

Learning from Non-Binary Constituency Trees via Tensor Decomposition

Daniele Castellana

University of Pisa, Pisa

daniele.castellana@di.unipi.it

Davide Bacciu

University of Pisa, Pisa

bacciu@di.unipi.it

Abstract

Processing sentence constituency trees in binarised form is a common and popular approach in literature. However, constituency trees are non-binary by nature. The binarisation procedure changes deeply the structure, furthering constituents that instead are close. In this work, we introduce a new approach to deal with non-binary constituency trees which leverages tensor-based models. In particular, we show how a powerful composition function based on the canonical tensor decomposition can exploit such a rich structure. A key point of our approach is the weight sharing constraint imposed on the factor matrices, which allows limiting the number of model parameters. Finally, we introduce a Tree-LSTM model which takes advantage of this composition function and we experimentally assess its performance on different NLP tasks.

1 Introduction

One of the fundamental problems in Natural Language Processing (NLP) is learning a distributed encoding of sentences, as this is the stepping stone for many NLP tasks, such as sentence classification, sentiment analysis and natural language inference. The multitude of approaches addressing this problem can be categorised according to how a sentence is represented.

The simpler sentence representation is bag-of-words, which depicts sentences as words multisets ignoring the word order. Despite the simple representation, it has been used to obtain meaningful sentence encodings (Iyyer et al., 2015; Wieting et al., 2016; Arora et al., 2017).

Sequence representation overcomes this limitation considering the sentence as an ordered sequence of words. It allows building models which progressively constructs a sentence encoding, processing one word at the time. Recurrent Neural Network (Elman, 1990) and Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) are probably the most famous models which use this representation.

A key aspect of sentences, which is missing in sequential processing, is compositionality. For example, the sentence *"The sky is blue and the grass is green"* is obtained by composing the two sub-phrases *"The sky is blue"* and *"the grass is green"* with the conjunction *"and"*. The intrinsic compositionality of sentences makes them suitable for a tree representation, where the whole sentence (the root) is built in terms of sub-phrases (the internal nodes) which in turn are defined in terms of smaller constituents; the base cases are words (the leaves) since they are the atomic piece of information. This representation takes the name of *constituency tree*. In Fig. 1a we show the constituency tree of the sentence *"Effective but too-tepid biopic"*: the leaves are the words while internal nodes represent syntactic categories which are the constituents of the whole sentence.

There are many models which compute a sentence encoding starting from its constituency tree. For our purposes, we restrict the discussion on bottom-up Recursive Neural Networks (RecNNs) (Goller and Kuchler, 1996; Frasconi et al., 1998). The parsing direction is constrained by the structure of constituency trees, having information (i.e. words) on leaf nodes. In this domain, we refer to the term *composition function* to indicate the state-transition function which computes the *representation* (i.e. the

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

represented as tensors. The exponential number of parameters with respect to the tree out-degree (i.e. the maximum number of children for each node in the tree) required by the full-tensorial approach can be controlled by applying tensor decomposition. The tensorial models outperform sum-based models, especially when the tree out-degree increases (Castellana and Bacciu, 2020a; Castellana and Bacciu, 2020b).

Within the scope of this paper, we unveil that non-binary constituency trees can be effectively exploited to improve predictive performance in NLP task, showing that more powerful composition functions are necessary to take advantages of such a rich representation. To this end, we introduce two new Tree-LSTM models which leverage canonical tensor decomposition: the former is suitable for binarised constituency trees, while the latter can process general non-binary constituency trees imposing weight sharing on the tensor decomposition factors. Finally, we test the quality of sentence encodings produced by our models on different NLP tasks, showing that the combination of a rich representation and a powerful composition function is able to outperform baseline models using the same number of parameters.

2 Related Models

In this section, we discuss three architectures from the literature that are related to the approach put forward in this paper and are used as baselines in our empirical analysis.

2.1 Tree-Structured LSTM

The Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) is one of the most popular neural architecture to process sequential data. Tai et al. (2015) proposed two extensions of such architecture to handle tree-structured information: the L -ary Tree-LSTM and the Child-Sum Tree-LSTM. Both models propagate the information along the input tree structure in a bottom-up fashion (i.e. from the leaves to the root). Hence, at each node, the tree-LSTM cell aggregates the hidden child states to compute its gates leveraging sum-based composition functions.

