
Modeling Hierarchical Structures with Continuous Recursive Neural Networks

Jishnu Ray Chowdhury¹ Cornelia Caragea¹

Abstract

Recursive Neural Networks (RvNNs), which compose sequences according to their underlying hierarchical syntactic structure, have performed well in several natural language processing tasks compared to similar models without structural biases. However, traditional RvNNs are incapable of inducing the latent structure in a plain text sequence on their own. Several extensions have been proposed to overcome this limitation. Nevertheless, these extensions tend to rely on surrogate gradients or reinforcement learning at the cost of higher bias or variance. In this work, we propose Continuous Recursive Neural Network (CRvNN) as a backpropagation-friendly alternative to address the aforementioned limitations. This is done by incorporating a continuous relaxation to the induced structure. We demonstrate that CRvNN achieves strong performance in challenging synthetic tasks such as logical inference (Bowman et al., 2015b) and ListOps (Nangia & Bowman, 2018). We also show that CRvNN performs comparably or better than prior latent structure models on real-world tasks such as sentiment analysis and natural language inference.¹

1. Introduction

Constructing a sentence representation is crucial for Natural Language Processing (NLP) tasks such as natural language inference, document retrieval, and text classification. In certain contexts, methods with a bias towards composing sentences according to their underlying syntactic structures (Chomsky, 1957) have been shown to outperform comparable structure-agnostic methods (Socher et al., 2013; Tai et al., 2015; Nangia & Bowman, 2018; Choi et al., 2018; Maillard et al., 2019; Havrylov et al., 2019; Shen et al., 2019a) and

¹Computer Science, University of Illinois at Chicago, United States. Correspondence to: Jishnu Ray Chowdhury <jraych2@uic.edu>, Cornelia Caragea <cornelia@uic.edu>.

some of these structure-aware methods (Shen et al., 2019a; Qian et al., 2020) also exhibit better *systematicity* (Fodor & Pylyshyn, 1988). Notably, even contemporary Transformer-based methods (Vaswani et al., 2017) have benefited from structural biases in multiple natural language tasks (Wang et al., 2019; Fei et al., 2020).

Recursive Neural Networks (RvNNs) (Pollack, 1990; Socher et al., 2013) are capable of composing sequences in any arbitrary order as defined by the structure of the input. However, RvNNs follow a hard structure when composing a given sequence. Latent-Tree models such as Gumbel-Tree-LSTM (Choi et al., 2018), despite being capable of adaptively making structural decisions, still follow a hard structure of composition. As such, these models are forced to either make discrete structural decisions or rely on external parsers. In the former case, the model is often forced to rely on surrogate gradients² (resulting in increased bias) or reinforcement learning (resulting in increased variance). The increased bias or variance can lead to poor performance (Nangia & Bowman, 2018) unless much care is taken (Havrylov et al., 2019).

In contrast to these RvNN-based methods, chart-based parsers (Le & Zuidema, 2015; Maillard et al., 2019) can follow a “soft-structure” by recursively applying weighted average on outputs from multiple possible structures. However, these methods require the construction and maintenance of a chart of vectors, which is comparatively expensive.

Given these challenges, in this work, we propose Continuous Recursive Neural Network (CRvNN), which incorporates a continuous relaxation to RvNNs making it end-to-end differentiable. We aim for the following features:

1. **Automatic task-specific latent tree induction.** Our model is not reliant on ground truth structural labels in any form.
2. **Backpropagation friendly.** This allows our model to be easily integrated as a module within a larger neural architecture and avoid the challenges of reinforcement learning or surrogate gradients (Havrylov et al., 2019).

²We use the term “surrogate gradients” in the same sense as Martins et al. (2019) to indicate methods that approximate gradients of a discrete argmax-style mapping. Straight-through estimation with Gumbel-Softmax is such an example.

3. **Parallelism.** Unlike most prior latent tree models, CRvNN can compose multiple constituents *in parallel* if they are in the same level of hierarchy. This implies that, ideally, CRvNN can recurse over the induced tree-depth, which typically is much shorter than the sequence length. As such, CRvNN can also alleviate difficulties with gradient vanishing or gradient exploding (Hochreiter, 1991; Bengio et al., 1994), which can be caused by applying recurrence or recursion over long-length sequences.

Although chart-based parsers technically share the above features to an extent, they parallelize over multiple paths of composition instead of straightforwardly composing multiple constituents in a single path. Overall, chart-based parsers still need to recurse over the entire sequence length. In contrast, CRvNN can instead halt early based on the induced tree-depth albeit at the cost of more greediness.

2. Preliminaries

In this section, we provide a brief introduction to Recursive Neural Networks.

2.1. Recursive Neural Network (RvNN)

An RvNN uses some composition function f to recursively compose two given children into their parent:

$$p = f(x_i, x_{i+1}), \quad (1)$$

Here, we consider a sequence $x_{1:n} = (x_1, x_2, \dots, x_n)$, and a composition function f that takes two child vectors of dimension d as input and returns a parent vector of the same dimension d . The recursion terminates when it composes every node into the root parent. As an example, consider the sequence $x_{1:6}$ with an underlying hierarchical structure expressed as:

$$((x_1, (x_2, x_3)), (x_4, (x_5, x_6))).$$

This kind of structured sequence can be also expressed in terms of a tree. The RvNN operates on the input as:

$$p = f(f(x_1, f(x_2, x_3)), f(x_4, f(x_5, x_6))). \quad (2)$$

3. Our Approach

In this section, we describe the technical details of CRvNN. However, before getting into the exact details, we first construct a new but equivalent reformulation of RvNN. This will make it easier to connect CRvNN to RvNN.

3.1. Recursive Neural Network: Reformulation

We now present the aforementioned reformulation of RvNN. Below we describe different components of our reformulated RvNN individually.

3.1.1. COMPOSITION FUNCTION

The notion of composition function is the same as in §2.1.

Example. Given a composition function (or recursive cell) $f(r_1, r_2)$, if one input (r_1) represents x_1 , and another input (r_2) represents (x_2, x_3) , then the output parent will represent $(x_1, (x_2, x_3))$.

3.1.2. EXISTENTIAL PROBABILITY

In an RvNN, once two child representations are composed into a parent, the child representations themselves will no longer be needed. The unneeded child representations can be treated as “non-existent”. Based on this idea, we introduce the notion of “existential probability.”

Definition 1. We define the notion of an “existential probability” e_i to denote whether a representation r_i requires further processing ($e_i = 1$) or it should be treated as “non-existing” and hence ignored ($e_i = 0$). Every position has an existential probability, which initially is 1 for all positions. For now, we only consider binary values $e_i \in \{0, 1\}$.

Example. If r_i from position i is composed with r_j from position j , then we may update position j with the composed parent representation of (r_i, r_j) and set the existential probability of position i as 0. Typically, we can just remove the unnecessary representations. However, in this formulation, we keep them with an existential probability. This is a key change that enables us to bring a continuous relaxation to RvNNs as we will discuss later.

