

Learning to Embed Sentences Using Attentive Recursive Trees

Jiaxin Shi,¹ Lei Hou,^{1*} Juanzi Li,¹ Zhiyuan Liu,¹ Hanwang Zhang²

¹Tsinghua University

²Nanyang Technological University

shijx12@gmail.com, {houlei,lijuanzi,liuzy}@tsinghua.edu.cn, hanwangzhang@ntu.edu.sg

Abstract

Sentence embedding is an effective feature representation for most deep learning-based NLP tasks. One prevailing line of methods is using recursive latent tree-structured networks to embed sentences with task-specific structures. However, existing models have no explicit mechanism to emphasize task-informative words in the tree structure. To this end, we propose an Attentive Recursive Tree model (AR-Tree), where the words are dynamically located according to their importance in the task. Specifically, we construct the latent tree for a sentence in a proposed important-first strategy, and place more attentive words nearer to the root; thus, AR-Tree can inherently emphasize important words during the bottom-up composition of the sentence embedding. We propose an end-to-end reinforced training strategy for AR-Tree, which is demonstrated to consistently outperform, or be at least comparable to, the state-of-the-art sentence embedding methods on three sentence understanding tasks.

Introduction

Along with the success of representation learning (e.g., word2vec (Mikolov et al. 2013)), sentence embedding, which maps sentences into dense real-valued vectors that represent their semantics, has received much attention. It is playing a critical role in many applications such as sentiment analysis (Socher et al. 2013), question answering (Wang and Nyberg 2015) and entailment recognition (Bowman et al. 2015).

There are three predominant approaches for constructing sentence embeddings. (1) Recurrent neural networks (RNNs) encode sentences word by word in sequential order (Dai and Le 2015; Hill, Cho, and Korhonen 2016). (2) Convolutional neural networks (CNNs) produce sentence embeddings in a bottom-up manner, moving from local n-grams to the global sentence as the receptive fields enlarge (Blunsom, Grefenstette, and Kalchbrenner 2014; Hu et al. 2014). However, the above two approaches cannot well encode linguistic composition of natural languages to some extent. (3) The last approach, on which this paper focuses, exploits tree-structured recursive neural networks (TreeRNNs) (Socher et al. 2011; 2013) to embed a sentence

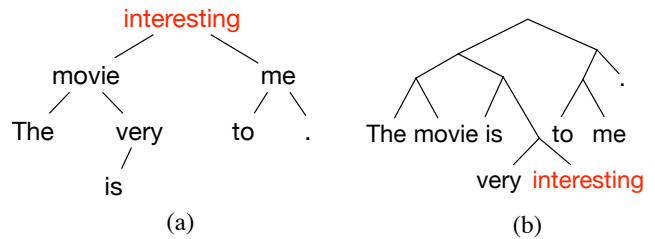


Figure 1: Two recursive trees for sentence *The movie is very interesting to me.* in the sentiment analysis task. Our AR-Tree (a) is constructed by recursively selecting the most informative word, e.g., *interesting*. However, other latent trees (b) are built by composing adjacent pairs, e.g., *very interesting* captured by (Choi, Yoo, and goo Lee 2017), which lacks the potential to emphasize words.

along its parsing tree. Tree-structured Long Short-Term Memory (Tree-LSTM) (Tai, Socher, and Manning 2015; Zhu, Sobhani, and Guo 2015) is one of the most renowned variants of TreeRNNs that is shown to be effective in learning task-specific sentence embeddings (Bowman et al. 2016).

Tree-LSTM models are motivated by the intuition that in human languages there are complicated hierarchical structures which contain rich semantics. Latent tree models (Yogatama et al. 2016; Maillard, Clark, and Yogatama 2017; Choi, Yoo, and goo Lee 2017; Williams, Drozdov, and Bowman 2017) can learn the optimal hierarchical structure, which may vary from tasks to tasks, without explicit structure annotations. The training signals to parse and embed sentences are both from certain downstream tasks. Existing models place all words in leaves equally and build the tree structure and the sentence embedding by composing adjacent node pairs bottom up (e.g., Figure 1b). This mechanism prevents the sentence embedding from focusing on the most informative words, resulting in a performance limitation on certain tasks (Shi et al. 2018).

To address this issue, we propose an Attentive Recursive Tree model (**AR-Tree**) for sentence embedding, which is a novel framework that incorporates task-specific attention mechanism into the latent tree structure learning (dos Santos et al. 2016). AR-Tree represents a sentence as a binary tree

*Corresponding author.

that contains one word in each leaf and non-leaf node, similar to the dependency parsing tree (Nivre 2003) but our AR-Tree does not depend on manual rules. To utilize the sequential information, we expect the tree’s in-order traversal preserves the word sequence, so that we can easily recover the original word sequence and obtain context of a word from its subtrees. As shown in Figure 1a, the key advantage of an AR-Tree is that those task-important words will be placed at those nodes near the root and will be naturally emphasized in tree-based embedding. This is attributed to our proposed top-down attention-first parsing strategy, inspired by easy-first parsing (Goldberg and Elhadad 2010). Specifically, we introduce a trainable scoring function to measure the word attention in a sentence with respect to a task. We greedily select the word with the highest score (e.g., *interesting*) as the root node and then recursively parse the remaining two subsequences (e.g., *The movie is* and *to me.*) to obtain two children of the parent node. After the tree construction, we embed the sentence using a modified Tree-LSTM unit (Tai, Socher, and Manning 2015; Zhu, Sobihani, and Guo 2015) in a bottom-up manner, i.e., the resultant embedding is obtained at the root node and is then applied in a downstream application. As the Tree-LSTM computes node vectors incrementally from leaf nodes to the root node, our model naturally pays more attention to those shallower words, i.e., task-informative words, meanwhile remaining advantages of the recursive semantic composition (Socher et al. 2013; Zhu, Sobihani, and Guo 2015).