L -ary Tree-LSTM (Tai et al., 2015) assumes that child nodes are ordered and then they can be indexed from 1 to L , where L denotes the maximum node output degree, i.e. the maximum number of children a node can have. Let v a generic node, we indicate with $h_{vj} \in \mathbb{R}^d$ and $c_{vj} \in \mathbb{R}^d$ respectively the hidden state and the memory cell of its j -th child node; hence, the L -ary Tree-LSTM cell computation is defined by the following equations:

$$\begin{aligned} t_v &= \sigma \left(W^t x_v + \sum_{j=1}^L U_j^t h_{vj} + b^t \right) \quad t \in \{i, o, u\}, & c_v &= i_v \odot u_v + \sum_{j=1}^L f_{vj} \odot c_{vj}, \\ f_{vk} &= \sigma \left(W^f x_v + \sum_{j=1}^L U_{kj}^f h_{vj} + b_k^f \right) \quad k \in [1, L], & h_v &= o_v \odot \tanh(c_v), \end{aligned} \quad (1)$$

where $c_v \in \mathbb{R}^d$, $h_v \in \mathbb{R}^d$ and $x_v \in \mathbb{R}^n$ are the memory cell state, the hidden state and the label of node v ; $\{i_v, o_v, u_v\} \in \mathbb{R}^d$ are the input gate, the output gate and the update value, respectively, and f_{vk} is the forget gate associated with k -th child of v . The symbol σ denotes the logistic sigmoid function and \odot denotes the elementwise product.

We apply this model on binary constituency tree. Hence, the input tree has labels (i.e. the words) attached only on leaf nodes. In particular, we consider x_v as the vector representation of a word in the sentence. Hence, internal nodes do not have any input labels (see Fig. 2a). In the rest of the paper, we refer to this model as Binary Sum-LSTM.

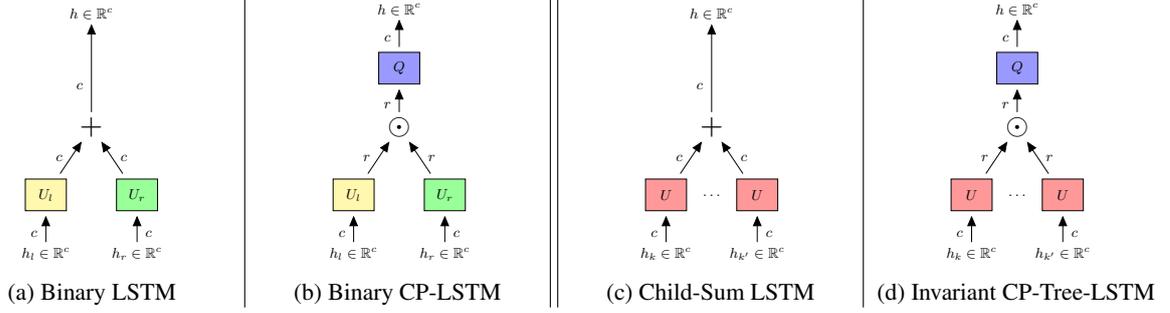


Figure 2: Graphical illustration of composition functions; for the sake of clarity, we removed the bias.

Child-Sum Tree-LSTM (Tai et al., 2015) assumes that there is no order among child nodes. Let v a generic node, we indicate with $\text{Ch}(v)$ the set of its child nodes; hence, its hidden state h_v is computed as:

$$\begin{aligned}
 t_v &= \sigma \left(W^t x_v + \sum_{k \in \text{Ch}(v)} U^t h_k + b^t \right) \quad t \in \{i, o, u\}, & c_v &= i_v \odot u_v + \sum_{k \in \text{Ch}(v)} f_{vk} \odot c_{vk}, \\
 f_{vk} &= \sigma \left(W^f x_v + U^f h_{vk} + b^f \right) \quad k \in \text{Ch}(v), & h_v &= o_v \odot \tanh(c_v).
 \end{aligned} \tag{2}$$

When comparing Eq. (2) and Eq. (1), we observe that the Child-Sum Tree-LSTM can be derived from the L -ary Tree-LSTM by imposing weight sharing across children positions. In fact, in Eq. (2), the parameter U is the same for each child, while in Eq. (1) the subscript j of the parameter U_j denotes that the weights are positional dependent.

In Tai et al. (2015), this model has been applied only on dependency trees. However, in this work, we use it on non-binary constituency trees. Again, we consider x_v as the vector representation of a word in the sentence and therefore are attached only on leaf nodes (see Fig. 2c). In the remainder of the paper, we refer this model with the name Child-Sum LSTM.

2.2 TreeNet

TreeNet (Cheng et al., 2018) has been introduced with the aim of learning from unconstrained tree-structured data (i.e. trees where out-degree not fixed). The idea is to process child nodes as a sequence and links only the rightmost child node to the parent node. Hence, each node compose the information of its left sibling and its rightmost child:

$$\begin{aligned}
 o_v &= \sigma \left(U_s^o h_{vs} + U_c^o h_{vc} + b^o \right), & c_v &= f_{vs} \odot c_{vs} + f_{vc} \odot c_{vc}, \\
 f_{vk} &= \sigma \left(U_{ks}^f h_{vs} + U_{kc}^f h_{vc} + b_k^f \right), \quad k \in \{s, c\}, & h_v &= o_v \odot \tanh(c_v),
 \end{aligned} \tag{3}$$

where $\{h_{vs}, c_{vs}\} \in \mathbb{R}^d$ are hidden state and memory cell of left sibling of v , while $\{h_{vc}, c_{vc}\} \in \mathbb{R}^d$ are hidden state and memory cell of the rightmost child node of v .