3.1.3. DECISION FUNCTION

Definition 2. We define a decision function D as a function dedicated to make structural decisions about which positions to compose together into a parent representation given the model state.

This definition generalizes both vanilla RvNNs and architectures such as Gumbel-Tree LSTMs. In case of vanilla RvNNs, D can be conceived of as a trivial algorithm which makes the appropriate decision at every step by simply looking at the ground truth. In case of Gumbel-Tree LSTM, the function D can represent its scoring function that scores all candidate representations to be chosen for composition in a particular step. In our generalized formulation, D makes structural decision by assigning a **composition probability** (c_i) to every position i in a sequence $r_{1:n}$. This composition probability determines which positions are to be replaced by the parent representation. For now, we only consider $c_i \in \{0, 1\}$. Note, unlike Gumbel-Tree LSTM, in our formulation multiple positions can be chosen for composition. In §3.1.5, we discuss the exact recursive rules to update the representations based on the composition probabilities.

Positions:	1	2	3	4	5	6
Iteration 1						
Sequence:	x_1	$x_2 \rightarrow$	x_3	x_4	$x_5 \rightarrow$	x_6
Existential Probability:	1	1	1	1	1	1
Composition Probability:	0	1	0	0	1	0
Iteration 2						
Sequence:	$x_1 \rightarrow$	x_2	(x_2, x_3)	$x_4 \rightarrow$	x_5	(x_5, x_6)
Existential Probability:	1	0	1	1	0	1
Composition Probability:	1	\emptyset	0	1	\emptyset	0
Iteration 3						
Sequence:	x_1	x_2	$(x_1, (x_2, x_3)) \rightarrow$	x_4	x_5	$(x_4, (x_5, x_6))$
Existential Probability:	0	0	1	0	0	1
Composition Probability:	\emptyset	\emptyset	1	\emptyset	\emptyset	0
Output after iteration 3						
Sequence:	x_1	x_2	$(x_1, (x_2, x_3))$	x_4	x_5	$((x_1, (x_2, x_3)), (x_4, (x_5, x_6)))$
Existential Probability:	0	0	0	0	0	1

Table 1. Simulation of an RvNN based on our new formulation. We use a strikethrough to illustrate “non-existing” representations (0 existential probability) as they are ignored for further computation. We use \emptyset for composition scores of non-existing representations. In each iteration, we bold the child with a composition probability of 1. We also use a \rightarrow symbol to indicate that the child will be sent over to the first right position with existential probability of 1 so that a composed parent representation can be formed at that position. The composition probabilities are predicted in each step by the decision function D .

3.1.4. LEFT AND RIGHT FUNCTIONS

Definition 3. Given a sequence of values $v_{1:n} = (v_1, v_2, \dots, v_n)$, we define a function *left* which takes as input some value v_i and returns the closest existing left value v_j (i.e., $v_j = \text{left}(v_i)$) such that $e_j = 1$ and $\forall l, j < l < i$, $e_l = 0$. Note that since position $i - 1$ can have an existential probability $e_{i-1} = 0$, $\text{left}(v_i)$ is not always v_{i-1} .

In a similar vein, we also define a function *right* which takes as input some value v_i and returns the closest existing right value v_j (i.e., $v_j = \text{right}(v_i)$) such that $e_j = 1$ and $\forall l, i < l < j$, $e_l = 0$. Similar to the *left* function, *right*(v_i) is not always v_{i+1} . Note that the sequence $v_{1:n}$ could be any sequence including a sequence of representations ($r_{1:n}$) or a sequence of composition probabilities ($c_{1:n}$).

3.1.5. RECURSIVE UPDATE

Now, we describe how the above components come together in each recursive update of our reformulated RvNN. We can formally express the rule for recursively updating some representation at position i , r_i in the recursive step k as:

$$r_i^{k+1} = \text{left}(c_i) \cdot f(\text{left}(r_i^k), r_i^k) + (1 - \text{left}(c_i)) \cdot r_i^k, \quad (3)$$

where f refers to the composition function (recursive cell) as before and c_i is the composition probability (as predicted by the function D) at any position i . Note that after this point, if $\text{left}(c_i) = 1$, $\text{left}(r_i^k)$ is no longer needed because it has been composed together with r_i^k and the composed parent representation already exists in position i (as r_i^{k+1}) according to the above rule. Thus, when $\text{left}(c_i) = c_j = 1$ we can set the existential probability (e_j) at position $\text{left}(i) = j$ as

0. Thus, we can express the update rule for e_j as:

$$e_j^{k+1} = e_j^k \cdot (1 - c_j). \quad (4)$$

In Table 1, we show how we can simulate an RvNN based on the rules as laid out above using a sequence $x_{1:6}$.³

3.2. Towards Continuous RvNN

In the above sections we consider only discrete binary values for composition probability and existential probability. As such, it is still a “hard RvNN” with a discrete order of composition. To transform it into a Continuous RvNN (soft RvNN), we simply incorporate a continuous relaxation to the decision function D allowing it to predict composition probabilities (c_i) in the interval $[0, 1]$. Similarly, we also allow existential probabilities (e_i) to be in $[0, 1]$.

As a result of this relaxation, we can simply use a neural network with sigmoid activation for D . Given that we accommodate for continuous values, we do not have to force binary decisions using some form of reparameterization or gradient approximation. Moreover, we can directly use the recursive update rules as defined in Eq. 3 and Eq. 4 without any changes because they are already compatible with real-valued c_i and e_i .

Below, we discuss our re-definition of *left* and *right* functions (neighbor retriever functions) under this new context.

³Note, in iteration 2 in Table 1, when we compose together the representation of x_1 (position 1) and the representation of (x_2, x_3) (position 3), the former is not at the immediate left (position 2) from the latter because position 2 has 0 existential probability.

Algorithm 1 Continuous Recursive Neural Network

Input: data $x_{1:n} = (x_1, x_2, \dots, x_n)$
 $r_i^1 \leftarrow \text{leafTransform}(x_i) \{i = 1, \dots, n\}$
 Initialize $e_i^1 \leftarrow 1$
for $k = 1$ **to** $n - 1$ **do**
 $c_{1:n} \leftarrow \text{decisionFunction}(r_{1:n}^k)$
 $\alpha \leftarrow \text{left}(c_i)$
 $h \leftarrow \text{left}(r_i^k)$
 $r_i^{k+1} \leftarrow \alpha \cdot \text{compose}(h, r_i^k) + (1 - \alpha) \cdot r_i^k$
 $e_i^{k+1} \leftarrow e_i^k \cdot (1 - c_i)$
 if $\text{dynamicHalt}(e_{1:n}^{k+1})$ **then**
 break
 end if
end for

3.2.1. NEIGHBOR RETRIEVER FUNCTIONS

In our reformulation presented in §3.1, an active role is played by the *left* and *right* functions. However, the previous definition (Def. 3) was made under the condition that the existential probabilities can only take binary values. Since now we have real-valued probabilities, the notion of “closest existing left (or right) value” is not well defined because existential probabilities can have non-binary values.