Training AR-Tree is challenging due to the non-differentiability caused by the dynamic decision-making procedure. To this end, we develop a novel end-to-end training strategy based on REINFORCE algorithm (Williams 1992). To make REINFORCE work for the structure inference, we equip it with a weighted reward which is sensitive to the tree structures and a macro normalization strategy for the policy gradients.

We evaluate our model on three benchmarking tasks: textual entailment, sentiment classification, and author profiling. We show that AR-Tree outperforms previous Tree-LSTM models and is comparable to other state-of-the-art sentence embedding models. Further qualitative analyses demonstrate that AR-Tree learns reasonable task-specific attention structures.

To sum up, the contributions of our work are as follows:

- We propose Attentive Recursive Tree (AR-Tree), a Tree-LSTM based sentence embedding model, which can parse the latent tree structure dynamically and emphasize informative words inherently.
- We design a novel REINFORCE algorithm for the training of discrete tree parsing.
- We demonstrate that AR-Tree outperforms previous Tree-LSTM models and is comparable to other state-of-the-art sentence embedding models in three benchmarks.

Related Work

Latent Tree-Based Sentence Embedding. (Bowman et al. 2016) build trees and compose semantics via a generic shift-

reduce parser, whose training relies on ground-truth parsing trees. In this paper, we are interested in latent trees that dynamically parse a sentence without syntax supervision. Combination of latent tree learning with TreeRNNs has been shown as an effective approach for sentence embedding as it jointly optimizes the sentence compositions and a task-specific objective. For example, (Yogatama et al. 2016) use reinforcement learning to train a shift-reduce parser without any ground-truth.

Maillard, Clark, and Yogatama use a CYK chart parser (Cocke 1970; Younger 1967; Kasami 1965) instead of the shift-reduce parser and make it fully differentiable with the help of the softmax annealing technique. However, their model suffers from both time and space issues as the chart parser requires $\mathcal{O}(n^3)$ time and space complexity. (Choi, Yoo, and goo Lee 2017) propose an easy-first parsing strategy, which scores each adjacent node pair using a query vector and greedily combines the best pair into one parent node at each step. They use Straight-Through Gumbel-Softmax estimator (Jang, Gu, and Poole 2016) to compute parent embedding in a hard categorical gating way and enable the end-to-end training. (Williams, Drodzov, and Bowman 2017) compare above-mentioned models on several datasets and demonstrate that (Choi, Yoo, and goo Lee 2017) achieve the best performance.

Attention-Based Sentence Embedding. Attention-based methods can be divided into two categories: inter-attention (dos Santos et al. 2016; Munkhdalai and Yu 2017b), which requires a pair of sentences to attend with each other, and intra-attention (Arora, Liang, and Ma 2016; Lin et al. 2017), which does not require extra inputs except the sentence; thus the latter is more flexible than the former. (Kim et al. 2017) incorporate structural distributions into attention networks using graphical models instead of recursive trees. Note that existing latent tree-based models treat all input words equally as leaf nodes and ignore the fact that different words make varying degrees of contributions to the sentence semantics, which is nevertheless the fundamental motivation of attention mechanism. To our best knowledge, AR-Tree is the first model that generates attentive tree structures and allows the TreeRNNs to focus on more informative words for sentence embeddings.

Attentive Recursive Tree

We represent an input sentence S of N words as $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, where \mathbf{x}_i is a D_x -dimensional word embedding vector. For each sentence, we build an *Attentive Recursive Tree* (AR-Tree) where the root and nodes are denoted by R and T , respectively. Each node $t \in T$ contains one word denoted as $t.index$ ($t.index = i$ means the i -th word of input sentence) and has two children denoted by $t.left \in T$ and $t.right \in T$ (*nil* for missing cases). Following previous work (Choi, Yoo, and goo Lee 2017), we discuss binary trees in this paper and leave the n-ary case for future work. To keep the important sequential information, we guarantee that the *in-order* traversal of T corresponds to S (i.e., all nodes in t ’s left subtree must contain an index less than $t.index$). The most outstanding property of AR-Tree is

that words with more task-specific information are closer to the root.

To achieve the property, we devise a scoring function to measure the degree of importance of words, and recursively select the word with the maximum score in a top-down manner. To obtain the sentence embedding, we apply a modified Tree-LSTM to embed the nodes bottom-up, *i.e.*, from leaf to root. The resultant sentence embedding is fed into downstream tasks.

