentences (separated by “—”).⁴ Notice that our prompt has been reasonably well engineered with role-playing and exemplar techniques, and that an LLM may not be fully unsupervised because it is likely to be pretrained with linguistically annotated corpora (such as PTB). That being said, its performance (line 7 in Table 2 and line 6 in Table 3) is less satisfactory and lower than Compound PCFG. The results suggest that, although LLMs are good at generating fluent text, it remains a challenge to discover linguistic structures with LLMs.

Regarding our method, we pretrained HRNN with several Transformer encoders:

- (a) BERT (Devlin et al. 2019), an encoder-only Transformer model used in our preliminary conference paper (Deshmukh et al. 2021).
- (b) BART (Lewis et al. 2020), chosen due to its pretrained decoder, which allows HRNN to be finetuned on Gigaword and MNLI-Gen datasets for generation tasks.
- (c) mBART (Liu et al. 2020b), a multilingual variant of BART, which is needed for the translation task.

⁴ Our prompt is: <s>[INST] <<SYS>>\n You’re a helpful assistant who is an expert in linguistics. The user will ask you to separate a given sentence into phrase chunks using ‘|’. In NLP, chunking refers to the process of segmenting and labeling multi-token sequences, such as phrases or sub-sentences. For example, for the sentence ‘The cat ate deeply fried fish’, the desired answer is ‘The cat | ate | deeply fried fish’. The user wants you to directly give the answer without explanation. \n<<SYS>>\n \n Please separate the sentence ‘%s’ into phrases. [/INST].

We see from lines 8–10 in Table 2 that all the pretrained HRNN variants achieve an improvement of more than 5 percentage points in phrase F1 based on the Compound PCFG chunker on the CoNLL-2000 dataset. BART-based and mBART-based HRNN is better than the BERT-based variant not only on the standard CoNLL-2000 dataset but also in transferring its learned chunking knowledge to three downstream datasets. The large margins between the Compound PCFG chunker and our pretrained HRNN model imply that HRNN can indeed smooth out the noise of heuristics while capturing meaningful chunking patterns.⁵ More importantly, our method drastically improves the transferability of unsupervised chunkers (lines 7–9 vs. lines 2–5) in Table 3.

We then finetuned the HRNN on some downstream task for syntactic structure discovery. To measure in-domain performance on Gigaword, MNLI-Gen, and WMT-14, we used the NLTK supervised chunker (Bird and Loper 2004) as pseudo-groundtruth labels. The results are shown in Table 3. We observe an improved chunking performance across all datasets, including Gigaword, MNLI-Gen, and WMT-14. As shown in lines 11–12 (Table 3), finetuning HRNN with MNLI-Gen and WMT-14 leads to a 1–2 percent increase in phrase F1 on the in-domain datasets (70.79 vs. 68.81; 71.85 vs. 71.24). Furthermore, as seen in Table 2, it leads to improvement in the general CoNLL-2000 corpus (69.68 vs. 68.57; 69.77 vs. 68.65) and English Web Treebank corpus (70.47 vs. 69.71; 69.57 vs. 67.53).

Notably, finetuning with the summarization task yields the most significant improvements, as indicated in line 11 in Table 2 and line 10 in Table 3. It achieves nearly an increase of 6 percentage points in phrase F1 on the in-domain Gigaword dataset (70.66 vs. 64.60), as well as an increase of 5 and 4 percentage points on the out-of-domain MNLI-Gen and WMT-14 datasets (73.96 vs. 68.81; 75.87 vs. 72.09). Furthermore, it brings another 2.2-point improvement to the CoNLL-2000 dataset. The superiority of chunk induction from the summarization task may be due to the inherent connection between summarization and chunking. This is in line with a summarization method known as phrasal extractive summarization (Muresan, Tzoukermann, and Klavans 2001; Xu and Zhao 2022), where the salient information (often key phrases) is extracted to generate summary text. Future work may further analyze the connection between linguistic structures and real-world tasks, perhaps with linguistic probing methods.

It should be emphasized that our findings are different from those in the multi-task transfer learning literature. For example, Phang, Févry, and Bowman (2018) suggest that pretraining a model with intermediate tasks can help the downstream task, but both are in the supervised regime; more interestingly, Phang et al. (2020) demonstrate that such transferability holds across different languages in a multilingual model. However, our chunker is never trained with linguistic annotations, but the meaningful linguistic structure emerges with the supervision signal of the downstream task only.

3.4 Detailed Analyses

In this subsection, we conduct comprehensive experiments to verify our proposed approach. We will first show an analysis of the proposed heuristic for chunk label induction from unsupervised parsers for pretraining HRNN. Then, we will show an ablation

⁵ It is understandable that a machine learning system may smooth out the noise of training labels. For example, our HRNN achieves an F1 score of 96.67% in a supervised setting; if we randomly flip 40% of the labels, the F1 only decreases to 91.52%. Such a smoothing effect is also observed in other domains, such as text generation (Li et al. 2020; Wen et al. 2024) and syntactic parsing (Li, Mou, and Keller 2019; Shayegh et al. 2024).

Table 4

Analysis of the heuristics for inducing chunking labels from Compound PCFG.

Chunking Heuristics	Phrase F1	Tag Acc.
1-word & 2-word chunks	55.72	75.14
Maximal right branching	40.83	69.28
Maximal left branching	62.89	81.64

study on the proposed HRNN structure. Finally, we will analyze the effectiveness of finetuning HRNN on downstream text generation tasks.

3.4.1 Analysis of the Left-Branching Chunking Heuristic. We provide a detailed analysis of our maximal left-branching chunking heuristic, where we used the CoNLL-2000 dataset as a testbed because we pretrained all HRNN variants on CoNLL-2000. Without loss of generality, we conducted this analysis with BERT-based HRNN.

Table 4 shows a comparison between different heuristics for inducing chunks from parse trees. We observe that our maximal left-branching heuristic outperforms right branching by 20 points in phrase F1. Additionally, we introduce a thresholding approach that only extracts one-word and two-word chunks since most of the groundtruth chunks contain one or two words. The performance of such a heuristic is better than right-branching but worse than our left-branching approach. The findings support our hypothesis that right-branching is a common structure in English and does not suggest meaningful chunks. Conversely, left-branching identifies closely related words and is an effective heuristic for inducing chunks from parse trees.