Precisely, both *left* and *right* functions are important for CRvNN, and either of the functions can be re-purposed as the other with minimal changes. Thus, for now, we only focus on the *right* function which can be easily adapted into a *left* function.

Definition 4. Given a sequence of values $v_{1:n} = (v_1, v_2, \dots, v_n)$, and for every v_i , given a sequence of probabilities $p_{i1:in} = (p_{i1}, p_{i2}, \dots, p_{in})$ such that p_{ij} indicates the probability that v_j is the closest existing value right to v_i , we define the function *right* as: $\text{right}(v_i) = \sum_{j=1}^n p_{ij} \cdot v_j$.

Essentially, the function *right* returns something analogous to an expected value of the immediate right existing representation. Note, however, that $\sum_j p_{ij}$ is not necessarily 1, rather ≤ 1 . This is because there is another possibility that there is no existing representation at the right at all.

In our implementation, we transform the existential probabilities $e_{1:n} = (e_1, e_2, \dots, e_n)$ into the sequence $p_{i1:in}$. Below, we formulate the precise rules that we use:

$$p_{ij} = \begin{cases} 0 & j \leq i \\ e_j & \sum_{l=i+1}^j e_l \leq 1 \\ \max(0, 1 - \sum_{l=i+1}^{j-1} e_l) & \sum_{l=i+1}^j e_l > 1 \end{cases} \quad (5)$$

That is, we re-purpose the existential probabilities e_j as p_{ij} . Naturally, we first set all non-right values from i to have 0 probability. Then we adjust them so that they sum up to ≤ 1 by zeroing out p_{ij} for any values after the position where the accumulated probability exceeds 1.

Algorithm 2 *dynamicHalt*

Input: data $e_{1:n} = (e_1, e_2, \dots, e_n)$
if $e_i < \epsilon$ **then**
 $d_i = 0 \{i = 1, \dots, n\}$
else
 $d_i = 1$
end if
if $\sum_{i=1}^n d_i = 1$ **then**
 return True
else
 return False
end if

To an extent, however, Eq. 5 is an engineering choice rather than being completely mathematically grounded. A more mathematically grounded choice would be to set $p_{ij} = e_j \cdot \prod_{l=i+1}^{j-1} (1 - e_l)$. We created a log-sum-exp based formulation for the equation (to avoid potential instabilities due to cumulative multiplication) as:

$$p_{ij} = e_j \cdot \exp \left(\sum_{l=i+1}^{j-1} \log(1 - e_l) \right) \quad (6)$$

We did some experiments with this on the synthetic logical inference (Bowman et al., 2015b) dataset. It performed reasonably but slightly worse than when we use Eq. 5. Overall both formulations (Eq. 5 or Eq. 6) should be equivalent in the discrete setting, and thus, both can push the model to approximate a discrete RvNN. Regardless, for focus, we only consider Eq. 5 in the rest of the paper.

Given these formulations we can further generalize both *left* and *right* functions.

Definition 5. Given a sequence $v_{1:n} = (v_1, v_2, \dots, v_n)$, we define $\text{left}_m(v_i)$ as the function to retrieve the expected value in the m^{th} left position from i . Formally:

$$\text{left}_1(r_i) = \text{left}(r_i), \quad (7)$$

$$\text{left}_m(r_i) = \text{left}(\text{left}_{m-1}(r_i)). \quad (8)$$

Similarly, we can define *right* _{m} . In the following subsection, we describe the main algorithm.

3.3. Continuous RvNN: Algorithm

The Continuous RvNN model is presented in Algorithm 1. We already explained the recursive update rules. Next, we describe the implementation details of all the functions in the algorithm.

3.3.1. LEAF TRANSFORMATION

We use the function *leafTransform* simply for an initial transformation of the embedded sequence. This can be

expressed as:

$$r_i^1 = LN(Wx_i + b). \quad (9)$$

Here, $x_i \in \mathbb{R}^{d_{embed} \times 1}$, $W \in \mathbb{R}^{d_h \times d_{embed}}$, and $b \in \mathbb{R}^{d_h \times 1}$. LN refers to layer normalization (Ba et al., 2016).

3.3.2. DECISION FUNCTION

The decision function is used to predict the composition probabilities c_i . For this, we take into account the local context using a convolutional layer. However, we want the “true local context” for which we need to use the $left_m$ and $right_m$ functions, as defined above. Given a locality window size of $2 \cdot \tau + 1$, we use the following function to get the initial un-normalized scores u_i for computing the composition probabilities c_i :

$$u_i = W_2 \text{GeLU} \left(\sum_{j=-\tau}^{\tau} W_{conv}^{j+\tau} L(r_i, j) + b_{conv} \right) + b_2, \quad (10)$$

$$L(r_i, j) = \begin{cases} left_{-j}(r_i) & j < 0, \\ r_i & j = 0, \\ right_j(r_i) & j > 0. \end{cases} \quad (11)$$

Here, we have a set of convolutional kernel weights $\{W_{conv}^0, W_{conv}^1, \dots, W_{conv}^{2 \cdot \tau}\}$, where any $W_{conv}^l \in \mathbb{R}^{d_h \times d_h}$; $L(r_i, j) \in \mathbb{R}^{d_h \times 1}$, $W_2 \in \mathbb{R}^{1 \times d_h}$, $b_{conv} \in \mathbb{R}^{d_h}$ and, $b_2 \in \mathbb{R}^1$. Note that u_i is now an un-normalized scalar real value. To turn it into a probability we can use *sigmoid*:

$$c_i = \text{sigmoid}(u_i) = \frac{\exp(u_i)}{\exp(u_i) + 1}. \quad (12)$$

The use of sigmoid allows multiple positions to have high composition probabilities (c_i). As such, multiple parents can be composed in the same recursion resulting in the parallelism that we alluded to before. However, pure sigmoid in this form is very unconstrained. That is, there is no constraint that prevents multiple contiguous (as defined by $left$ and $right$ functions) positions from having high composition probabilities (c_i). Concurrently composing contiguous representations with their expected right representation will violate the tree-structure. It can also cause information being propagated to positions that, at the same time, lose their own existential probability due to propagating their information rightwards. As such, some propagated information can henceforth be ignored due to landing in an area that loses its existential probability. Thus, to prevent contiguous positions from having high existential probabilities, we **modulate** the sigmoid formulation with the scores from the neighbors as follows:

$$c_i = \frac{\exp(u_i)}{\exp(u_i) + \exp(left(u_i)) + \exp(right(u_i)) + 1}. \quad (13)$$

We refer to this function as **modulated sigmoid**.