Top-Down AR-Tree Construction

We feed the input sentence into a bidirectional LSTM and obtain a context-aware hidden vector for each word:

$$\begin{aligned} \vec{\mathbf{h}}_i, \vec{\mathbf{c}}_i &= \overrightarrow{\text{LSTM}}(\mathbf{x}_i, \vec{\mathbf{h}}_{i-1}, \vec{\mathbf{c}}_{i-1}), \\ \overleftarrow{\mathbf{h}}_i, \overleftarrow{\mathbf{c}}_i &= \overleftarrow{\text{LSTM}}(\mathbf{x}_i, \overleftarrow{\mathbf{h}}_{i+1}, \overleftarrow{\mathbf{c}}_{i+1}), \\ \mathbf{h}_i &= [\vec{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_i], \\ \mathbf{c}_i &= [\vec{\mathbf{c}}_i; \overleftarrow{\mathbf{c}}_i], \end{aligned} \quad (1)$$

where \mathbf{h}, \mathbf{c} denote the hidden states and the cell states respectively. We utilize \mathbf{h} for scoring and let $S = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N\}$. Based on these context-aware word embeddings, we design a trainable scoring function to reflect the importance of each word:

$$\text{Score}(\mathbf{h}_i) = \text{MLP}(\mathbf{h}_i; \theta), \quad (2)$$

where MLP can be any multi-layer perceptron parameterized by θ . In particular, we use a 2-layer MLP with 128 hidden units and ReLU activation. Traditional tf-idf is a simple and intuitive method to reflect the degree of importance of words, however, it is not designed for specific tasks. We will use it as a baseline.

Algorithm 1 Recursive AR-Tree construction

Input: Sentence hidden vectors $S = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N\}$, beginning index b and ending index e

Output: root node $R_{S[b:e]}$ of sequence $S[b : e]$

```

procedure BUILD( $S, b, e$ )
   $R \leftarrow \text{nil}$ 
  if  $e = b$  then
     $R \leftarrow \text{new Node}$ 
     $R.\text{index} \leftarrow b$ 
     $R.\text{left}, R.\text{right} \leftarrow \text{nil}, \text{nil}$ 
  else if  $e > b$  then
     $R \leftarrow \text{new Node}$ 
     $R.\text{index} \leftarrow \text{argmax}_{i=b}^e \text{Score}(\mathbf{h}_i)$ 
     $R.\text{left} \leftarrow \text{BUILD}(S, b, R.\text{index} - 1)$ 
     $R.\text{right} \leftarrow \text{BUILD}(S, R.\text{index} + 1, e)$ 
  end if
  return  $R$ 
end procedure

```

We use a recursive top-down attention-first strategy to construct AR-Tree. Given an input sentence S and the scores for all the words, we select the word with the maximum score as the root R and recursively deal with the remaining two subsequences (before and after the selected word)

to obtain its two children. Algorithm 1 gives the procedure of constructing AR-Tree for sequence $S[b : e] = \{\mathbf{h}_b, \mathbf{h}_{b+1}, \dots, \mathbf{h}_e\}$. We can obtain the whole sentence's AR-Tree by calling $R = \text{BUILD}(S, 1, N)$ and obtain T by the traversal of all nodes. In the parsed AR-Tree, each node is most informative among its rooted subtree. Note that we do not use any extra information during the construction, thus AR-Tree is generic for any sentence embedding task.

Bottom-Up Tree-LSTM Embedding

After the AR-Tree construction, we use Tree-LSTM (Tai, Socher, and Manning 2015; Zhu, Sobihani, and Guo 2015), which introduces *cell state* into TreeRNNs to achieve better information flow, as the composition function to compute parent representation from its children and corresponding word in a bottom-up manner (*i.e.*, Figure 2). Because the original word sequence is kept in the in-order traversal of the AR-Tree, Tree-LSTM units can utilize both the sequential and the structural information to compose semantics.

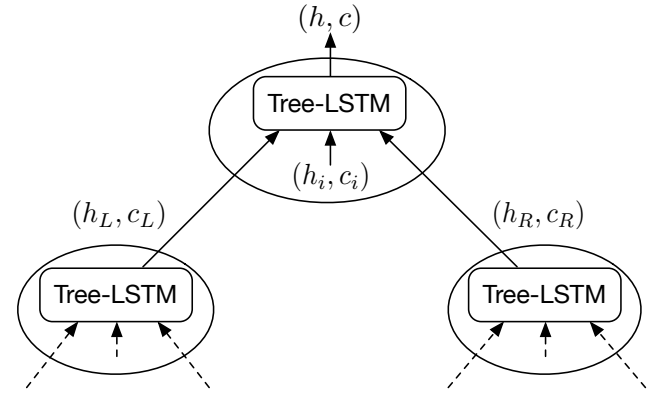


Figure 2: Our Tree-LSTM unit composes semantics of left child $(\mathbf{h}_L, \mathbf{c}_L)$, right child $(\mathbf{h}_R, \mathbf{c}_R)$ and current word $(\mathbf{h}_i, \mathbf{c}_i)$ to obtain the node embedding (\mathbf{h}, \mathbf{c}) .

The complete Tree-LSTM composition function in our model is as follows:

$$\begin{aligned} \begin{bmatrix} \mathbf{i} \\ \mathbf{f}_L \\ \mathbf{f}_R \\ \mathbf{f}_i \\ \mathbf{o} \\ \mathbf{g} \end{bmatrix} &= \begin{bmatrix} \sigma \\ \sigma \\ \sigma \\ \sigma \\ \sigma \end{bmatrix} \left(\mathbf{W}_c \begin{bmatrix} \mathbf{h}_L \\ \mathbf{h}_R \\ \mathbf{h}_i \end{bmatrix} + \mathbf{b}_c \right), \\ \mathbf{c} &= \mathbf{f}_L \odot \mathbf{c}_L + \mathbf{f}_R \odot \mathbf{c}_R + \mathbf{f}_i \odot \mathbf{c}_i + \mathbf{i} \odot \mathbf{g}, \\ \mathbf{h} &= \mathbf{o} \odot \tanh(\mathbf{c}), \end{aligned} \quad (3)$$

where $(\mathbf{h}_L, \mathbf{c}_L)$, $(\mathbf{h}_R, \mathbf{c}_R)$ and $(\mathbf{h}_i, \mathbf{c}_i)$ come from left child, right child, and bidirectional LSTM, respectively. For those nodes missing some inputs, such as the leaf nodes or nodes with only one child, we fill the missing inputs with zeros.