If the node v is a leaf, its hidden state depends solely on the input label x_v :

$$\begin{aligned}
 i_v &= \sigma \left(W^i x_v + \hat{b}^i \right), & o_v &= \sigma \left(W^o x_v + \hat{b}^o \right), & u_v &= \sigma \left(W^u x_v + \hat{b}^u \right), \\
 c_v &= i_v \odot u_v, & h_v &= o_v \odot \tanh(c_v).
 \end{aligned} \tag{4}$$

Comparing Eq. (3) and Eq. (1), we can observe that TreeNet cell is a binary Tree-LSTM cell without the input and the output gates. In fact, in both cells, all the gates are computed composing two constituents. The TreeNet define the constituents of a node as its left sibling and its rightmost child, while the binary Tree-LSTM used directly its left child node and right child node. Hence, we argue that the difference between the two lies solely on how the tree is binarised, rather than on how the tree is processed. In Fig. 1d, we show an example on how a constituency tree is binarised according to the TreeNet; the constituent *ADJP* of the original constituency tree (see Fig. 1a) is composed of three words: "Effective",

”but”, ”too-tepid”. The TreeNet breaks this ternary relation processing one words at a time; as we can see in Fig. 1d, the node which has the first word ”Effective” is combined with a bottom node since it does not have a left sibling. The result is then fused with the word ”but”. Finally, also the word ”too-tepid” is combined with the result of the previous composition, obtaining the encoding of the constituent *ADJP*. Hence, the original ternary relation is broken into a sequence of three binary relations (one for each word) each of them combines the composition of previous words with the new word.

3 Canonical Tree-LSTM

This section aims to introduce new Tree-LSTM models that extend the recurrent approach to tensor-based processing tailored for constituency trees. These two models rely on Tensor Tree-LSTM (Castellana and Bacciu, 2020a) and canonical tensor decomposition. We focus on such decomposition since it can be combined with a weight sharing constraint, developing a composition function which can exploit non-binary constituency trees without increasing the model parameters number.

Since these models are applied on constituency trees, we define them only on the nodes which apply composition functions (i.e. the internal nodes). The input label x_v , which is attached only on leaf nodes, is processed using the same computation described in Eq. 4.

3.1 Canonical Decomposition

The canonical decomposition (usually denoted by CP) factorises a tensor into a sum of component rank-one tensors, i.e. tensors that are obtained by the outer product of vectors. For example, let $\mathbf{T} \in \mathbb{R}^{d_1 \times d_2 \times \dots \times d_L}$ a general L -th order tensor, the CP decomposition is defined by (Kolda and Bader, 2009):

$$\mathbf{T}(j_1, \dots, j_L) = \sum_{i=1}^r U_1(j_1, i)U_2(j_2, i) \dots U_L(j_L, i), \quad (5)$$

where brackets are used to denote entries of vectors, matrices and tensors (e.g. $W(i, j)$ denotes the entry in the i -th row and j -th column of W). $U_1 \in \mathbb{R}^{d_1 \times r}, \dots, U_L \in \mathbb{R}^{d_L \times r}$ are the factor matrices of the decomposition. Each factor matrix contains r vectors, which are the basis of the rank-one tensors. The value of r indicates the number of such tensors which are summed to obtain the original tensor \mathbf{T} and it is denoted as tensor rank.

Following Castellana and Bacciu (2020a), we apply the decomposition on a tensor which represents a multi-affine map. Note that in this case the tensor that should be decomposed is not known, since it is the parameter of the recursive model and it is learned from data. Instead, we assume that such tensor is already decomposed, making the decomposition factors the new recursive model parameters. Let $\phi^{\mathbf{T}} : \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_L} \rightarrow \mathbb{R}^K$ a multi-affine map which parameter is the tensor $\mathbf{T} \in \mathbb{R}^{(d_1+1) \times \dots \times (d_L+1) \times K}$, applying the CP decomposition on \mathbf{T} we obtain:

$$\begin{aligned} a(k) &= \phi^{\mathbf{T}}(a_1, \dots, a_L)(k) = \sum_{j_1=1}^{d_1+1} \dots \sum_{j_L=1}^{d_L+1} \mathbf{T}(j_1, \dots, j_L, k) \bar{a}_1(j_1) \dots \bar{a}_L(j_L) = \\ &= \sum_{j_1=1}^{d_1+1} \dots \sum_{j_L=1}^{d_L+1} \sum_{i=1}^r U_1(j_1, i) \dots U_L(j_L, i) Q(k, i) \bar{a}_1(j_1) \dots \bar{a}_L(j_L) = \\ &= \sum_{i=1}^r Q(k, i) \left(\sum_{j_1=1}^{d_1+1} U_1(j_1, i) \bar{a}_1(j_1) \right) \dots \left(\sum_{j_L=1}^{d_L+1} U_L(j_L, i) \bar{a}_L(j_L) \right), \quad (6) \end{aligned}$$

where $\bar{a} = [a; 1]$ denotes the homogeneous coordinate of vector a . From the equation, we can observe that the CP decomposition define a multi-affine map which applies each factor matrix to the corresponding input vector (i.e. the one on the same mode), obtaining a vector $e_j \in \mathbb{R}^r$, for each mode $j \in \{1, \dots, L\}$. Then, these vectors are element-wise multiplied and the result mapped to the output space \mathbb{R}^K thanks to matrix Q , i.e. the factor matrices on the output (the last) mode.