3.4.2 Ablation Study on the HRNN Architecture. Table 5 presents an ablation study on the HRNN model. Here, we only consider pretraining HRNN from unsupervised parsers, while excluding the finetuning phase on downstream tasks. This allows us to eliminate distracting factors in this ablation study.

We experimented with various neural architectures for the chunker, which learns from Compound PCFG with the maximal left-branching heuristic. We see all variants outperform the Compound PCFG chunker (lines 2–9 vs. line 1 in Table 5), demonstrating

Table 5

Ablation study of the HRNN model.

#	Method	Phrase F1	Tag Acc.
1	Compound PCFG chunker	62.89	81.64
2	→ 1-layer HRNN	63.40	81.70
3	→ 2-layer HRNN	64.59	82.32
4	→ HRNN only	65.01	82.22
5	→ BERT tagger	66.67	83.34
6	→ BERT+1-layer HRNN	67.19	83.86
7	→ BERT+2-layer HRNN	66.53	83.34
8	→ BERT+HRNN (hard)	67.90	83.80
9	→ BERT+HRNN	68.12	83.90

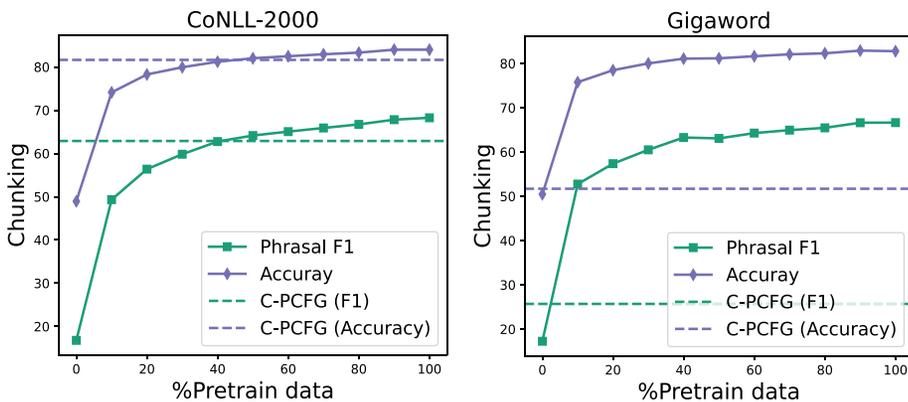


Figure 2
Analysis of the size of the pretraining dataset.

that machine learning models can effectively smooth out noise and mitigate the imperfection of chunk heuristics.

Among different neural models, we see our HRNN outperforms traditional 1-layer and 2-layer HRNNs (lines 2–4), suggesting that the HRNN’s switching gate provides useful autoregressive information during the chunking process. Moreover, the underlying BERT model provides valuable contextual information and can largely boost the performance when the number of HRNN layers is controlled (lines 2–3 vs. 6–7). Overall, our HRNN with BERT yields the highest performance, justifying our architecture design.

3.4.3 Analysis of the Size of Pretraining Data. We investigate the effect of the dataset size for pretraining our HRNN. We report the learning curves on two datasets: (1) CoNLL-2000, which is in the same domain as the pretraining data but with a different data split (test vs. train), and (2) Gigaword, which is considered an out-of-domain corpus in the pretraining stage. We incrementally increased the percentage of pretraining data from 0% to 100%, and the resulting pretrained HRNN performance is shown in Figure 2. They are compared with the chunking performance achieved by the Compound PCFG chunker (dashed lines), which is the source of our pretraining data.

For both CoNLL-2000 and Gigaword datasets, the initial increase in pretraining data size leads to a significant boost in the model’s chunking performance, but the improvement slows down when the pretraining dataset is large. It is worth noting that the HRNN outperforms the Compound PCFG on Gigaword, even with a small pretraining dataset (~10%). This highlights that, compared with heuristic chunk induction, our approach is able to improve the generalizability of unsupervised chunking.

3.4.4 Analysis of Finetuning. Chunking structure discovery with finetuning HRNN on downstream tasks requires several treatments: (1) We need to pretrain the HRNN, whose architecture has been analyzed in Section 3.4.2, and (2) We need to apply an auxiliary loss to push HRNN’s predicted chunking probability away from 0.5, as presented in Equation (15). We analyze the effect of pretraining and the auxiliary loss in this part.

Figure 3 presents the learning curve on the downstream dataset Gigaword. As seen, the HRNN with both pretraining and auxiliary loss is the only configuration that achieves an improved chunking performance (green curves with dots), whereas other

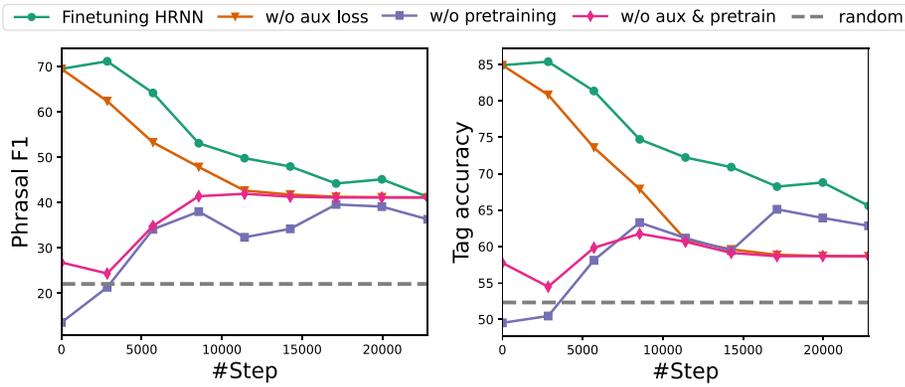


Figure 3 Ablation study of our HRNN finetuning method. We plot the learning curves in terms of phrasal F1 (top) and tag accuracy (bottom).