3.3.3. COMPOSITION FUNCTION

We use the same recursive gated cell as introduced by Shen et al. (2019a) for our composition function. This is originally inspired from the feedforward functions of Transformers (Vaswani et al., 2017).

$$\begin{bmatrix} z_i \\ h_i \\ c_i \\ u_i \end{bmatrix} = W_2 \text{GeLU} \left(W_1^{Cell} \begin{bmatrix} left(r_i) \\ r_i \end{bmatrix} + b_1 \right) + b_2 \quad (14)$$

$$o_i = LN(\sigma(z_i) \odot left(r_i) + \sigma(h_i) \odot r_i + \sigma(c_i) \odot u_i) \quad (15)$$

Here, σ is *sigmoid*; o_i is the output parent $\in \mathbb{R}^{d_h \times 1}$; $r_i, left(r_i) \in \mathbb{R}^{d_h \times 1}$; $W_1^{Cell} \in \mathbb{R}^{d_{cell} \times 2 \cdot d_h}$; $b_1 \in \mathbb{R}^{d_{cell} \times 1}$; $W_2 \in \mathbb{R}^{d_h \times d_{cell}}$; $b_2 \in \mathbb{R}^{d_h \times 1}$. Different from (Shen et al., 2019a), we use GeLU (Hendrycks & Gimpel, 2016) instead of ReLU as the activation function.

3.3.4. DYNAMIC HALT

Ideally, CRvNN can learn to simultaneously process all children at the same hierarchy level. Thus, CRvNN will only need to recurse over the tree depth. This means that we need a mechanism to detect if the induced tree-depth has been traversed and if we can halt. Following our framework, we only need to look at the sequence of existential probabilities $e_{1:n} = (e_1, e_2, \dots, e_n)$. Near the ideal time for halting, all existential probabilities except the last should be close to 0 (0 indicates that it is no longer in need of further processing). The last position has nothing to the right of it to be composed. So we enforce the last position to always have 0 composition probability and thus, 1 existential probability. Overall, to determine when to halt, we simply check if all existential probabilities except the last are less than some small threshold ϵ . Algorithm 2 shows an implementation of this mechanism.

3.4. Extra components

Here we describe some extra, potentially less essential, components that we use in our implementation of CRvNN.

3.4.1. SPECIAL EMBEDDINGS

We prepend all sequences with a special <START> embedding, and we append all sequences with a special <END> embedding. Both are trainable vectors of size d_{embed} . These embeddings are enforced to always have a composition probability 0 and an existential probability of 1. When constructing the local context, these embeddings can provide more explicit information about how close to the starting or ending boundary a representation r_i is.

Model	Number of Operations						Systematicity		
	7	8	9	10	11	12	A	B	C
<i>(Sentence representation models + ground truths)</i>									
Tree-LSTM*	94	92	92	88	87	86	91	84	76
Tree-Cell*	98	96	96	95	93	92	95	95	90
Tree-RNN*	98	98	97	96	95	96	94	92	86
<i>(Inter-sentence interaction models)</i>									
Transformer*	51	52	51	51	51	48	53	51	51
Universal Transformer*	51	52	51	51	51	48	53	51	51
<i>(Sentence representation models)</i>									
LSTM*	88	84	80	78	71	69	84	60	59
RRNet*	84	81	78	74	72	71	—	—	—
ON-LSTM*	91	87	85	81	78	75	70	63	60
Ordered Memory*	98 ₀	97 ₄	96 ₅	94 ₈	93 ₅	92 ₁₁	94	91	81
<i>(Our model)</i>									
CRvNN	98 ₁	97 ₃	96 ₂	95 ₆	94 ₈	93 ₅	99 ₁	98 ₃	92 ₂₂

Table 2. Accuracy on the Synthetic Logical Inference dataset (Bowman et al., 2015b) for different number of operations after training on samples with ≤ 6 operations. We also show results of testing systematicity of the models in specially designed splits: A, B, and C. * indicates that the results were taken from (Shen et al., 2019a). Our models were run 5 times on different random seeds. We show its mean and standard deviation. Specifically, subscript represents standard deviation. As an example, $90_1 = 90 \pm 0.1$

3.4.2. TRANSITION FEATURES

In an attempt to enhance the decision function, we also concatenate features to provide the explicit information if a position i was composed and updated in the last recursion or not. For that, we construct a set of features $T \in \mathbb{R}^{d_s}$ as:

$$T = \text{left}(c_i) \cdot C + (1 - \text{left}(c_i)) \cdot \widehat{C}. \quad (16)$$

Here, C and \widehat{C} are both trainable vector parameters $\in \mathbb{R}^{d_s}$. C can be conceived to be representing a set a features indicating that in the last iteration the concerned position was composed with the left child and updated, whereas \widehat{C} can be conceived to be representing the contrary. Here, we use the $\text{left}(c_i)$ values from the last recursion (initially 0).

3.4.3. HALT PENALTY

As discussed before, during halting, in the sequence of existential probabilities ($e_{1:n}$) all but the last existential probabilities should be close to 0. While this is ideally supposed to happen if all the positions are processed properly, there is no guarantee that it will indeed happen. To encourage this property, we use an auxiliary loss ($A(e_{1:n})$), which we define as:

$$A(e) = -\log \left(\frac{e_n}{\sum_{j=1}^n e_j} \right). \quad (17)$$

This can be conceived as cross entropy between the ideal final sequence of existential probabilities (all but the last being 0) and the actual final sequence of the same after normalization. The overall optimization function can be formulated as:

$$\min_{\theta} \mathcal{L}(\theta) + \gamma \cdot A(e), \quad (18)$$

where $\mathcal{L}(\theta)$ is the main cross entropy loss and γ is a Lagrangian multiplier for the auxiliary objective.

4. Experiments and Results

In this section, we discuss our experiments and results. We evaluate our model on logical inference (Bowman et al., 2015b), list operations (ListOps) (Nangia & Bowman, 2018), sentiment analysis—two datasets, SST2 and SST5 (Socher et al., 2013), and natural language inference—two datasets, SNLI (Bowman et al., 2015a) and MNLI (Williams et al., 2018b). For implementation details, refer to the appendix.

4.1. Logical Inference

In the logical inference dataset, we focus on two particular generalization properties separately - length generalization and compositional generalization (i.e., systematicity). We compare CRvNN with Tree-LSTM (Tai et al., 2015), Tree-Cell (Shen et al., 2019a) Tree-RNN (Bowman et al., 2015b), Transformer (Vaswani et al., 2017), Universal Transformer (Dehghani et al., 2019), LSTM (Hochreiter & Schmidhuber, 1997), RRNet (Jacob et al., 2018), ON-LSTM (Shen et al., 2019b), Ordered Memory (Shen et al., 2019a) (see Table 2).

4.1.1. LENGTH GENERALIZATION

To evaluate CRvNN for length generalization, as in prior work, we train the model only on samples with ≤ 6 operations whereas we test it on samples with higher unseen number of operations (≥ 7). In Table 2, discounting the performance of RvNN-based models with ground-truth access, our model, along with ordered memory, achieves the best performance on length generalization.

Model	Sequence length ranges (ListOps)							
	200 – 300	300 – 400	400 – 500	500 – 600	600 – 700	700 – 800	800 – 900	900 – 1000
CRvNN	98.51±1.1	98.46±1.3	98.04±1.3	97.95±1.1	97.17±1.6	97.84±1.7	96.94±1.6	96.78±1.9

Table 3. Extrapolation of CRvNN on ListOps after training on samples of length ≤ 100 . We used the ListOps extrapolation test set from (Havrylov et al., 2019).