3.2 Binary Canonical Tree-LSTM

The Binary Tensor Tree-LSTM (Castellana and Bacciu, 2020a) computes the hidden state of an internal node v by:

$$\begin{aligned} i_v &= \sigma(\phi^{\mathbf{I}}(h_{vl}, h_{vr})), & f_{vk} &= \sigma(\phi^{\mathbf{F}^k}(h_{vl}, h_{vr})) \text{ with } k \in \{l, r\}, \\ o_v &= \sigma(\phi^{\mathbf{O}}(h_{vl}, h_{vr})), & c_v &= i_v \odot u_v + f_{vl} \odot c_{vl} + f_{vr} \odot c_{vr}, \\ u_v &= \sigma(\phi^{\mathbf{U}}(h_{vl}, h_{vr})), & h_v &= o_v \odot \tanh(c_v). \end{aligned} \quad (7)$$

All the terms (except hidden state and memory cell) are computed by applying a multi-affine map $\phi^{\mathbf{T}}(\cdot)$ on the left and right child hidden states, i.e. h_{vl} and h_{vr} , respectively. The superscript \mathbf{T} indicates the parameters of such multi-affine map; hence, the parameters $\{\mathbf{I}, \mathbf{O}, \mathbf{U}, \mathbf{F}_l, \mathbf{F}_r\}$ define the LSTM cell computation.

The Binary Canonical Tree-LSTM (Binary CP-LSTM) exploits the canonical decomposition on tensor parameters (see Fig. 2b):

$$e_l = U_l^t h_{vl} + b_l^t, \quad e_r = U_r^t h_{vr} + b_r^t, \quad e = a_l \odot a_r, \quad \phi^{\mathbf{T}}(h_{vl}, h_{vr}) = Q^t e + b^t, \quad (8)$$

where $\{U_l^t, U_r^t\} \in \mathbb{R}^{d \times r}$, $Q^t \in \mathbb{R}^{r \times d}$, $\{b_l^t, b_r^t\} \in \mathbb{R}^r$, $b \in \mathbb{R}^d$ define the decomposition's factor matrices (the bias is made explicit) and the value r is the decomposition rank. We use the superscript $t \in \{i, o, u, f_l, f_r\}$ to indicate different parameters of each multi-affine maps. The number of parameters required by the Binary CP-LSTM is $O(dr)$.

3.3 Invariant Canonical Tree-LSTM

The Invariant Canonical Tree-LSTM (Invariant CP-LSTM) allows to encode information from constituency trees which are not binarised. Let v an internal node, its hidden state it is computed by:

$$\begin{aligned} i_v &= \sigma(\phi^{\mathbf{I}}(\text{Ch}(v))), & f_{vk} &= \sigma(U^f h_{vk} + b^f) \text{ with } k \in \text{Ch}(v), \\ o_v &= \sigma(\phi^{\mathbf{O}}(\text{Ch}(v))), & c_v &= i_v \odot u_v + \sum_{k \in \text{Ch}(v)} f_{vk} \odot c_{vk}, \\ u_v &= \sigma(\phi^{\mathbf{U}}(\text{Ch}(v))), & h_v &= o_v \odot \tanh(c_v). \end{aligned} \quad (9)$$

As in Eq. (2), we indicate with $\text{Ch}(v)$ the set of child nodes of v . $\{\mathbf{I}, \mathbf{O}, \mathbf{U}\}$ are the parameters of the multi-affine maps to compute the input gate, output gate and the update value respectively.

Again, we exploit the canonical decomposition on parameter tensors to define the Invariant CP-LSTM. Moreover, we impose the weight sharing among input factor matrices, obtaining (see Fig. 2d):

$$e_k = U^t h_k + b^t \text{ with } k \in \text{Ch}(v), \quad e = \bigodot_{k \in \text{Ch}(v)} e_k, \quad \phi^{\mathbf{T}}(\text{Ch}(v)) = Q^t e + q^t. \quad (10)$$

where $U^t \in \mathbb{R}^{d \times r}$, $b^t \in \mathbb{R}^r$ is the factor matrix shared on all input nodes and $Q^t \in \mathbb{R}^{r \times d}$, $q^t \in \mathbb{R}^d$ is the factor matrix on the last mode; the value r is the decomposition rank. Hence, the number of parameters required by Invariant CP-LSTM is still $O(dr)$. Thanks to the weight sharing, we are able to process non-binary trees without adding new parameters. As we highlighted in Sec. 2.1, the weight sharing makes the model *invariant* to child nodes order, thus its name is Invariant CP-LSTM.