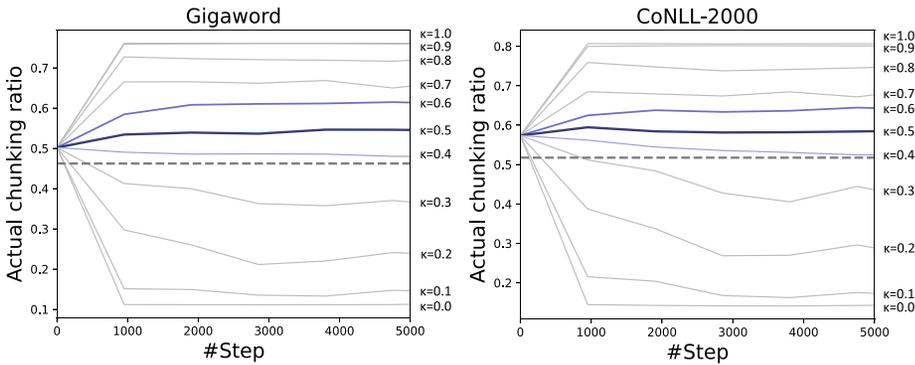


Figure 4 This analysis examines the effect of the chunking strength hyperparameter κ on the actual chunking ratio. The dashed line represents the groundtruth chunking ratio in the dataset. The purple lines indicate that the finetuned chunking performance outperforms the pretrained model. Additionally, the color depth and width of these lines indicate the ranking of the chunking performance achieved, with deeper and wider lines indicating better performance.

settings (without pretraining and/or without auxiliary loss) yield worse performance. This ablation study underscores the significance of both the pretraining strategy and auxiliary loss for HRNN’s chunking structure discovery in downstream tasks.

Moreover, the auxiliary loss provides the opportunity to control the chunking granularity. This can be shown by varying the chunking strength hyperparameter κ in Equation (15) from 0 to 1 in increments of 0.1. The results, presented in Figure 4, confirm that κ is indeed able to control the ratio of predicted cuts (i.e., chunking granularity), although the percentages may not match as there could be other factors contributing to the model’s behavior. Interestingly, closely aligning the predicted and groundtruth chunking ratios does not lead to the best performance. The optimal κ value turns out to be 0.5 in our experiments, even though the predicted ratio slightly exceeds the groundtruth ratio.

We also notice an intriguing phenomenon in Figure 3: All models converge to a phrase F1 score of approximately 40. By manually checking their predictions, we realize

that the models will eventually predict a cut for almost all steps, for example, every word is a chunk by itself. Even with our full approach (green curves), the chunking performance is only improved within limited training steps, after which it drops as well. This raises a curious question about the dynamics of the chunking structure during the process of downstream-task finetuning.

To further investigate this, we show in Figure 5 the learning curves of both chunking performance and the downstream-task performance, evaluated by Rouge1 score (Lin 2004). The learning curves show that the emerging of meaningful chunking structures is a transient phenomenon only during the early stage of downstream-task learning. The chunking performance drops after several thousand iterations, while the downstream-task performance continues to increase monotonically.

We hypothesize that chunking information (or linguistic information in general) is a meaningful abstraction of human languages. When our exquisitely designed model has not mastered the downstream NLP task well (e.g., at an early stage), it will produce structures akin to linguistically defined chunks, which serve as a convenient intermediate step for the downstream task. As training proceeds, however, such linguistically meaningful structures are abandoned by the model, as deep neural networks are known to achieve high performance with end-to-end unexplainable features (Wu et al. 2023). Our work sheds light on future research in linguistics.

Nevertheless, the drop of chunking performance can be easily alleviated by blending the pretraining and finetuning objectives. As seen in Figure 6, this treatment achieves similar peak chunking performance, which is mostly maintained during the finetuning stage. However, we do not adopt it as our main method. When we were developing our approach, our initial thought was that the pretrained chunking structures may not be perfect, so we refrained from adding the pretraining loss. Our finetuning allows the model to deviate from the pretraining structure, which unexpectedly leads to the intriguing phenomenon (i.e., the emerging and abandoning of chunking structures).

3.4.5 Effect of Pretraining on the Downstream Task. We study the effect of pretraining on the downstream text generation performance. We take summarization as the testbed and compare the pretrained HRNN with a randomly initialized one. Figure 7 shows the learning curves for both chunking and summarization performance. Interestingly, we find that the pretrained HRNN demonstrates slightly slower training progress in the text generation task than a randomly initialized model. This is understandable because our pretrained HRNN serves as a linguistic regularizer, whereas a random initialization is more flexible. The phenomenon is consistent with the explanation in Li, Mou, and Keller (2019): There are multiple local optimal latent structures that can accomplish the downstream task, but with meaningful initialization, the downstream task may recover the linguistically plausible structures (such as chunking and parsing).

3.4.6 Discussion on the Use of Labeled Validation Data. It is worth mentioning that we used labeled validation data in our main experiments. This is a common practice in previous unsupervised syntactic structure induction studies (Shen et al. 2019; Drozdov et al. 2019; Shayegh et al. 2024), with some authors claiming that such unsupervised methods are not “fully unsupervised in the strictest sense of the term” (Kim, Dyer, and Rush 2019). Our work falls in the same category.

Certain researchers believe that using labeled validation sets is problematic for unsupervised grammar induction (Shi, Livescu, and Gimpel 2020). We argue that not using a labeled validation set (e.g., using an unsupervised signal for validation) will eventually resort to the test set for hyperparameter tuning and model selection, because