Finally, we use \mathbf{h} of the root R as the embedding of the sentence S and feed it into downstream tasks. The sentence embedding will focus on those informative words as they are closer to root and their semantics is emphasized naturally.

End-to-end Training Using REINFORCE

Our overall training loss \mathcal{L} combines the loss of the downstream task \mathcal{L}_{task} (e.g., the cross-entropy loss for classification tasks), the tree construction loss $\mathcal{L}_{tree} = -J(\theta)$ (discussed soon), and an L2 regularization term on all trainable parameters ϕ :

$$\mathcal{L} = \mathcal{L}_{task} + \alpha \mathcal{L}_{tree} + \lambda \|\phi\|_2^2, \quad (4)$$

where α and λ are trade-off hyperparameters.

We train the model only according to the downstream task and do not incorporate any structure supervision or pre-trained parser, leading to non-differentiability as the AR-Tree construction is a discrete decision-making process. Specifically, the scoring function $Score(\mathbf{h})$ cannot be learned in an end-to-end manner when optimizing \mathcal{L}_{task} . Inspired by (Yogatama et al. 2016), we employ reinforcement learning, whose objective corresponds to \mathcal{L}_{tree} , to train the scoring function.

We consider the construction of AR-Tree as a recursive decision-making process where each action selects a word for a node. For node $t \in T$, we define the **state** s_t as its corresponding sequence $S[b_t : e_t]$, where b_t and e_t respectively represent the index of beginning and ending position. The **action space** A_t is $\{b_t, b_t + 1, \dots, e_t\}$. We feed scores of candidate words into a softmax layer as our **policy network**, which outputs a probability distribution over the action space:

$$\pi(a_t = i | s_t; \theta) = \frac{\exp(Score(\mathbf{h}_i; \theta))}{\sum_{j=b_t}^{e_t} \exp(Score(\mathbf{h}_j; \theta))}, \quad (5)$$

where $i \in A_t$. Different from Algorithm 1 which is greedy and deterministic, at training, we construct AR-Trees randomly in terms of π , to explore more structures and bring greater gain in the long run. After the **action** a_t is sampled based on $\pi(a_t | s_t; \theta)$, the sequence is split into $S[b_t : a_t - 1]$ and $S[a_t + 1 : e_t]$, which are used respectively as two children’s states.

As for the **reward** r_t , we consider the performance metric on a downstream task. For simplicity, we discuss a classification task in this paper and leave further explorations in future work. After the whole tree T is recursively sampled based on π , we can obtain the sentence embedding following Section , feed it into the downstream classifier and obtain a predicted label. The sampled tree is considered good if the prediction is correct, and bad if the prediction is wrong. A simple rewarding strategy is to give $r_t = 1$ for all t in a good tree, and $r_t = -1$ for all t in a bad tree.

However, we consider that an early decision from a longer sequence has a greater impact on the tree structure (e.g., the selection of root is more important than leaves). So we multiply the reward value by $|A_t|$, i.e., $r_t = |A_t|$ for all t in a good tree and $r_t = -|A_t|$ for all t in a bad tree.

We use REINFORCE (Williams 1992), a widely-used policy gradient method in reinforcement learning, to learn parameters of the policy network. The goal of the learning algorithm is to maximize the expected long-term reward:

$$J(\theta) = \mathbb{E}_{\pi(a_t | s_t; \theta)} r_t. \quad (6)$$

Following (Sutton et al. 2000), the gradient w.r.t. the parameters of policy network can be derived as:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi} [r_t \nabla_{\theta} \log \pi(a_t | s_t; \theta)]. \quad (7)$$

It is prohibitively expensive to calculate the accurate $\nabla_{\theta} J(\theta)$ by iterating over all possible trees. Following (Yu et al. 2017), we apply Monte Carlo search to estimate the expectation. Specifically, we sample M trees for sentence S , denoted as T_1, \dots, T_M , each containing N nodes. Then we can simplify $\nabla_{\theta} J(\theta)$ by averaging rewards among all these $M \times N$ nodes (micro average):

$$\nabla_{\theta} J(\theta) = \frac{1}{MN} \sum_{k=1}^M \sum_{t \in T_k} r_t \nabla_{\theta} \log \pi(a_t | s_t; \theta). \quad (8)$$

However, we observed that frequent words (e.g., *the, is*) were assigned high scores if we used Formula 8 to train θ . We think the reason is that the scoring function takes one single word embedding as input, meaning that frequent words will contribute more to its training gradient if rewards are averaged among all tree nodes, which is harmful to the score estimation of low-frequency words.

To eliminate the influence caused by the word frequency, we integrate B input sentences as a mini-batch, sample M trees for each of them, and normalize the gradient in word-level (macro average) rather than node-level:

$$\nabla_{\theta} J(\theta) = \frac{1}{|W|} \sum_{w \in W} \frac{1}{|T^w|} \sum_{t \in T^w} r_t \nabla_{\theta} \log \pi(a_t | s_t; \theta), \quad (9)$$

where W represents all words of the mini-batch, T^w represents all nodes whose selected word is w in all $B \times M$ sampled trees. Figure 3 gives an example.