4 Experiments

We test the models introduced in the previous section on two tasks: sentence classification and semantic textual similarity.

Sentence Classification. The goal of this task is to predict the class of the given input sentence. Hence, we use a Tree-LSTM to encode the constituency tree of the input sentence in a succinct representation (h_{root}) and then we feed it to a single-layer Neural Network to predict the class $y \in \{1 \dots m\}$:

$$s = \text{ReLU}(W' h_{\text{root}} + b') \quad p(y) = \text{softmax}(W'' s + b''), \quad (11)$$

where $s \in \mathbb{R}^s$ is the hidden representation of the classifier and $W' \in \mathbb{R}^{d \times s}$, $b \in \mathbb{R}^s$, $W'' \in \mathbb{R}^{s \times m}$, $b'' \in \mathbb{R}^m$ are the classifiers parameters. Following Tai et al. (2015), we use dropout (Srivastava et al., 2014) with rate 0.5 on both h_{root} and h_s . We test our models on three different classification datasets:

- **SST-5:** Stanford Sentiment Treebank (SST) dataset (Socher et al., 2013) contains sentences that are classified with a fine-grained sentiment which goes from 1 to 5 (very negative, negative, neutral, positive and very positive);
- **SST-2:** identical to SST-5, but with binary sentiment class; neutral sentences are removed and all negative (positive) sentences are collapsed in one cluster;
- **TREC:** TREC dataset (Li and Roth, 2002) contains questions that are annotated with six classes which indicates a question type.

Semantic Textual Similarity. The goal of this task is to predict the semantic similarity between two sentences. Let a and b the two sentences, we produce their encodings h_a and h_b applying a Tree-LSTM on both sentence constituency trees and taking the hidden state of the root. Then, we compute the similarity score $y_r \in [1, m]$ as in Tai et al. (2015):

$$s = \sigma(W^+ |h_a - h_b| + W^\times (h_a \odot h_b) + b), \quad p_r = \text{softmax}(W^r s + b^r), \quad y_r = r^\top p_r \quad (12)$$

where $s \in \mathbb{R}^s$ is the hidden representation of the classifier, $\{W^+, W^\times\} \in \mathbb{R}^{d \times s}$, $b \in \mathbb{R}^s$, $W^r \in \mathbb{R}^{s \times m}$, $b^r \in \mathbb{R}^m$ are the classifiers parameters and $r^\top = [1, 2, \dots, m]$.

It also common to represent the semantic similarity attaching an entailment class to each pair of sentences. In this case, we predict the entailment class $y_e = \arg \max(p_e)$ starting from the distribution $p_e = \text{softmax}(W^e s + b^e)$; s is computed as in Eq. (12). We test our models on two different datasets:

- **SICK-R:** Sentences Involving Compositional Knowledge (SICK) dataset (Marelli et al., 2014), contains sentence pairs annotated with a relatedness score between 1 and 5;
- **SICK-E:** identical to SICK-R, but with entailment classes instead of relatedness scores; the entailment class indicates whether one sentence entails or contradicts the other (neutral, entailment, contradiction).

4.1 Implementation and training details

We have implemented all the models using PyTorch (Paszke et al., 2019) and Deep Graph Library (Wang et al., 2019). The source code to reproduce the model and the experiments is released here.¹ Constituency trees are built using the PCFG constituency parser of the Stanford Core NLP (Manning et al., 2014). Also, we binarise them computing the Chomsky Normal Form available in the Natural Language Tool Kit (Bird et al., 2009). To facilitate the learning, we collapse all unary relations. In each task, we perform a grid search to find the best hyper-parameters configuration (see Appendix A for further details). For the training, we use AdaDelta (Zeiler, 2012) algorithm and Adam (Kingma and Ba, 2015) optimiser. In all experiments, we initialised our word representations using 300-dimensional Glove vectors (Pennington et al., 2014). We fine-tune the words representation only on the SST dataset.

5 Results

Table 1 reports the results obtained in our experiments. Also, we report the results obtained by our implementation of the TreeNet (Chen et al., 2018) which differ from the results published in the original

¹<https://github.com/danielecastellana22/tensor-tree-nn>

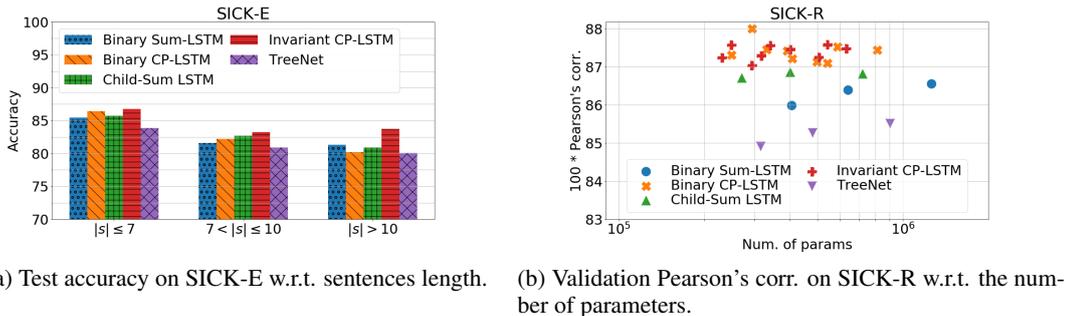


Figure 3: Analyses of the results obtained on SICK-E and SICK-R dataset.