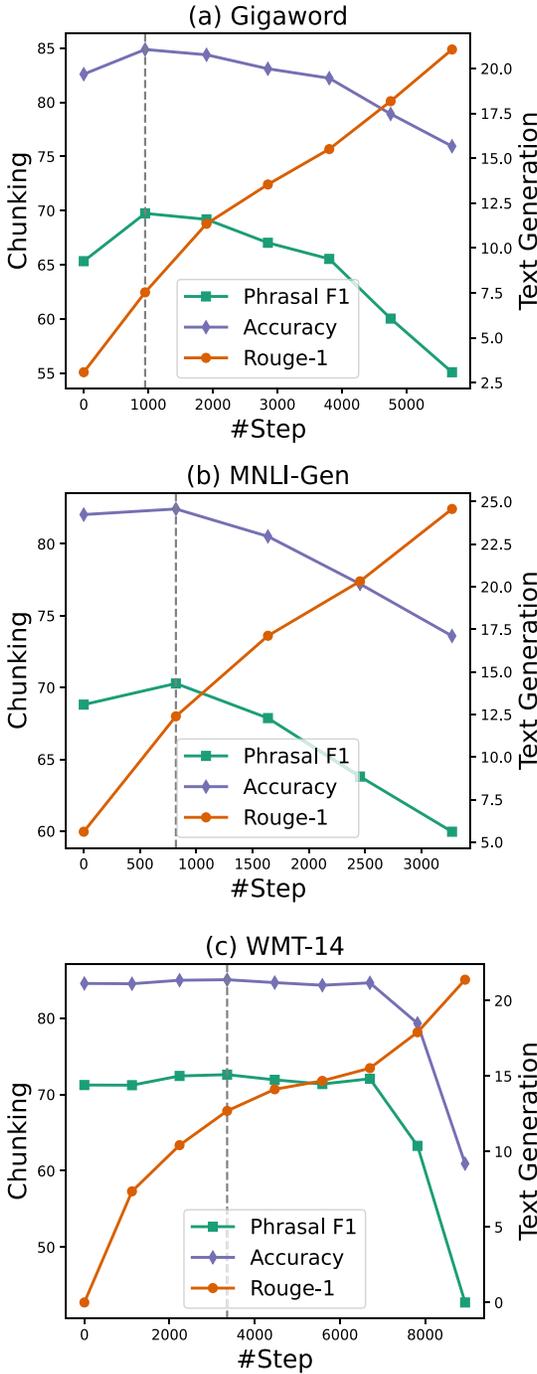


Figure 5 Learning curves of finetuning HRNN with text generation tasks. Phrasal F1 and accuracy reflect the chunking performance, and the Rouge1 score measures the text generation performance. The gray dashed lines mark the steps when HRNN achieves the highest chunking performance.

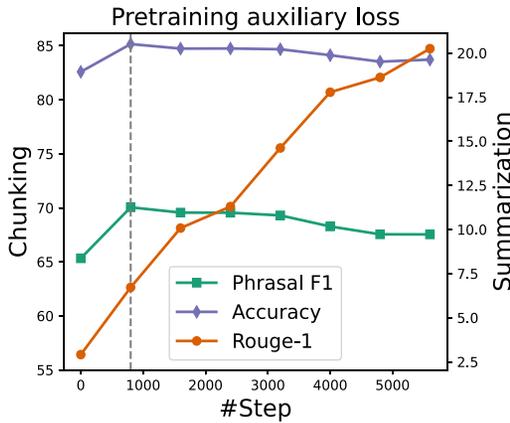


Figure 6 Using our pretraining objective as the auxiliary loss during the finetuning of the downstream summarization task.

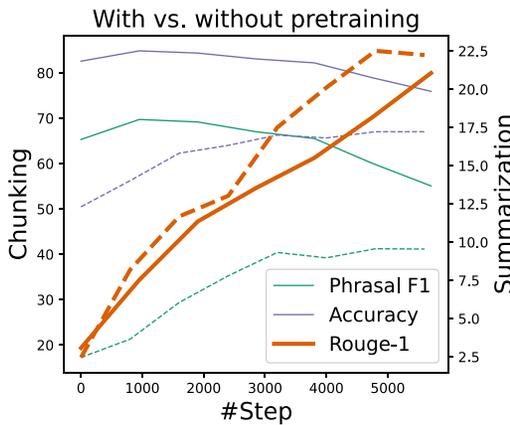


Figure 7 Dashed lines represent the performance of finetuning a randomly initialized HRNN model on the downstream summarization task. Solid lines are from our full method. The chunking performance is evaluated with the in-domain source sentences.

authors will have to repeatedly check their test performance to determine whether each of their validation strategies is effective or not.

We nevertheless compare our unsupervised approach with using some labeled validation samples for training, by splitting the provided validation set into a training set and an actual validation set (Shi, Livescu, and Gimpel 2020). In particular, we vary the number of labeled samples used, and take 70% for training and 30% for actual validation.

As shown in Figure 8, the performance of the unsupervised HRNN improves slightly when the validation set size increases from 10 to 100 samples, as the F1 score increases from 69.97 to 70.66 and accuracy increases from 84.68 to 85.25; however, the improvement plateaus beyond this point.

By contrast, the supervised HRNN does not show significant gains with an increase in training samples from 10 to 100, but its performance consistently improves as more

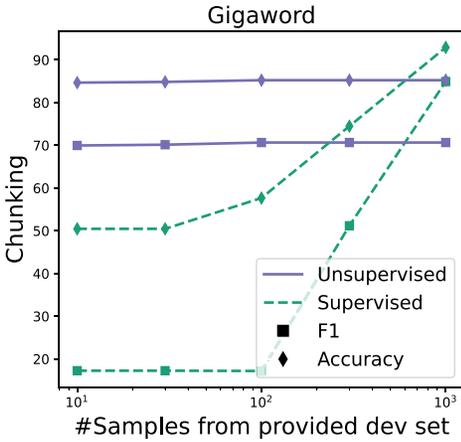


Figure 8 The impact of validation set size on the performance of unsupervised and supervised HRNNs in the Gigaword chunking task.

data are added. Notably, while the supervised HRNN still underperforms the unsupervised model at 300 samples, its performance surpasses the unsupervised counterpart when the labeled dataset size reaches 1,000 samples.

These results indicate that our unsupervised method is robust to the validation set size, and works well with an extremely small validation set such as 10 samples. This also suggests that our approach could be useful in low-resource languages where large labeled datasets are not available. In such scenarios, even experts visually checking model outputs could be sufficient for effective model selection.

3.4.7 Case Studies. In Figure 9, we show examples of chunking structures generated by Compound PCFG, pretrained HRNN, and finetuned HRNN in the summarization task.

As seen, our pretrained HRNN produces longer noun phrases than Compound PCFG, such as *Carlyle Group* and *a merchant banking concern* in Example (c), and

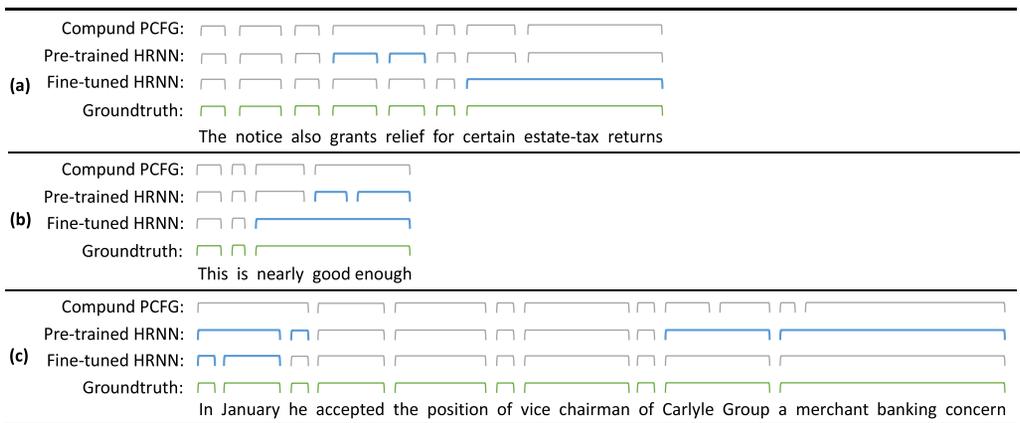


Figure 9 Examples of chunking structures generated by Compound PCFG, pretrained HRNN and finetuned HRNN are provided. Differences are highlighted in thick blue. Groundtruth chunks are also displayed in green for reference.

finetuning HRNN can further detect longer noun phrases—such as *certain estate-tax returns* in Example (a)—which are usually considered as a chunk by groundtruth labels.