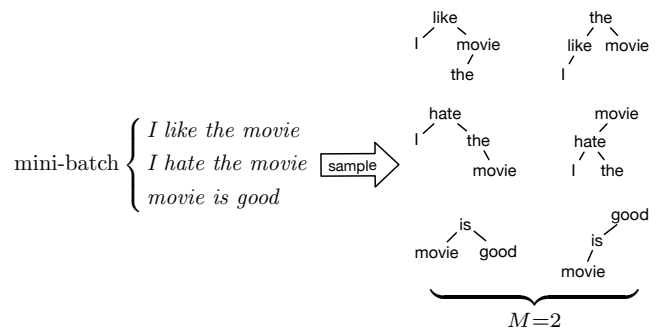


Figure 3: Sampled results of a mini-batch. We have three sentences in a mini-batch ($B = 3$) and for each sentence we sample two trees ($M = 2$). Totally we get six sampled trees. Micro average is to average gradients over all nodes of these six trees. Macro average is to first average gradients over nodes of the same word (e.g., average over 6 nodes containing *movie* to get the gradient of *movie*), and then average gradients of these words to obtain the final training signals.

Experiments

We evaluate the proposed AR-Tree on three tasks: natural language inference, sentence sentiment analysis, and author profiling.

Experiment	D_x	D_h	D_c	Finetune	Dropout	Bn	Batch size	M	Optimizer
SNLI-100D	100	100	200	✓	0.1	✓	128	2	Adam (Kingma and Ba 2014)
SNLI-300D	300	300	1024		0.1	✓	128	2	Adam
SST-2	300	300	300	✓	0.5		32	3	Adadelta (Zeiler 2012)
SST-5	300	300	1024	✓	0.5		64	3	Adadelta
Age prediction	300	600	2000	✓	0.3		50	3	Adam

Table 1: Experimental settings. Dropout: dropout probability. Bn: whether using batch normalization.

Model	# params.	Acc. (%)
100D Latent Syntax Tree-LSTM (Yogatama et al. 2016)	500k	80.5
100D CYK Tree-LSTM (Maillard, Clark, and Yogatama 2017)	231k	81.6
100D Gumbel Tree-LSTM (Choi, Yoo, and goo Lee 2017)	262k	82.6
100D Tf-idf Tree-LSTM (Ours)	343k	82.3
100D AR-Tree (Ours)	356k	82.8
300D SPINN (Bowman et al. 2016)	3.7m	83.2
300D NSE (Munkhdalai and Yu 2017a)	6.3m	84.8
300D NTI-SLSTM-LSTM (Munkhdalai and Yu 2017b)	4.0m	83.4
300D Gumbel Tree-LSTM (Choi, Yoo, and goo Lee 2017)	2.9m	85.0
300D Self-Attentive (Lin et al. 2017)	4.1m	84.4
300D Tf-idf Tree-LSTM (Ours)	3.5m	84.5
300D AR-Tree (Ours)	3.6m	85.5
600D Gated-Attention BiLSTM (Chen et al. 2017)	11.6m	85.5
300D Decomposable attention (Parikh et al. 2016)	582k	86.8
300D NTI-SLSTM-LSTM global attention (Munkhdalai and Yu 2017b)	3.2m	87.3
300D Structured Attention (Kim et al. 2017)	2.4m	86.8

Table 2: Test accuracy and the number of parameters (excluding word embeddings) on the SNLI dataset. The above two sections list results of Tree-LSTM and other baseline models grouped by the dimension. The bottom section contains state-of-the-art inter-attention models on SNLI dataset.

We set $\alpha = 0.1$, $\lambda = 1e - 5$ in Eq. 4 through all experiments. For fair comparisons, we followed the experimental settings in (Choi, Yoo, and goo Lee 2017) on language inference and sentence sentiment analysis. For the author profiling task whose dataset is provided by (Lin et al. 2017), we followed their settings by contacting the authors. We considered their model, which is self-attentive but without tree structures, as a baseline, to show the effect of latent trees. We conducted *Tf-idf Tree-LSTM* experiment, which replaces the scoring function with tf-idf value while retaining all other settings, as one of our baselines. For all experiments, we saved the model that performed best on the validation set as our final model and evaluated it on the test set. The implementation is made publicly available.¹

Natural Language Inference

The natural language inference is a task of predicting the semantic relationship between two sentences, a premise, and a hypothesis. We evaluated our model using the Stanford Natural Language Inference corpus (SNLI; (Bowman et al. 2015)), which aims to predict whether two sentences are *entailment*, *contradiction*, or *neutral*. SNLI consists of 549,367/9,842/9,824 premise-hypothesis pairs for train/validation/test sets respectively.

Following (Bowman et al. 2016; Mou et al. 2016), we ran AR-Tree separately on two input sentences to obtain their embeddings \mathbf{h}_{pre} and \mathbf{h}_{hyp} . Then we constructed a feature

vector \mathbf{v} for the pair by the following equation:

$$\mathbf{v} = \begin{bmatrix} \mathbf{h}_{pre} \\ \mathbf{h}_{hyp} \\ |\mathbf{h}_{pre} - \mathbf{h}_{hyp}| \\ \mathbf{h}_{pre} \odot \mathbf{h}_{hyp} \end{bmatrix}, \quad (10)$$

and fed the feature into a neural network, *i.e.*, a multi-layer perceptron (MLP) which has a D_c -dimensional hidden layer with *ReLU* activation function and a *softmax* layer.