	SST-5	SST-2	SICK-E	SICK-R	TREC
Binary Sum-LSTM*	51.5 (0.7)	87.9 (0.2)	82.3 (0.5)	84.3 (0.7)	92.3 (0.8)
Binary CP-LSTM (our)	50.3 (0.8)	88.0 (0.3)	82.4 (0.3)	85.9 (0.1)	90.4 (0.2)
Binary LSTM (Tai et al., 2015)	51.00 (0.5)	88.0 (0.3)	-	85.8 (0.4)	-
Child-Sum LSTM*	49.4 (0.6)	85.5 (0.8)	82.6 (0.4)	84.9 (0.2)	91.9 (1.0)
Invariant CP-LSTM (our)	48.3 (0.8)	85.3 (0.3)	84.2 (0.4)	86.4 (0.1)	90.0 (0.7)
TreeNet*	48.4 (1.5)	87.0 (0.5)	81.2 (0.3)	84.5 (0.3)	91.3 (1.1)

Table 1: Results obtained on different task All the values are accuracy except SICK-R, whose score is Pearson’s correlation multiplied by 100. The superscript * indicates we re-implement the model.

paper on the TREC dataset. However, to be sure that there are no errors in our implementation, we run the original code published by Cheng et al. (2018) with our experimental settings (i.e. model selection on the validation set and risk assessment on the test set) and we obtain results comparable with the one published in this table. In the next paragraphs, we analyse in details the results obtained on each dataset.

SST. The results obtained on the SST dataset (both with fine-grained and binary labels) do not show any improvements using non-binary constituency trees. However, the comparison is unfair since the original dataset provides labels on internal nodes of binary constituency trees. By removing the binarisation, it is no longer possible to leverage such information during training. As a reference, note that binary constituency trees data contains 119.413 labels, while the non-binary constituency trees data contains only 91.536 labels.

SICK. The results obtained on the SICK dataset show the advantage of combining a rich representation (such as non-binary trees) and a more powerful composition function. In fact, the Invariant CP-LSTM outperforms all the other models in both the entailment (SICK-E) and relatedness (SICK-R) task. It is worth to point out that the Invariant CP-LSTM is the only model which benefits from the non-binary representation. In Fig. 3a we report the test accuracy of each model with respect to the input length (we consider the maximum between the length of each sentence in the input pairs). Observing the plot, it is clear that most of the models struggle with long sentences. The Invariant CP-LSTM model, instead, reaches an accuracy of approximately 84%, while other models stop around 81%. In Fig. 3b, we show the validation results on SICK-R obtained by all the models against the number of parameters they require. Observing the plot, it is clear that thanks to the canonical decomposition and the weight sharing, we can build powerful composition function using the same number of parameters of sum-based functions.

TREC. On the TREC dataset, the sum-based composition functions seems to be advantageous over the canonical counterpart. In fact, both on binary and non-binary trees, the sum-based Tree-LSTMs outperforms the canonical-based Tree-LSTMs. Also, the Tree-Net (which is based on summation), reaches results comparable with the Child-Sum LSTM. We argue that the summation is preferable for the question classification task, probably due to the intrinsic characteristics of question sentences.

5.1 Qualitative analysis on SICK-E

In this section, we analyse in details the prediction of Child-Sum LSTM and Invariant CP-LSTM on the SICK-E dataset. In particular, we report the prediction on example #3991 of the test set. The input pair

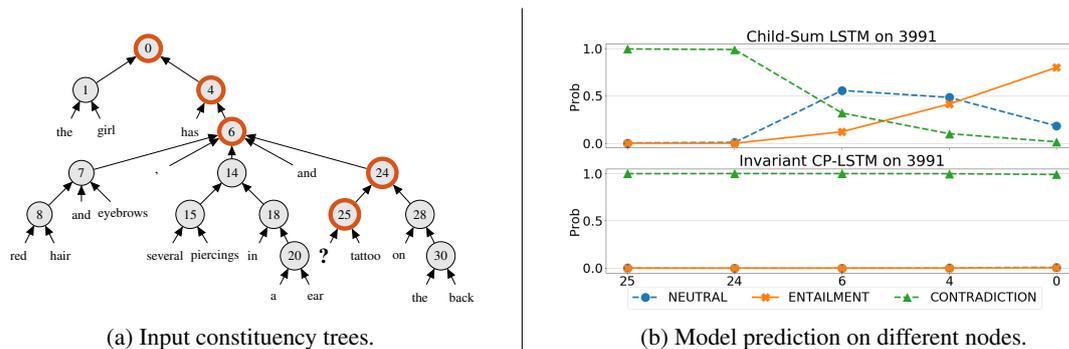


Figure 4: Comparison between Invariant CP-LSTM and Child-Sum LSTM prediction on input #3991 taken from SICK (Marelli et al., 2014) test set.