Furthermore, our approach is able to correct nonsensical chunks produced by Compound PCFG. In Example (a), the two-word chunk *grants relief* is split into *grants* and *relief* as they are not in the same semantic unit. In Example (c), the chunk *In January he* is first split into *In January* and *he*, which makes more sense. Then the finetuned HRNN further splits *In January* into *In* and *January*, which agree more with the groundtruth chunks.

In general, the pretrained HRNN not only effectively learns the chunking patterns from Compound PCFG, but also can smooth out its noise. The finetuned HRNN model further improves the pretrained model and achieves better performance of unsupervised chunking.

4. Related Work

This section briefly discusses recent advances in chunking and unsupervised grammar induction approaches, followed by recent literature on weakly supervised neuro-symbolic methods.

4.1 The Chunking Task

Chunking is a prominent task in natural language processing with various applications, such as named entity recognition (Wang, Che, and Manning 2013), syntactic analysis (Zhou, Qu, and Zhang 2012), and fine-grained sentiment analysis (Yang and Cardie 2013). Chunking is also referred to as shallow parsing, where the goal is to separate a text into non-overlapping chunks that are syntactically connected (Abney 1991).

The CoNLL-2000 shared task (Tjong Kim Sang and Buchholz 2000), derived from the Penn Treebank (Marcus, Santorini, and Marcinkiewicz 1993), has provided a platform for system comparison of English chunking research. In the early stage of research, various classification models, such as support vector machines (Kudo and Matsumoto 2001) and the Winnow algorithm (Zhang, Damerau, and Johnson 2002), are applied to accomplish the chunking task. Additionally, sequential labeling models are extensively used for chunking; examples include the HMM (Molina and Pla 2002) and conditional random fields (Sha and Pereira 2003; McDonald, Crammer, and Pereira 2005). Sun et al. (2008) adopt latent-dynamic conditional random fields to the chunking task, which leads to further performance improvement.

Zhou, Qu, and Zhang (2012) propose to use chunk-level features for the chunking task, outperforming most previous research based on word-level features (Kudo and Matsumoto 2001; Zhang, Damerau, and Johnson 2002; Molina and Pla 2002; McDonald, Crammer, and Pereira 2005; Sun et al. 2008). Inspired by this, we design a hierarchical RNN that captures chunk features that are formed dynamically. We further leverage contextual word representations from a language model and enhance performance through weak supervision from downstream tasks.

4.2 Unsupervised Syntactic Structure Induction

Detecting syntactic structures without linguistic supervision has attracted significant interest due to its applications in scenarios with limited resources (Clark 2001; Klein 2005).

Unsupervised parsing is one of the most widely explored areas in syntactic structure detection. It can be broadly divided into three areas: unsupervised constituency parsing (Klein and Manning 2002; Haghighi and Klein 2006; Reichart and Rappoport 2008; Clark 2001), which aims to derive a constituency parse tree for a given sentence; unsupervised dependency parsing (Seginer 2007; Paskin 2001), which focuses on identifying the grammatical relationships between words in a sentence; and unsupervised joint parsing (Klein and Manning 2004; Shen et al. 2018), which seeks to combine both constituency and dependency parsing into a unified framework.

Among these settings, unsupervised constituency parsing has been explored most. Klein and Manning (2002) propose an approach that utilizes an expectation-maximization algorithm to model the constituency of each span. Haghighi and Klein (2006) propose using a PCFG based on manually designed features. Kim, Dyer, and Rush (2019) introduce the Compound PCFG for unsupervised parsing. On the other hand, methods have been developed to obtain tagged parse trees without supervision. Reichart and Rappoport (2008) perform clustering by syntactic features to obtain tagged parse trees. Clark (2001) clusters sequences of tags based on their local mutual information to build tagged parse trees. Overall, these early studies mostly rely on heuristics, linguistic knowledge, and manually designed features.

Unsupervised parsing has revived interest in the deep learning era. Socher et al. (2011) propose a recursive autoencoder that constructs a binary tree by greedily minimizing the reconstruction loss. Other methods for learning recursive tree structures in an unsupervised manner include CYK-style marginalization (Drozdov et al. 2019) and Gumbel-softmax (Choi, Yoo, and Lee 2018). Yogatama et al. (2017) present a shift-reduce parser learned by reinforcement learning to improve performance on a downstream task such as textual entailment. However, evidence shows the above approaches do not yield linguistically plausible trees (Williams, Drozdov, and Bowman 2018).

Li, Mou, and Keller (2019) propose to transfer knowledge between multiple unsupervised parsers to improve their performance. Our pretraining method draws inspiration from such knowledge transfer, but we introduce insightful heuristics that generate chunk labels from unsupervised parsers. Additionally, we have developed HRNN that learns from the chunk labels generated by our heuristics, and can further be finetuned with the downstream tasks for performance improvement.

Our proposed HRNN shares some similarities with the stack-LSTM (Dyer et al. 2015), as both models aim to capture hierarchical structures of language. While the stack-LSTM can generate deep tree structures, our approach is restricted to two layers and is more suited to the chunking task.

Regarding unsupervised chunking, the speech processing community has addressed the task as a component of speech detection systems and utilized acoustic information to determine the chunks (Pate and Goldwater 2011; Barrett et al. 2018). Our work only relies on textual information and views unsupervised chunking as a new task of inducing syntactic structure.

4.3 Weakly Supervised Neuro-symbolic Methods

In recent years, neuro-symbolic approaches have gained significant attention from the AI and NLP communities for interpreting deep learning models (Liu, Lu, and Mou 2023). Many of these approaches treat neuro-symbolic methods as combining deep learning and symbolic representations. For example, Lei, Barzilay, and Jaakkola (2016) extract key phrases that rationalize the neural model’s predictions on classification

tasks. Liu et al. (2018) model text classification as a sequential decision process, and propose a policy network that is able to skip reading and make decisions. Liang et al. (2017) and Mou et al. (2017) perform SQL-like execution based on input text for semantic parsing.