Model	Accuracy
<i>(Models with ground truth)</i>	
Tree-LSTM \ddagger	98.7
<i>(Models without ground truth)</i>	
Transformer*	57.4±0.4
Universal Transformer*	71.5±7.8
LSTM \dagger	71.5±1.5
RL-SPINN \dagger	60.7±2.6
Gumbel-Tree LSTM \dagger	57.6±2.9
(Havrylov et al., 2019) \dagger	99.2 ±0.5
Ordered Memory*	99.97±0.014
<i>(Our model)</i>	
CRvNN	99.6±0.3

Table 4. Accuracy on ListOps. Results with * were taken from (Shen et al., 2019a). \ddagger indicates that the results were taken from (Nangia & Bowman, 2018). \dagger indicates that the results were taken from (Havrylov et al., 2019). Our models were run 5 times on different random seeds. We show its mean and standard deviation.

4.1.2. SYSTEMATICITY

Following Shen et al. (2019a), we create three different train-test splits on the logical inference dataset: A, B, and C (with increased level of difficulty from A to C, A being the easiest). For each split, we filter all samples with a specific compositional pattern from the training set and put them into the test set. Then we check whether our model can generalize to unseen patterns (unseen combinations of operands and operators). In set A, we filter samples according to the pattern $*(and(not\ a))*$, in B, we filter according to the pattern $*(and(not\ *))*$, and in C, we filter according to $*(\{and, or\}(not\ *))*$. As evident in Table 2, CRvNN exhibits exceptional capability for compositional generalization and outperforms all prior reported results.

4.2. ListOps

ListOps is a challenging synthetic task which explicitly demands capability for hierarchical modeling. Several prior latent-tree models were shown to perform poorly on it (Nangia & Bowman, 2018). As shown in Table 4, CRvNN gets close to perfect accuracy, demonstrating its capability to capture underlying structures without structural supervision. We also show the length generalization capability of CRvNN on ListOps in Table 3. CRvNN still achieves high accuracy in much higher sequence lengths (400 – 1000) even when it is originally trained on sequences of lengths ≤ 100 .

Model	SST2	SST5	SNLI	MNLI
RL-SPINN \ddagger	—	—	82.3	67.4
Gumbel-Tree-LSTM $\dagger\dagger$	90.7	53.7	85.6	—
Gumbel-Tree-LSTM \ddagger	—	—	83.7	69.5
Gumbel-Tree-LSTM \dagger	90.3 ₅	51.6 ₈	84.9 ₁	—
(Havrylov et al., 2019) \dagger	90.2 ₂	51.5 ₄	85.1 ₂	70.7 ₃
Ordered Memory*	90.4	52.2	—	—
CRvNN	88.3 ₆	51.4 ₁₃	85.1 ₂	72.9 ₂

Table 5. Accuracy on multiple natural language datasets. * indicates that the results were taken from (Shen et al., 2019a). \dagger indicates that the results were taken from (Havrylov et al., 2019). \ddagger indicates that the results were taken from (Williams et al., 2018a). $\dagger\dagger$ indicates that the results were taken from (Choi et al., 2018). Our models were run 5 times (except, on MNLI it was run 3 times) on different random seeds. We show its mean and standard deviation. Subscript represents standard deviation. E.g., $90_1 = 90 \pm 0.1$.

4.3. Natural Language Datasets

In Table 5, we also evaluate CRvNN on natural language datasets (SST2, SST5, SNLI, MNLI). Consistent with prior work (Havrylov et al., 2019), for MNLI we augment the training data with SNLI training data and report the test result on the matched test set. In the real-world tasks, CRvNN obtains mixed results, but overall the results are comparable to prior work in the similar latent-tree context. Particularly, CRvNN performs comparatively weak on SST2, but on contrary it performs significantly better than prior work (besides OM) on MNLI, which is a harder task. We also ran a CRvNN model with an enforced balanced tree structure induction on MNLI but its performance was comparatively worse (71.5 ± 0.4). For Gumbel-Tree LSTM Choi et al. (2018), we show results reported from different works because the results fluctuate from one paper to another (Havrylov et al., 2019; Williams et al., 2018a).

5. Analysis

In this section, we contrast CRvNN with Ordered Memory in terms of speed and show an ablation study of CRvNN.

5.1. Speed Test

Ordered memory (OM) is a close competitor for CRvNN. Both of the models are end-to-end differentiable. Both can simulate an RvNN without ground truth structures. CRvNN neither consistently nor conclusively outperforms OM. How-

Model	Sequence length ranges				
	81 – 100	101 – 200	201 – 500	501 – 700	701 – 1000
Ordered Memory	3.28 min	5.54 min	13.23 min	25 min	35.34 min
CRvNN	0.20 min	0.52 min	0.30 min	0.38 min	1.08 min

Table 6. Training time taken by Ordered Memory and CRvNN models to be run on 50 samples on various sequence length ranges. We used the publicly available code to run Ordered Memory.

Model	Logical inference			ListOps
	Number of Operations			
	10	11	12	
CRvNN	95.0 ₆	93.7 ₈	93.1 ₅	99.6 ₃
– modulated sigmoid	95.0 ₃	94.3 ₄	92.8 ₅	90.8 ₁₆₃
– transition features	95.0 ₈	94.4 ₇	92.1 ₅	99.4 ₅
– structure	87.2 ₁₁	86.1 ₁₀	80.9 ₂₀	83.4 ₁₁
– halt penalty	95.2 ₅	94.6 ₄	93.1 ₇	99.4 ₂
– GeLU + ReLU	95.1 ₇	94.1 ₅	92.7 ₆	99.2 ₄
– Cell + LSTMCell	89.4 ₇	89.2 ₇	86.5 ₁₁	71.2 ₄₁

Table 7. Cell represents Gated Recursive Cell. All the models were run 5 times on different random seeds. Subscript represents standard deviation. For example, $90_1 = 90 \pm 0.1$

Parsing Examples

(((i did) not) like) (((a single) minute) ((of this) (film .)))
 (((roger dodger) ((is (one of)) (the most)))
 (compelling (variations (of (this theme))))))

Table 8. Parsing examples obtained using CRvNN trained on MNLI. The example sentences are from (Socher et al., 2013).

ever, there is one crucial advantage for CRvNN. Particularly, it can achieve a degree of parallelism by processing multiple positions concurrently. At the same time, it can do an early-exit using dynamic halt. In contrast, OM not only has an outer loop over the whole sequence length, but it also has an inner loop where it recurse over its memory slots, adding significant overhead. The inner loop recursion of OM involves heavy sequential matrix operations. As such, in practice, CRvNN can be much faster than OM. In Table 6, we compare the training run time for both models. We generated synthetic ListOps samples for different ranges of sequence lengths. For each range of sequence lengths, we trained both the models on 50 samples for 1 epoch and 1 batch size on an AWS P3.2x instance (Nvidia V100). As we can see from the table, CRvNN is substantially faster than OM. Although in other settings with different tasks and different batch sizes, the gap between the speed of OM and CRvNN may not be as high as in Table 6, we still noticed CRvNN to be roughly 2 – 4 times faster than OM even when using higher batch sizes for OM (taking advantage of its lower memory complexity). We also tried running CYK-LSTM (Maillard et al., 2019) but faced memory issues when running it for longer sequences (> 200 length).