We conducted SNLI experiments with two settings: 100D ($D_x = 100$) and 300D ($D_x = 300$). In both experiments, we initialized the word embedding matrix with GloVe pre-trained vectors (Pennington, Socher, and Manning 2014), added a dropout (Srivastava et al. 2014) after the word embedding layer and added batch normalization layers (Ioffe and Szegedy 2015) followed by dropout to the input and the output of the MLP. Details can be found in Table 1. The training on an NVIDIA GTX1080 Ti needs about 30 hours, slower than Gumbel Tree-LSTM (Choi, Yoo, and goo Lee 2017) because our tree construction is implemented for every single sentence instead of the whole batch.

Table 2 summarizes the results. We can see that our 100D and 300D models perform best among the Tree-LSTM models. State-of-the-art inter-attention models get the highest performance on SNLI, because they incorporate inter-information between the sentence pairs to boost the performance. However, inter-attention is limited for paired inputs and lacks flexibility. Our 300D model outperforms self-attentive (Lin et al. 2017), the state-of-the-art intra-attention

¹<https://github.com/shijx12/AR-Tree>

Model	SST-2 (%)	SST-5 (%)
LSTM (Tai, Socher, and Manning 2015)	84.9	46.4
Bidirectional LSTM (Tai, Socher, and Manning 2015)	87.5	49.1
RNTN (Socher et al. 2013)	85.4	45.7
DMN (Kumar et al. 2016)	88.6	52.1
NSE (Munkhdalai and Yu 2017a)	89.7	52.8
BCN+Char+CoVe (McCann et al. 2017)	90.3	53.7
byte-mLSTM (Radford, Jozefowicz, and Sutskever 2017)	91.8	52.9
Constituency Tree-LSTM (Tai, Socher, and Manning 2015)	88.0	51.0
Latent Syntax Tree-LSTM (Yogatama et al. 2016)	86.5	-
NTI-SLSTM-LSTM (Munkhdalai and Yu 2017b)	89.3	53.1
Gumble Tree-LSTM (Choi, Yoo, and goo Lee 2017)	90.1	52.5
Tf-idf Tree-LSTM (Ours)	88.9	51.3
AR-Tree (Ours)	90.4	52.7

Table 3: Results of SST experiments. The bottom section contains results of Tree-LSTM models and the top section contains other baseline and state-of-the-art models.

model, by 1.1%, demonstrating its effectiveness.

Sentiment Analysis

We used Stanford Sentiment Treebank (SST) (Socher et al. 2013) to evaluate the performance of our model. The sentences in SST dataset are parsed into binary trees with the Stanford parser, and each subtree corresponding to a phrase is annotated with a sentiment score. It includes two subtasks: SST-5, classifying each phrase into 5 classes, and SST-2, preserving only 2 classes.

Following (Choi, Yoo, and goo Lee 2017), we used all phrases for training but only the entire sentences for evaluation. We used an MLP with a D_c -dimensional hidden layer as the classifier. For both SST-2 and SST-5, we initialized the word embeddings with GloVe 300D pre-trained vectors, and added dropout to the word embedding layer and the input and the output of the MLP. Table 1 lists the parameter details.

Table 3 shows the results of SST experiments. Our model on SST-2 outperforms all Tree-LSTM models and other state-of-the-art models except Byte-mLSTM (Radford, Jozefowicz, and Sutskever 2017), a byte-level language model trained on a very large corpus. (McCann et al. 2017) obtains the highest performance on SST-5 due to the help of pretraining and character n-gram embeddings. Without the help of character-level information, our model can still get comparable results on SST-5.

Author Profiling

The Author Profiling dataset consists of Twitter tweets and some annotations about age and gender of the user writing the tweet. Following (Lin et al. 2017) we used English tweets as input to predict the age range of the user, including 5 classes: 18-24, 25-34, 35-49, 50-64 and 65+. The age prediction dataset consists of 68,485/4,000/4,000 tweets for train/validation/test sets.

We applied GloVe and dropout as in the SST experiments. Table 1 describes detailed settings, which are the same as (Lin et al. 2017)’s published implementation except for the optimizer (they use SGD but we find Adam converges better).

Model	Acc. (%)
BiLSTM+MaxPooling (Lin et al. 2017)	77.40
CNN+MaxPooling (Lin et al. 2017)	78.15
Gumble Tree-LSTM (Choi, Yoo, and goo Lee 2017)	80.23
Self-Attentive (Lin et al. 2017)	80.45
Tf-idf Tree-LSTM (Ours)	80.20
AR-Tree (Ours)	80.85

Table 4: Results of age prediction experiments.

Results of the age prediction experiments are shown in Table 4. We can see that our model outperforms all other baseline models. Compared to self-attentive model, our AR-Tree model obtains higher performance in the same experimental settings, indicating that latent structures are helpful to sentence understanding.

Qualitative Analysis

We conducted experiments to observe structures of the learned trees. We select 2 sentences from the test set of three experiment datasets respectively and show their attentive trees in Figure 4.

The left column is a sentence pair with relationship *contradiction* from SNLI. Figure 4a and 4b both focus on the predicate word *chased* firstly, then focus on its subject and object respectively. The middle column is from SST-2, the sentiment analysis dataset. Both Figure 4c and 4d focus on emotional adjectives such as *embarrassing*, *amusing* and *enjoyable*. The right column is from the age prediction dataset, predicting the author’s age based on the tweet. Figure 4e attends to *@Safety_Ist*, a baby production, indicating that the author is probably a young parent. Figure 4f focuses on *lecturers* which suggests that the author is likely to be a college student.