However, training a neuro-symbolic model is more complicated than an end-to-end neural model, because the explainable discrete latent structure lacks ground truth for direct supervision. As a result, these approaches are typically trained in a weakly supervised manner, where the training signals only exist at the end of the entire model. Reinforcement learning or its relaxation, such as Gumbel-softmax (Jang, Gu, and Poole 2017), is often used in the training procedure due to the non-differentiable latent structure. Recently, Wu et al. (2023) propose neural fuzzy logic to explain the predictions of the natural language inference model, which is trained end-to-end. Our HRNN model uses switching gates as the intermediary of the weak supervision from text generation for chunk induction, which allows it to be trained with different downstream datasets.

5. Limitations and Future Work

One limitation of this article is that we have only applied our approach to the English language, as our HRNN chunker is pretrained by unsupervised parsers (which are mainly built for English) and our maximal left-branching heuristic for chunk induction utilizes the right-branching prior of English. To generalize unsupervised grammar induction to other languages, we may develop models that incorporate language-specific prior, e.g., left branching for Japanese (Martin 2003). Another viable path is to translate a sentence into English, perform unsupervised syntactic structure prediction on the English translation, and map the structure back to the original language. Extending this, multilingual unsupervised grammar induction is also a promising research direction, since our results show that translation as a downstream task helps the discovery of syntactic structures.

We acknowledge that the current unsupervised chunking performance is still much lower than the supervised one, but with persistent research efforts, we expect it will be considerably improved in the near future. A recent study in our team addresses unsupervised parsing and has achieved close performance to supervised methods, especially for domain-shift scenarios (Shayegh et al. 2024; Shayegh, Wen, and Mou 2024).

Our long-term vision is to build an interpretable NLP system, trained in an end-to-end fashion, that can detect semantic units (chunks), determine their relationships, perform task-specific reasoning, and ultimately achieve high performance for the task of interest (Liu, Lu, and Mou 2023). The main obstacle to this ambitious goal is the difficulty of end-to-end training: This article performs chunking during text generation, but the obtained chunking structure is later abandoned by the model; another study of ours performs chunk detection and alignment by heuristics, after which a differentiable fuzzy logic mechanism is designed to perform logical reasoning (Wu et al. 2023). In future work, we plan to put these components together and train our HRNN in the NLI reasoning task; since the granularity of chunks is usually ambiguous, training an end-to-end chunker in a downstream task may yield chunks that are more suited to the task. Our goal is to make traditional NLP pipelines (Manning et al. 2014) end-to-end learnable. We anticipate that, with proper development, such a neuro-pipeline approach will achieve high performance in end tasks similar to black-box neural networks, while excelling in terms of explainability and interpretability.

6. Conclusion

In this article, we propose a framework to induce chunks in an unsupervised manner as syntactic structure discovery. Specifically, we propose HRNN with switching gates to learn from the chunk labels induced by a state-of-the-art unsupervised parser. The HRNN is further finetuned with various downstream text generation tasks to achieve better chunking results.

The experiments show that our approach largely improves unsupervised chunking. Additionally, we provide comprehensive analyses of our HRNN, the pretraining stage using unsupervised parser-induced labels, as well as the finetuning stage in a downstream task. We also point out limitations and discuss future directions.

Acknowledgments

We thank the reviewers and editors for their efforts and valuable comments. The research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) under grant no. RGPIN2020-04465, an Alberta Innovates Project, the Amii Fellow Program, the Canada CIFAR AI Chair Program, a UAHJIC project, a donation from DeepMind, and the Digital Research Alliance of Canada (alliancecan.ca).

References

- Abney, Steven P. 1991. Parsing by chunks. In *Principle-based Parsing*. Springer, pages 257–278. https://doi.org/10.1007/978-94-011-3474-3_10
- Baggenstoss, Paul M. 2001. A modified Baum-Welch algorithm for hidden Markov models with multiple observation spaces. *IEEE Transactions on Speech and Audio Processing*, pages 411–416. <https://doi.org/10.1109/89.917686>, PubMed: 36189398
- Barrett, Maria, Ana Valeria González-Garduño, Lea Frermann, and Anders Søgaard. 2018. Unsupervised induction of linguistic categories with records of reading, speaking, and writing. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2028–2038. <https://doi.org/10.18653/v1/N18-1184>
- Bies, Ann, Mott Justin, Warner Colin, and Kulick Seth. 2012. English web treebank. Linguistic Data Consortium.
- Bird, Steven and Edward Loper. 2004. NLTK: The natural language toolkit. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217. <https://doi.org/10.3115/1219044.1219075>
- Bojar, Ondřej, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amant, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. Findings of the 2014 Workshop on Statistical Machine Translation. In *Proceedings of the Workshop on Statistical Machine Translation*. <https://doi.org/10.3115/v1/W14-3302>
- Borgeaud, Sebastian, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. Improving language models by retrieving from trillions of tokens. In *Proceedings of the International Conference on Machine Learning*, pages 2206–2240.
- Choi, Jihun, Kang Min Yoo, and Sang-goo Lee. 2018. Learning to compose task-specific tree structures. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5094–5101. <https://doi.org/10.1609/aaai.v32i1.11975>
- Clark, Alexander. 2001. Unsupervised induction of stochastic context-free grammars using distributional clustering. In *Proceedings of the Workshop on Computational Natural Language Learning*. <https://doi.org/10.3115/1117822.1117831>
- Deshmukh, Anup Anand, Qianqiu Zhang, Ming Li, Jimmy Lin, and Lili Mou. 2021. Unsupervised chunking as syntactic structure induction with a knowledge-transfer approach. In *Findings of the Association for Computational Linguistics: EMNLP*, pages 3626–3634. <https://doi.org/10.18653/v1/2021.findings-emnlp.307>
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019.

- BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Drozdo, Andrew, Patrick Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. 2019. Unsupervised latent tree induction with deep inside-outside recursive auto-encoders. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1129–1141. <https://doi.org/10.18653/v1/N19-1116>
- Dyer, Chris, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. Transition-based dependency parsing with stack long short-term memory. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 334–343. <https://doi.org/10.3115/v1/P15-1033>
- Firoozeh, Nazanin, Adeline Nazarenko, Fabrice Alizon, and Béatrice Daille. 2020. Keyword extraction: Issues and methods. *Natural Language Engineering*, 26:259–291. <https://doi.org/10.1017/S1351324919000457>
- Gillenwater, Jennifer, Kuzman Ganchev, João Graça, Fernando Pereira, and Ben Taskar. 2010. Sparsity in dependency grammar induction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 194–199.
- Haghighi, Aria and Dan Klein. 2006. Prototype-driven grammar induction. In *Proceedings of the International Conference on Computational Linguistics and Annual Meeting of the Association for Computational Linguistics*, pages 881–888. <https://doi.org/10.3115/1220175.1220286>
- Izacard, Gautier, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. Few-shot learning with retrieval augmented language models. *arXiv preprint arXiv:2208.03299*.
- Jang, Eric, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*.
- Jurafsky, Daniel and James H. Martin. 2000. *Speech and Language Processing*. Prentice Hall.
- Kim, Taeuk, Jihun Choi, Daniel Edmiston, and Sang goo Lee. 2020. Are pre-trained language models aware of phrases? Simple but strong baselines for grammar induction. In *International Conference on Learning Representations*.
- Kim, Yoon, Chris Dyer, and Alexander Rush. 2019. Compound probabilistic context-free grammars for grammar induction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 2369–2385. <https://doi.org/10.18653/v1/P19-1228>
- Kingma, Diederik P. and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Klein, Dan. 2005. *The Unsupervised Learning of Natural Language Structure*. Stanford University.
- Klein, Dan and Christopher Manning. 2004. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 478–485. <https://doi.org/10.3115/1218955.1219016>
- Klein, Dan and Christopher D. Manning. 2002. A generative constituent-context model for improved grammar induction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 128–135. <https://doi.org/10.3115/1073083.1073106>
- Kudo, Taku and Yuji Matsumoto. 2001. Chunking with support vector machines. In *Proceedings of the North American Chapter of the Association for Computational Linguistics*, pages 1–8. <https://doi.org/10.3115/1073336.1073361>
- Lei, Tao, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 107–117. <https://doi.org/10.18653/v1/D16-1011>
- Lewis, Mike, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880. <https://doi.org/10.18653/v1/2020.acl-main.703>
- Li, Bowen, Lili Mou, and Frank Keller. 2019. An imitation learning approach to

- unsupervised parsing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 3485–3492. <https://doi.org/10.18653/v1/P19-1338>
- Li, Jingjing, Zichao Li, Lili Mou, Xin Jiang, Michael Lyu, and Irwin King. 2020. Unsupervised text generation by learning from search. In *Advances in Neural Information Processing Systems*, pages 10820–10831.
- Li, Zichao, Xin Jiang, Lifeng Shang, and Qun Liu. 2019. Decomposable neural paraphrase generation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 3403–3414. <https://doi.org/10.18653/v1/P19-1332>
- Liang, Chen, Jonathan Berant, Quoc Le, Kenneth D. Forbus, and Ni Lao. 2017. Neural symbolic machines: Learning semantic parsers on Freebase with weak supervision. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 23–33. <https://doi.org/10.18653/v1/P17-1003>
- Lin, Chin Yew. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- Liu, Puyuan, Chenyang Huang, and Lili Mou. 2022. Learning non-autoregressive models from search for unsupervised sentence summarization. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 7916–7929. <https://doi.org/10.18653/v1/2022.acl-long.545>
- Liu, Xianggen, Zhengdong Lu, and Lili Mou. 2023. Weakly supervised reasoning by neuro-symbolic approaches. In *Compendium of Neurosymbolic Artificial Intelligence*. IOS Press, pages 665–692. <https://doi.org/10.3233/FAIA230162>
- Liu, Xianggen, Lili Mou, Haotian Cui, Zhengdong Lu, and Sen Song. 2018. JUMPER: Learning when to make classification decisions in reading. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 4237–4243. <https://doi.org/10.24963/ijcai.2018/589>
- Liu, Xianggen, Lili Mou, Fandong Meng, Hao Zhou, Jie Zhou, and Sen Song. 2020a. Unsupervised paraphrasing by simulated annealing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 302–312. <https://doi.org/10.18653/v1/2020.acl-main.28>
- Liu, Yang. 2019. Fine-tune BERT for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Liu, Yinhan, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742. <https://doi.org/10.1162/tacl.a.00343>
- Manning, Christopher, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60. <https://doi.org/10.3115/v1/P14-5010>
- Marcus, Mitchell P., Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330. <https://doi.org/10.21236/ADA273556>
- Martin, Samuel E. 2003. *A Reference Grammar of Japanese*. University of Hawaii Press.
- McDonald, Ryan, Koby Crammer, and Fernando Pereira. 2005. Flexible text segmentation with structured multilabel classification. In *Proceedings of the Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 987–994. <https://doi.org/10.3115/1220575.1220699>
- Molina, Antonio and Ferran Pla. 2002. Shallow parsing using specialized HMMs. *Journal of Machine Learning Research*, 2:595–613.
- Mou, Lili, Zhengdong Lu, Hang Li, and Zhi Jin. 2017. Coupling distributed and symbolic execution for natural language queries. In *Proceedings of the International Conference on Machine Learning*, pages 2518–2526.
- Muresan, Smaranda, Evelyne Tzoukermann, and Judith L. Klavans. 2001. Combining linguistic and machine learning techniques for email summarization. In *Proceedings of the ACL Workshop on Computational Natural Language Learning*. <https://doi.org/10.3115/1117822.1117837>
- Niklaus, Christina, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2018. A survey on open information extraction. In *Proceedings of the International Conference on Computational Linguistics*, pages 3866–3878.
- Paskin, Mark. 2001. Grammatical bigrams. In *Advances in Neural Information Processing Systems*. <https://doi.org/10.7551/mitpress/1120.003.0016>