5.2. Ablation Study

In Table 7, we show an ablation study on CRvNN. We find that replacing modulated sigmoid with simple sigmoid significantly degrades and destabilizes the performance in ListOps. Removing “structure” or the structural bias by simply using the composition function as left-to-right recurrent network, again, significantly harms the performance of the model. The gated recursive cell itself is also crucial for the performance. Replacing it with an LSTM cell (Hochreiter & Schmidhuber, 1997) causes severe degradation. To an extent, this is consistent with the findings of Shen et al. (2019a). They replaced the gated recursive cell in ordered memory (Shen et al., 2019a) with an RNN cell, and observed substantial degradation. However, even with an LSTM cell, CRvNN performs better, in logical inference, than any of the other reported models which do not have the gated recursive cell. Replacing the activation function, removing the transition features (§3.4.2), or eliminating the halt penalty (§3.4.3) makes little difference.

5.3. Parsing Results

The internal composition scores of each layer of CRvNN can be used to parse induced trees. While there can be multiple ways to convert the composition scores to extract trees, one method is to simply treat any particular position of a given sequence at a particular iteration in the algorithm as having a composition probability of 1 whenever the cumulative composition probability at that position over all the iterations thus far is ≥ 0.5 . Otherwise, we treat the position at the particular iteration as having a composition probability of 0. Once the scores are binarized, it is straightforward to extract the trees following the ideas discussed in §3.1. Table 1 (in §3.1) shows an example on how binary and discrete composition probabilities relate to a particular structure of composition. In Table 8, we show parsing examples obtained with CRvNN using the above procedure.

6. Related Work

Several early approaches focused on adapting neural models for simulating pushdown automata or for grammar induction in general (Sun, 1990; Giles et al., 1990; Das et al., 1992; Mozer & Das, 1993; Zeng et al., 1994; Grefenstette et al., 2015). Also, recently, there are multiple works that focus on structure induction based on language modeling

objectives (Yogatama et al., 2018; Shen et al., 2018; 2019b; Li et al., 2019; Kim et al., 2019; Drozdov et al., 2019; 2020; Shen et al., 2021). In this work, we focus on models with structural bias that are used for general downstream tasks including NLP tasks such as classification and NLI.

Pollack (1990) presented RvNN as a recursive architecture to compose tree and list-like data-structures. Socher et al. (2010; 2013) aimed at capturing syntactic and semantic patterns of linguistic phrases and showed the effectiveness of RvNNs for sentiment analysis. Originally, RvNNs relied on external parsers to provide tree-structured inputs. Several works focused on augmenting RvNNs so that they can automatically induce the tree structure from plain text. To this end, Le & Zuidema (2015) used a chart-based parsing algorithm (Cocke–Younger–Kasami algorithm (Sakai, 1961)) in a neural framework with a convolutional composition function. Maillard et al. (2019) extended it with an attention mechanism to pool over multiple subtree candidates and a Tree-LSTM (Tai et al., 2015) composition function. Bowman et al. (2016) presented a stack-augmented RNN (SPINN) to simulate the functions of a Tree-RNN based on the principles of shift-reduce parsing (Schützenberger, 1963; Knuth, 1965) but relied on ground truth structure annotations for supervision. Yogatama et al. (2017) augmented SPINN with reinforcement learning (RL-SPINN) allowing unsupervised structure induction. Maillard & Clark (2018) enhanced shift-reduce parsing-based stack-RNNs with beam search. Munkhdalai & Yu (2017) introduced Neural Tree Indexers and gained strong performance by using full binary trees on natural language tasks. Similarly, Shi et al. (2018) showed that trivial trees (e.g., binary balanced trees) can be used to get competitive performances on several natural language tasks. Dong et al. (2019) proposed a neural model to learn boolean logic rules and quantifications to reason over relational data with multiple arities. Choi et al. (2018) used Gumbel Softmax (Jang et al., 2017) to adaptively make discrete decisions at every recursive step in choosing a parent to compose with a tree-LSTM. Jacob et al. (2018) proposed a combination of recursive and recurrent neural structure through reinforcement learning for hierarchical composition. Havrylov et al. (2019) extended Gumbel Tree-LSTM (Choi et al., 2018) through disjoint co-operative training of the parser (with reinforcement learning) and the composition function. Shen et al. (2019a) developed ordered memory (OM) by synthesizing the principles of ordered neurons (Shen et al., 2019b) with a stack-augmented RNN. While OM is a strong contender to CRvNN, we discussed the advantages of CRvNN as a more parallelized sequence processor in §5.1.

In another direction, Liu & Lapata (2018) applied structured attention (Kim et al., 2017) based on induced dependency trees to enhance downstream performance. Niculae et al. (2018) introduced SparseMAP to apply marginalization over

a sparsified set of latent structures whereas Corro & Titov (2019) did a Monte-Carlo estimation by stochastically sampling a projective dependency tree structure for their encoder. These methods, similar to ours, are end-to-end differentiable, but do not focus on hierarchical composition of a sequence vector (root node) from the elements (leaf nodes) in the sequence following its latent constituency structure. As such, they do not address the challenges we face here.

In recent years, Transformers (Vaswani et al., 2017) have also been extended either to better support tree structured inputs (Shiv & Quirk, 2019; Ahmed et al., 2019) or to have a better inductive bias to induce hierarchical structures by constraining self-attention (Wang et al., 2019; Nguyen et al., 2020; Shen et al., 2021) or by pushing intermediate representations to have constituent information (Fei et al., 2020). However, the fundamental capability of Transformers for composing sequences according to their latent structures in a length-generalizable manner is shown to be lacking (Tran et al., 2018; Shen et al., 2019a; Hahn, 2020).

7. Conclusion and Future Directions

We proposed a reformulation of RvNNs to allow for a continuous relaxation of its structure and order of composition. The result is CRvNN, which can dynamically induce structure within data in an end-to-end differentiable manner. One crucial difference from prior work is that it can parallelly process multiple positions at each recursive step and also, dynamically halt its computation when needed. We evaluated CRvNN on six datasets, and obtain strong performance on most of them. There are, however, several limitations of the model. First, the neighbor retriever functions in CRvNN, construct an $n \times n$ matrix (n is the sequence length) which can be memory intensive. This is similar to the memory limitations of a Transformer. Another limitation is that it is a greedy model. It is also not explicitly equipped to handle structural ambiguities in natural language. In future work, we will address these limitations. To mitigate the memory limitations, we plan to constrain the neighbor retriever functions to look at only k left or right candidates so that we only need an $n \times k$ matrix. We can compress contiguous positions with low existential probabilities so that we do not need to look beyond k . To handle its greedy nature, we will extend CRvNN to follow multiple paths concurrently.