Furthermore, we applied parsers trained on different tasks to the same sentence, and show results in Figure 5. The parser of SNLI focuses on *partially* (Figure 5a), as SNLI is an inference dataset and pays more attention to words which may be different in two sentences to reflect the contradiction relationship (e.g., *partially* v.s. *totally*). The parser of SST-2, the sentiment classification task, focuses on sentimental

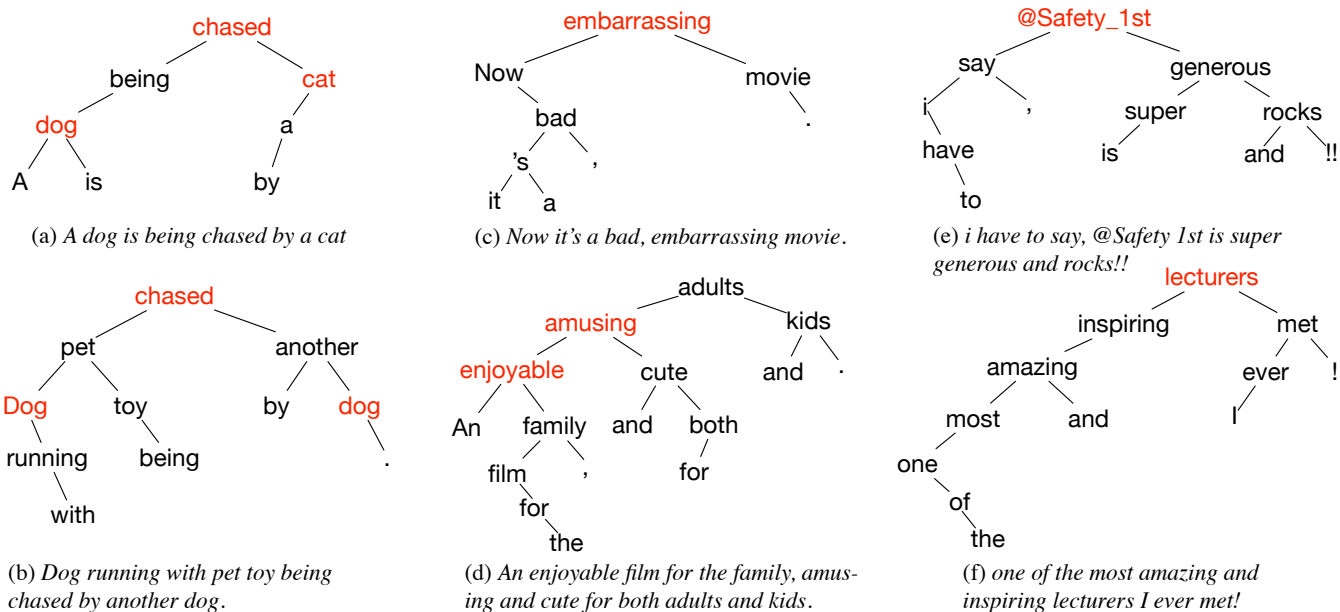


Figure 4: Examples of our produced attentive trees. The caption of each subfigure is the input sentence. The left, middle and right columns are from SNLI, SST-2 and age prediction respectively. We can see that our AR-Tree can place task-informative words at shallow nodes.

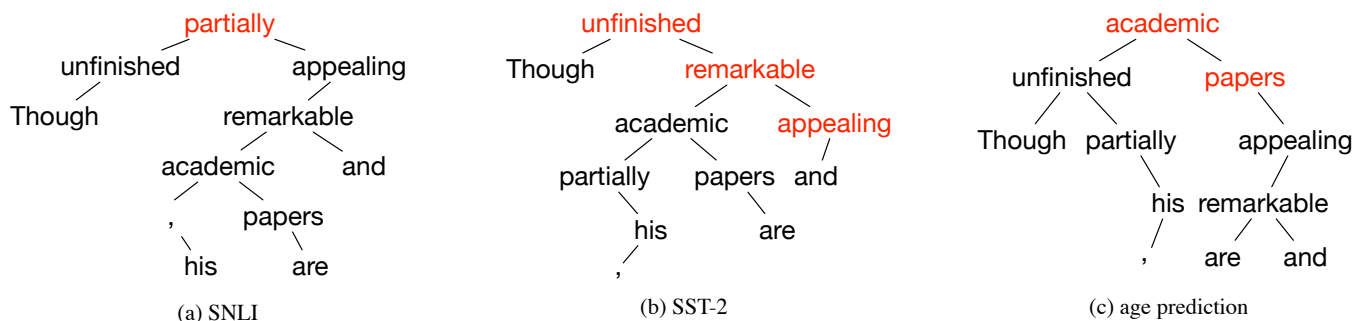


Figure 5: Different structures from different trained parsers for the same sentence *Though unfinished partially, his academic papers are remarkable and appealing*. We can see that words are emphasized adaptively based on the target task.

words (Figure 5b) as we have expected. In the parsed results of age prediction, *academic* and *papers* are emphasized (Figure 5c) because they are more likely to be discussed by college students, and are more informative to the age prediction task than other words.

Our model is able to pay attention to task-specific critical words for different tasks and learn interpretable structures, which is beneficial to the sentence understanding.