- Pate, John and Sharon Goldwater. 2011. Unsupervised syntactic chunking with acoustic cues: Computational models for prosodic bootstrapping. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*.
- Phang, Jason, Iacer Calixto, Phu Mon Htut, Yada Pruksachatkun, Haokun Liu, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. English intermediate-task training improves zero-shot cross-lingual transfer too. In *Proceedings of the Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 557–575. <https://doi.org/10.18653/v1/2020.aacl-main.56>
- Phang, Jason, Thibault Févry, and Samuel R. Bowman. 2018. Sentence encoders on STILTs: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- Pradhan, Sameer S., Wayne H. Ward, Kadri Hacioglu, James H. Martin, and Dan Jurafsky. 2004. Shallow semantic parsing using support vector machines. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 233–240.
- Rabiner, L. R. 1989. A tutorial on hidden Markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, pages 257–286. <https://doi.org/10.1109/5.18626>
- Ramshaw, Lance and Mitch Marcus. 1995. Text chunking using transformation-based learning. In *Workshop on Very Large Corpora*.
- Reichart, Roi and Ari Rappoport. 2008. Unsupervised induction of labeled parse trees by clustering with syntactic features. In *Proceedings of the International Conference on Computational Linguistics*, pages 721–728. <https://doi.org/10.3115/1599081.1599172>
- Rush, Alexander M., Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 379–389. <https://doi.org/10.18653/v1/D15-1044>
- Sato, Motoki, Hiroyuki Shindo, Ikuya Yamada, and Yuji Matsumoto. 2017. Segment-level neural conditional random fields for named entity recognition. In *Proceedings of the International Joint Conference on Natural Language Processing*, pages 97–102.
- Seginer, Yoav. 2007. Fast unsupervised incremental parsing. In *Proceedings of the Annual Meeting of the Association of Computational Linguistics*, pages 384–391.
- Sha, Fei and Fernando Pereira. 2003. Shallow parsing with conditional random fields. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 213–220. <https://doi.org/10.3115/1073445.1073473>
- Shayegh, Behzad, Yanshuai Cao, Xiaodan Zhu, Jackie C. K. Cheung, and Lili Mou. 2024. Ensemble distillation for unsupervised constituency parsing. In *International Conference on Learning Representations*.
- Shayegh, Behzad, Yuqiao Wen, and Lili Mou. 2024. Ensemble-based unsupervised discontinuous constituency parsing by tree averaging. *arXiv preprint arXiv:2403.00143*. <https://doi.org/10.18653/v1/2024.acl-long.808>
- Shen, Yikang, Zhouhan Lin, Chin wei Huang, and Aaron Courville. 2018. Neural language modeling by jointly learning syntax and lexicon. In *International Conference on Learning Representations*.
- Shen, Yikang, Shawn Tan, Alessandro Sordani, and Aaron Courville. 2019. Ordered neurons: Integrating tree structures into recurrent neural networks. In *International Conference on Learning Representations*.
- Shen, Yikang, Yi Tay, Che Zheng, Dara Bahri, Donald Metzler, and Aaron Courville. 2021. StructFormer: Joint unsupervised induction of dependency and constituency structure from masked language modeling. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing*, pages 7196–7209. <https://doi.org/10.18653/v1/2021.acl-long.559>
- Shi, Haoyue, Karen Livescu, and Kevin Gimpel. 2020. On the role of supervision in unsupervised constituency parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 7611–7621. <https://doi.org/10.18653/v1/2020.emnlp-main.614>
- Socher, Richard, Jeffrey Pennington, Eric H. Huang, Andrew Y. Ng, and Christopher D. Manning. 2011. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 151–161.

- Sun, Xu, Louis-Philippe Morency, Daisuke Okanohara, Yoshimasa Tsuruoka, and Jun'ichi Tsujii. 2008. Modeling latent-dynamic in shallow parsing: A latent conditional model with improved inference. In *Proceedings of the International Conference on Computational Linguistics*, pages 841–848. <https://doi.org/10.3115/1599081.1599187>
- Tjong Kim Sang, Erik F. and Sabine Buchholz. 2000. Introduction to the CoNLL-2000 shared task chunking. In *Conference on Computational Natural Language Learning and Learning Language in Logic Workshop*. <https://doi.org/10.3115/1117601.1117631>
- Touvron, Hugo, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrubti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned Chat models. *arXiv preprint arXiv:2307.09288*.
- Van de Cruys, Tim. 2011. Two multivariate generalizations of pointwise mutual information. In *Proceedings of the Workshop on Distributional Semantics and Compositionality*, pages 16–20.
- Wang, Mengqiu, Wanxiang Che, and Christopher D. Manning. 2013. Joint word alignment and bilingual named entity recognition using dual decomposition. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1073–1082.
- Wen, Yuqiao, Behzad Shayegh, Chenyang Huang, Yanshuai Cao, and Lili Mou. 2024. EBBS: An ensemble with bi-level beam search for zero-shot machine translation. *arXiv preprint arXiv:2403.00144*.
- Williams, Adina, Andrew Drozdov, and Samuel R. Bowman. 2018. Do latent tree learning models identify meaningful structure in sentences? *Transactions of the Association for Computational Linguistics*, 6:253–267. <https://doi.org/10.1162/tacl.a.00019>
- Williams, Adina, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1112–1122. <https://doi.org/10.18653/v1/N18-1101>
- Wu, Zijun, Zi Xuan Zhang, Atharva Naik, Zhijian Mei, Mauajama Firdaus, and Lili Mou. 2023. Weakly supervised explainable phrasal reasoning with neural fuzzy logic. In *International Conference on Learning Representations*.
- Xu, Wang and Tiejun Zhao. 2022. Jointly learning guidance induction and faithful summary generation via conditional variational autoencoders. In *Findings of the Association for Computational Linguistics: NAACL*, pages 2340–2350. <https://doi.org/10.18653/v1/2022.findings-naacl.180>
- Yang, Bishan and Claire Cardie. 2013. Joint inference for fine-grained opinion extraction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1640–1649.
- Yogatama, Dani, Phil Blunsom, Chris Dyer, Edward Grefenstette, and Wang Ling. 2017. Learning to compose words into sentences with reinforcement learning. In *International Conference on Learning Representations*.
- Younger, Daniel H. 1967. Recognition and parsing of context-free languages in time n^3 . *Information and Control*, 10:189–208. [https://doi.org/10.1016/S0019-9958\(67\)80007-X](https://doi.org/10.1016/S0019-9958(67)80007-X)
- Zhang, Tong, Fred Damerau, and David Johnson. 2002. Text chunking based on a generalization of Winnow. *Journal of Machine Learning Research*, pages 615–637. <https://doi.org/10.1023/A:1012498226479>
- Zhou, Junsheng, Weiguang Qu, and Fen Zhang. 2012. Exploiting chunk-level features to improve phrase chunking. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 557–567.