is composed of the following sentences:

- A: *The girl has red hair and eyebrows, several piercings in a ear and **a tattoo** on the back.*
 B: *The girl has red hair and eyebrows, several piercings in a ear and **no tattoo** on the back.*

The two sentences are exactly the same, unless the sub-phrase *a tattoo* in sentence A which becomes *no tattoo* in sentence B. Hence, the expected output is "contradiction". In Fig. 4a we plot the constituency tree of the input sentence, indicating with ? the position where the two sentences differ. Also, we highlight all the nodes that are in the path between the ? and the root. These are the only nodes which have a different sub-tree in the two sentences. To analyse how the two models predict the final label, we study how the prediction changes going up through the structure. In Fig. 4b we report the output of the classifier fed with the hidden states pair (h_{ai}, h_{bi}) , where h_{ai} (h_{bi}) is the hidden state of node i computed by the tree model on sentence A (B). On node 25, both models predict correctly the contradiction. However, going up through the structure, the Child-Sum Tree-LSTM changes the prediction to entailment, which will be the final output on node 0. The change of the output label start on node 6, which is the node where most of the information is aggregated; it seems that the Child-Sum Tree-LSTM performs a sort of average which soften the contribution of node 24 (the only one that instead should be taken into account). On the contrary, the Invariant CP-LSTM propagates correctly the information through the structure. Even if the model process sub-phrases which are identical, their contribution does not influence the output. In fact, observing Fig. 4b, we can notice that the class predicted by the Invariant CP Tree-LSTM is always "contradiction" in all nodes in the path between ? and the root.

6 Conclusion

In this paper, we show that using non-binary constituency trees can be beneficial, especially in semantic similarity tasks. Moreover, we highlight the need of powerful composition function to exploit such a rich representation. To this end, we have introduced a new Tree-LSTM model which leverages tensor canonical decomposition and weight sharing to process non-binary trees without adding new parameters.

Such results pave the way to the definition of new tensor models which leverage suitable tensor decomposition to take advantage of non-binary constituency trees. To this end, the next step would be the application of other tensor decompositions. Among the others, the tensor train decomposition seems to be promising to define new composition functions which are sensitive to child nodes order.

Ultimately, we would like to test multiple tensor-based models on different NLP tasks, studying the relation between the bias introduced by each different tensor decomposition and the intrinsic property of the task.

Acknowledgements

This work has been partially supported by the MIUR SIR 2014 LIST-IT project (grant n. RBSI14STDE) and by TAILOR, a project funded by EU Horizon 2020 programme under GA No 952215.

References

- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough Baseline for Sentence Embeddings. *International Conference on Learning Representations (ICLR)*, 15:416–424.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Daniele Castellana and Davide Bacciu. 2020a. Generalising recursive neural models by tensor decomposition. In *Proceedings of the 2020 IEEE World Congress on Computational Intelligence*.
- Daniele Castellana and Davide Bacciu. 2020b. Tensor decompositions in recursive neural networks for tree-structured data. In *Proceedings of the The 28th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*.
- Xinyun Chen, Chang Liu, and Dawn Song. 2018. Tree-to-tree Neural Networks for Program Translation. *Workshop track in International Conference on Learning Representations (ICLR)*.
- Zhou Cheng, Chun Yuan, Jiancheng Li, and Haiqin Yang. 2018. TreeNet: Learning sentence representations with unconstrained tree structure. *IJCAI International Joint Conference on Artificial Intelligence*, 2018-July:4005–4011.
- Jeffrey L. Elman. 1990. Finding structure in time. *Cognitive Science*, 14(2):179–211.
- Paolo Frasconi, Marco Gori, and Alessandro Sperduti. 1998. A general framework for adaptive processing of data structures. *IEEE Transactions on Neural Networks*, 9(5):768–786.
- C. Goller and A. Kuchler. 1996. Learning task-dependent distributed representations by backpropagation through structure. *Proceedings of International Conference on Neural Networks (ICNN’96)*, 1:347–352.
- Sepp Hochreiter and Jurgen Schmidhuber. 1997. Long short term memory. *Neural Computation*, 9(8):1735–1780.
- Minlie Huang, Qiao Qian, and Xiaoyan Zhu. 2017. Encoding syntactic knowledge in neural networks for sentiment classification. *ACM Transactions on Information Systems*, 35(3).
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep Unordered Composition Rivals Syntactic Methods for Text Classification. 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)*, pages 1681–1691, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Taeuk Kim, Jihun Choi, Daniel Edmiston, Sanghwan Bae, and Sang-goo Lee. 2019. Dynamic Compositionality in Recursive Neural Networks with Structure-Aware Tag Representations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:6594–6601.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR, 12.
- Tamara G. Kolda and Brett W. Bader. 2009. Tensor Decompositions and Applications. *SIAM Review*, 51(3):455–500.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *Proceedings of the 19th international conference on Computational linguistics -*, volume 1, pages 1–7, Morristown, NJ, USA. Association for Computational Linguistics.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2017. Dynamic Compositional Neural Networks over Tree Structure. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pages 4054–4060, California, 8. International Joint Conferences on Artificial Intelligence Organization.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Stroudsburg, PA, USA. Association for Computational Linguistics.