8. Acknowledgment

We would like to sincerely thank our reviewers for their constructive feedback that helped us to improve our paper greatly. We also thank Adrian Silvescu for helpful discussions. This research is supported in part by NSF CAREER award #1802358, NSF CRI award #1823292, and an award from UIC Discovery Partners Institute. The computation for this project was performed on Amazon Web Services.

References

- Ahmed, M., Samee, M. R., and Mercer, R. E. You only need attention to traverse trees. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 316–322, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1030. URL <https://www.aclweb.org/anthology/P19-1030>.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bengio, Y., Simard, P., and Frasconi, P. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
- Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. Algorithms for hyper-parameter optimization. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, pp. 2546–2554, Red Hook, NY, USA, 2011. Curran Associates Inc. ISBN 9781618395993.
- Bergstra, J., Yamins, D., and Cox, D. D. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML’13*, pp. I-115–I-123. JMLR.org, 2013.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2015a.
- Bowman, S. R., Manning, C. D., and Potts, C. Tree-structured composition in neural networks without tree-structured architectures. In *Proceedings of the 2015th International Conference on Cognitive Computation: Integrating Neural and Symbolic Approaches - Volume 1583, COCO’15*, pp. 37–42, Aachen, DEU, 2015b. CEUR-WS.org.
- Bowman, S. R., Gauthier, J., Rastogi, A., Gupta, R., Manning, C. D., and Potts, C. A fast unified model for parsing and sentence understanding. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1466–1477, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1139. URL <https://www.aclweb.org/anthology/P16-1139>.
- Choi, J., Yoo, K. M., and Lee, S. Learning to compose task-specific tree structures. In McIlraith, S. A. and Weinberger, K. Q. (eds.), *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pp. 5094–5101. AAAI Press, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16682>.
- Chomsky, N. *Syntactic structures*. Walter de Gruyter, 1957.
- Corro, C. and Titov, I. Learning latent trees with stochastic perturbations and differentiable dynamic programming. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5508–5521, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1551. URL <https://www.aclweb.org/anthology/P19-1551>.
- Das, S., Giles, C. L., and Sun, G.-Z. Learning context-free grammars: Capabilities and limitations of a recurrent neural network with an external stack memory. In *Proceedings of The Fourteenth Annual Conference of The Cognitive Science Society*, pp. 14, 1992.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, L. Universal transformers. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HyzdRiR9Y7>.
- Dong, H., Mao, J., Lin, T., Wang, C., Li, L., and Zhou, D. Neural logic machines. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=B1xY-hRctX>.
- Drozdov, A., Verga, P., Yadav, M., Iyyer, M., and McCallum, A. Unsupervised latent tree induction with deep inside-outside recursive auto-encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 1129–1141, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1116. URL <https://www.aclweb.org/anthology/N19-1116>.
- Drozdov, A., Rongali, S., Chen, Y.-P., O’Gorman, T., Iyyer, M., and McCallum, A. Unsupervised parsing with S-DIORA: Single tree encoding for deep inside-outside recursive autoencoders. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4832–4845, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.392. URL <https://www.aclweb.org/anthology/2020.emnlp-main.392>.

- Fei, H., Ren, Y., and Ji, D. Retrofitting structure-aware transformer language model for end tasks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2151–2161, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.168. URL <https://www.aclweb.org/anthology/2020.emnlp-main.168>.
- Fodor, J. A. and Pylyshyn, Z. W. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1):3 – 71, 1988. ISSN 0010-0277. doi: [https://doi.org/10.1016/0010-0277\(88\)90031-5](https://doi.org/10.1016/0010-0277(88)90031-5). URL <http://www.sciencedirect.com/science/article/pii/0010027788900315>.
- Giles, C., Sun, G.-Z., Chen, H.-H., Lee, Y.-C., and Chen, D. Higher order recurrent networks and grammatical inference. In Touretzky, D. (ed.), *Advances in Neural Information Processing Systems*, volume 2, pp. 380–387. Morgan-Kaufmann, 1990. URL <https://proceedings.neurips.cc/paper/1989/file/cb70ab375662576bd1ac5aaf16b3fca4-Paper.pdf>.
- Grefenstette, E., Hermann, K. M., Suleyman, M., and Blunsom, P. Learning to transduce with unbounded memory. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 28, pp. 1828–1836. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/b9d487a30398d42ecff55c228ed5652b-Paper.pdf>.
- Hahn, M. Theoretical limitations of self-attention in neural sequence models. *Transactions of the Association for Computational Linguistics*, 8:156–171, 2020. doi: 10.1162/tacl_a_00306. URL https://doi.org/10.1162/tacl_a_00306.
- Havrylov, S., Kruszewski, G., and Joulin, A. Cooperative learning of disjoint syntax and semantics. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 1118–1128, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1115. URL <https://www.aclweb.org/anthology/N19-1115>.
- Hendrycks, D. and Gimpel, K. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *ArXiv*, abs/1606.08415, 2016. URL <http://arxiv.org/abs/1606.08415>.
- Hochreiter, S. Untersuchungen zu dynamischen neuronalen netzen. diploma thesis. *Diploma thesis, TU Munich*, 1991.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Jacob, A. P., Lin, Z., Sordani, A., and Bengio, Y. Learning hierarchical structures on-the-fly with a recurrent-recursive model for sequences. In *Proceedings of The Third Workshop on Representation Learning for NLP*, pp. 154–158, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-3020. URL <https://www.aclweb.org/anthology/W18-3020>.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=rkE3y85ee>.
- Kim, Y., Denton, C., Hoang, L., and Rush, A. M. Structured attention networks. *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=HkE0Nvqlg>.
- Kim, Y., Rush, A., Yu, L., Kuncoro, A., Dyer, C., and Melis, G. Unsupervised recurrent neural network grammars. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 1105–1117, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1114. URL <https://www.aclweb.org/anthology/N19-1114>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Knuth, D. E. On the translation of languages from left to right. *Information and Control*, 8(6):607 – 639, 1965. ISSN 0019-9958. doi: [https://doi.org/10.1016/S0019-9958\(65\)90426-2](https://doi.org/10.1016/S0019-9958(65)90426-2). URL <http://www.sciencedirect.com/science/article/pii/S0019995865904262>.
- Le, P. and Zuidema, W. The forest convolutional network: Compositional distributional semantics with a neural chart and without binarization. In *Proceedings*