Conclusions and Future Work

We propose Attentive Recursive Tree (AR-Tree), a novel yet generic latent Tree-LSTM sentence embedding model, learning to learn task-specific structural embedding guided by word attention. Results on three different datasets demonstrate that AR-Tree learns reasonable attentive tree structures and outperforms previous Tree-LSTM models.

Moving forward, we are going to design a batch-mode tree construction algorithm, *e.g.*, asynchronous parallel recursive tree construction, to make the full exploitation of distributed and parallel computing power. Therefore, we may be able to learn an AR-Forest to embed paragraphs.

Acknowledgments

The work is supported by National Key Research and Development Program of China (2017YFB1002101), NSFC key project (U1736204, 61533018), and THUNUS NExT Co-Lab.

References

Arora, S.; Liang, Y.; and Ma, T. 2016. A simple but tough-to-beat baseline for sentence embeddings.

- Blunsom, P.; Grefenstette, E.; and Kalchbrenner, N. 2014. A convolutional neural network for modelling sentences. In *ACL*.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*.
- Bowman, S. R.; Gauthier, J.; Rastogi, A.; Gupta, R.; Manning, C. D.; and Potts, C. 2016. A fast unified model for parsing and sentence understanding. In *ACL*.
- Chen, Q.; Zhu, X.; Ling, Z.-H.; Wei, S.; Jiang, H.; and Inkpen, D. 2017. Recurrent neural network-based sentence encoder with gated attention for natural language inference. In *RepEval*.
- Choi, J.; Yoo, K. M.; and goo Lee, S. 2017. Learning to compose task-specific tree structures. In *AAAI*.
- Cocke, J. 1970. Programming languages and their compilers: Preliminary notes.
- Dai, A. M., and Le, Q. V. 2015. Semi-supervised sequence learning. In *NIPS*.
- dos Santos, C. N.; Tan, M.; Xiang, B.; and Zhou, B. 2016. Attentive pooling networks. *CoRR*, abs/1602.03609.
- Goldberg, Y., and Elhadad, M. 2010. An efficient algorithm for easy-first non-directional dependency parsing. In *NAACL*.
- Hill, F.; Cho, K.; and Korhonen, A. 2016. Learning distributed representations of sentences from unlabelled data. In *NAACL*.
- Hu, B.; Lu, Z.; Li, H.; and Chen, Q. 2014. Convolutional neural network architectures for matching natural language sentences. In *NIPS*.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Kasami, T. 1965. An efficient recognition and syntaxanalysis algorithm for context-free languages. Technical report.
- Kim, Y.; Denton, C.; Hoang, L.; and Rush, A. M. 2017. Structured attention networks. *arXiv preprint arXiv:1702.00887*.
- Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kumar, A.; Irsoy, O.; Ondruska, P.; Iyyer, M.; Bradbury, J.; Gulrajani, I.; Zhong, V.; Paulus, R.; and Socher, R. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *ICML*.
- Lin, Z.; Feng, M.; Santos, C. N. d.; Yu, M.; Xiang, B.; Zhou, B.; and Bengio, Y. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.
- Maillard, J.; Clark, S.; and Yogatama, D. 2017. Jointly learning sentence embeddings and syntax with unsupervised tree-lstms. *arXiv preprint arXiv:1705.09189*.
- McCann, B.; Bradbury, J.; Xiong, C.; and Socher, R. 2017. Learned in translation: Contextualized word vectors. In *NIPS*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Mou, L.; Men, R.; Li, G.; Xu, Y.; Zhang, L.; Yan, R.; and Jin, Z. 2016. Natural language inference by tree-based convolution and heuristic matching. In *ACL*.
- Munkhdalai, T., and Yu, H. 2017a. Neural semantic encoders. In *ACL*.
- Munkhdalai, T., and Yu, H. 2017b. Neural tree indexers for text understanding. In *ACL*.
- Nivre, J. 2003. An efficient algorithm for projective dependency parsing. In *Proceedings of the 8th International Workshop on Parsing Technologies (IWPT)*.
- Parikh, A.; Täckström, O.; Das, D.; and Uszkoreit, J. 2016. A decomposable attention model for natural language inference. In *EMNLP*.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Radford, A.; Jozefowicz, R.; and Sutskever, I. 2017. Learning to generate reviews and discovering sentiment. *arXiv preprint arXiv:1704.01444*.
- Shi, H.; Zhou, H.; Chen, J.; and Li, L. 2018. On tree-based neural sentence modeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4631–4641.
- Socher, R.; Pennington, J.; Huang, E. H.; Ng, A. Y.; and Manning, C. D. 2011. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *EMNLP*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*.
- Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*.
- Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*.
- Tai, K. S.; Socher, R.; and Manning, C. D. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *ACL*.
- Wang, D., and Nyberg, E. 2015. A long short-term memory model for answer sentence selection in question answering. In *ACL*.
- Williams, A.; Drozdov, A.; and Bowman, S. R. 2017. Learning to parse from a semantic objective: It works. is it syntax? *arXiv preprint arXiv:1709.01121*.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*.
- Yogatama, D.; Blunsom, P.; Dyer, C.; Grefenstette, E.; and Ling, W. 2016. Learning to compose words into sentences with reinforcement learning. *arXiv preprint arXiv:1611.09100*.
- Younger, D. H. 1967. Recognition and parsing of context-free languages in time n^3 . *Information and control*.
- Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*.
- Zeiler, M. D. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Zhu, X.; Sobihani, P.; and Guo, H. 2015. Long short-term memory over recursive structures. In *ICML*.