- Marco Marelli, Luisa Bentivogli, Marco Baroni, Raffaella Bernardi, Stefano Menini, and Roberto Zamparelli. 2014. SemEval-2014 Task 1: Evaluation of Compositional Distributional Semantic Models on Full Sentences through Semantic Relatedness and Textual Entailment. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, number 1, pages 1–8, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary Devito Facebook, A I Research, Zeming Lin, Alban Desmaison, Luca Antiga, Orobix Srl, and Adam Lerer. 2019. Automatic differentiation in PyTorch. In *Advances in Neural Information Processing Systems 32*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 1532–1543.
- Gehui Shen, Zhi-Hong Deng, Ting Huang, and Xi Chen. 2020. Learning to compose over tree structures via POS tags for sentence representation. *Expert Systems with Applications*, 141:112917, 3.
- Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *EMNLP-CoNLL 2012 - 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, Proceedings of the Conference*.
- Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. *EMNLP 2013 - 2013 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, 8(9):1631–1642, 12.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. *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)*, pages 1556–1566.
- Zhiyang Teng and Yue Zhang. 2017. Head-Lexicalized Bidirectional Tree LSTMs. *Transactions of the Association for Computational Linguistics*, 5:163–177.
- Minjie Wang, Lingfan Yu, Da Zheng, Quan Gan, Yu Gai, Zihao Ye, Mufei Li, Jinjing Zhou, Qi Huang, Chao Ma, Ziyue Huang, Qipeng Guo, Hao Zhang, Haibin Lin, Junbo Zhao, Jinyang Li, Alexander Smola, and Zheng Zhang. 2019. Deep Graph Library: Towards Efficient and Scalable Deep Learning on Graphs. 9.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016. Towards universal paraphrastic sentence embeddings. In *4th International Conference on Learning Representations, ICLR 2016*.
- Matthew D. Zeiler. 2012. ADADELTA: An Adaptive Learning Rate Method.

A Appendix A. Experiments Details

For each model, we select the best hyper-parameters through a grid search (see Table 2). The hyper-parameters validated are: the batch size bs , the hidden representation of the Tree-LSTM d and the hidden representation of the classifier s . The models which use the canonical decomposition have also the rank r as hyper-parameter. Moreover, we train all the models using the AdaDelta algorithm, except for the TreeNet on SICK-E, SICK-R and TREC dataset where we use Adam optimiser. When we use Adam, we validate the learning rate lr and we fix the batch size to 25.

	Model	bs or lr	d	r	s
SST	Binary Sum-LSTM	[5, 10, 25]	[100, 200, 300]	-	[0, 500, 1000]
	Binary CP-LSTM	[5, 10, 25]	[100, 200, 300]	[50, 100, 150]	[0, 500, 1000]
	Child-Sum LSTM	[5, 10, 25]	[100, 200, 300]	-	[0, 500, 1000]
	Invariant CP-LSTM	[5, 10, 25]	[100, 200, 300]	[50, 100, 150]	[0, 500, 1000]
	TreeNet	[5, 10, 25]	[100, 200, 300]	-	[0, 500, 1000]
SICK	Binary Sum-LSTM	[10, 25, 40]	[150, 200, 300]	-	[50, 100, 200]
	Binary CP-LSTM	[10, 25, 40]	[150, 200, 300]	[30, 50, 100]	[50, 100, 200]
	Child-Sum LSTM	[10, 25, 40]	[150, 200, 300]	-	[50, 100, 200]
	Invariant CP-LSTM	[10, 25, 40]	[150, 200, 300]	[30, 50, 100]	[50, 100, 200]
	TreeNet*	[0.001, 0.005, 0.008]	[150, 200, 300]	-	[50, 100, 200]
TREC	Binary Sum-LSTM	[10, 25, 40]	[150, 200, 300]	-	[0, 50, 100]
	Binary CP-LSTM	[10, 25, 40]	[150, 200, 300]	[30, 50, 100]	[0, 50, 100]
	Child-Sum LSTM	[10, 25, 40]	[150, 200, 300]	-	[0, 50, 100]
	Invariant CP-LSTM	[10, 25, 40]	[150, 200, 300]	[30, 50, 100]	[0, 50, 100]
	TreeNet*	[0.001, 0.005, 0.008]	[150, 200, 300]	-	[0, 50, 100]

Table 2: Grids used for the model selection. The models with * validate the learning rate lr rather than the batch size bs .