- of the 2015 Conference on Empirical Methods in Natural Language Processing, pp. 1155–1164, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1137. URL <https://www.aclweb.org/anthology/D15-1137>.
- Li, B., Mou, L., and Keller, F. An imitation learning approach to unsupervised parsing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3485–3492, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1338. URL <https://www.aclweb.org/anthology/P19-1338>.
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgz2aEKDr>.
- Liu, Y. and Lapata, M. Learning structured text representations. *Transactions of the Association for Computational Linguistics*, 6:63–75, 2018. doi: 10.1162/tacl.a.00005. URL <https://www.aclweb.org/anthology/Q18-1005>.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Maillard, J. and Clark, S. Latent tree learning with differentiable parsers: Shift-reduce parsing and chart parsing. In *Proceedings of the Workshop on the Relevance of Linguistic Structure in Neural Architectures for NLP*, pp. 13–18, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-2903. URL <https://www.aclweb.org/anthology/W18-2903>.
- Maillard, J., Clark, S., and Yogatama, D. Jointly learning sentence embeddings and syntax with unsupervised tree-lstms. *Natural Language Engineering*, 25(4):433–449, 2019. doi: 10.1017/S1351324919000184.
- Martins, A. F. T., Mihaylova, T., Nangia, N., and Niculae, V. Latent structure models for natural language processing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pp. 1–5, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-4001. URL <https://www.aclweb.org/anthology/P19-4001>.
- Mozer, M. C. and Das, S. A connectionist symbol manipulator that discovers the structure of context-free languages. In Hanson, S., Cowan, J., and Giles, C. (eds.), *Advances in Neural Information Processing Systems*, volume 5, pp. 863–870. Morgan-Kaufmann, 1993. URL <https://proceedings.neurips.cc/paper/1992/file/45645a27c4f1adc8a7a835976064a86d-Paper.pdf>.
- Munkhdalai, T. and Yu, H. Neural tree indexers for text understanding. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pp. 11–21, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/E17-1002>.
- Nangia, N. and Bowman, S. ListOps: A diagnostic dataset for latent tree learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pp. 92–99, New Orleans, Louisiana, USA, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-4013. URL <https://www.aclweb.org/anthology/N18-4013>.
- Nguyen, X.-P., Joty, S., Hoi, S., and Socher, R. Tree-structured attention with hierarchical accumulation. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HJxK5pEYvr>.
- Niculae, V., Martins, A., Blondel, M., and Cardie, C. SparseMAP: Differentiable sparse structured inference. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3799–3808, Stockholm, Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/niculae18a.html>.
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543, 2014. URL <http://www.aclweb.org/anthology/D14-1162>.
- Pollack, J. B. Recursive distributed representations. *Artificial Intelligence*, 46(1):77 – 105, 1990. ISSN 0004-3702. doi: [https://doi.org/10.1016/0004-3702\(90\)90005-K](https://doi.org/10.1016/0004-3702(90)90005-K). URL <http://www.sciencedirect.com/science/article/pii/000437029090005K>.
- Qian, L., An, S., Lou, J.-G., Bei, C., Lin, Z., Gao, Y., Bin, Z., Nanning, Z., and Dongmei, Z. Compositional generalization by learning analytical expressions. In *Advances in Neural Information Processing Systems 33*. 2020. URL <https://papers.nips.cc/paper/2020/file/>

- 83adc9225e4deb67d7ce42d58fe5157c-Paper.pdf.
- Sakai, I. *Syntax in universal translation*. 1961.
- Schützenberger, M. On context-free languages and push-down automata. *Information and Control*, 6(3):246 – 264, 1963. ISSN 0019-9958. doi: [https://doi.org/10.1016/S0019-9958\(63\)90306-1](https://doi.org/10.1016/S0019-9958(63)90306-1). URL <http://www.sciencedirect.com/science/article/pii/S0019995863903061>.
- Shen, Y., Lin, Z., wei Huang, C., and Courville, A. Neural language modeling by jointly learning syntax and lexicon. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rkgOLb-0W>.
- Shen, Y., Tan, S., Hosseini, A., Lin, Z., Sordoni, A., and Courville, A. C. Ordered memory. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 5037–5048. Curran Associates, Inc., 2019a. URL <http://papers.nips.cc/paper/8748-ordered-memory.pdf>.
- Shen, Y., Tan, S., Sordoni, A., and Courville, A. Ordered neurons: Integrating tree structures into recurrent neural networks. In *International Conference on Learning Representations*, 2019b. URL <https://openreview.net/forum?id=B1l6qiR5F7>.
- Shen, Y., Tay, Y., Zheng, C., Bahri, D., Metzler, D., and Courville, A. Structformer: Joint unsupervised induction of dependency and constituency structure from masked language modeling, 2021. URL https://openreview.net/forum?id=Sm_4MDxPWxf.
- Shi, H., Zhou, H., Chen, J., and Li, L. On tree-based neural sentence modeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 4631–4641, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1492. URL <https://www.aclweb.org/anthology/D18-1492>.
- Shiv, V. L. and Quirk, C. Novel positional encodings to enable tree-structured transformers, 2019. URL <https://openreview.net/forum?id=SJerEhR5Km>.
- Socher, R., Manning, C. D., and Ng, A. Y. Learning continuous phrase representations and syntactic parsing with recursive neural networks. In *In Proceedings of the NIPS-2010 Deep Learning and Unsupervised Feature Learning Workshop*, 2010.
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., and Potts, C. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/D13-1170>.
- Sun, G. Connectionist pushdown automata that learn context-free grammars. In *Proceedings of International Joint Conference on Neural Networks*, volume 1, pp. 577–580, 1990.
- Tai, K. S., Socher, R., and Manning, C. D. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1556–1566, Beijing, China, July 2015. Association for Computational Linguistics. doi: 10.3115/v1/P15-1150. URL <https://www.aclweb.org/anthology/P15-1150>.
- Tran, K., Bisazza, A., and Monz, C. The importance of being recurrent for modeling hierarchical structure. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 4731–4736, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1503. URL <https://www.aclweb.org/anthology/D18-1503>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Wang, Y., Lee, H.-Y., and Chen, Y.-N. Tree transformer: Integrating tree structures into self-attention. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 1061–1070, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1098. URL <https://www.aclweb.org/anthology/D19-1098>.
- Williams, A., Drozdov, A., and Bowman, S. Do latent tree learning models identify meaningful structure in sentences? *Transactions of the Association for Computational Linguistics*, 6(0):253–267, 2018a. ISSN

2307-387X. URL <https://transacl.org/ojs/index.php/tacl/article/view/1281>.

Williams, A., Nangia, N., and Bowman, S. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122. Association for Computational Linguistics, 2018b. URL <http://aclweb.org/anthology/N18-1101>.

Wright, L. Ranger - a synergistic optimizer. <https://github.com/lessw2020/Ranger-Deep-Learning-Optimizer>, 2019.

Yogatama, D., Blunsom, P., Dyer, C., Grefenstette, E., and Ling, W. Learning to compose words into sentences with reinforcement learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Skvgqgqxe>.

Yogatama, D., Miao, Y., Melis, G., Ling, W., Kuncoro, A., Dyer, C., and Blunsom, P. Memory architectures in recurrent neural network language models. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SkFqf01AZ>.

Zeng, Z., Goodman, R. M., and Smyth, P. Discrete recurrent neural networks for grammatical inference. *IEEE Transactions on Neural Networks*, 5(2):320–330, 1994. doi: 10.1109/72.279194.

Zhang, M., Lucas, J., Ba, J., and Hinton, G. E. Lookahead optimizer: k steps forward, 1 step back. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 9597–9608. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/90fd4f88f588ae64038134f1eeaa023f-Paper.pdf>.