

ORDERED NEURONS: INTEGRATING TREE STRUCTURES INTO RECURRENT NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Recurrent neural network (RNN) models are widely used for processing sequential data governed by a latent tree structure. Previous work shows that RNN models (especially Long Short-Term Memory (LSTM) based models) could learn to exploit the underlying tree structure. However, its performance consistently lags behind that of tree-based models. This work proposes a new inductive bias *Ordered Neurons*, which enforces an order of updating frequencies between hidden state neurons. We show that the ordered neurons could explicitly integrate the latent tree structure into recurrent models. To this end, we propose a new RNN unit: ON-LSTM, which achieve good performances on four different tasks: language modeling, unsupervised parsing, targeted syntactic evaluation, and logical inference.

1 INTRODUCTION

Natural language is usually presented in a sequential format, but the underlying structure of language is not strictly sequential. Linguists agree on a set of rules, or *syntax*, that governs this structure (Sandra & Taft, 2014), and the structure also dictates how the words compose to form components of sentences. This structure is usually tree-like, despite its presented form. Despite being discovered by linguistics, the real origin of the latent structure is unclear. Some theories point out that this could be related to an underlying mechanism of human cognition (Chomsky & Lightfoot, 2002). This possibility brings more interest in studying the latent structure with artificial neural network approaches, which are inspired by information processing and communication patterns in biological nervous systems.

From a practical point of view, integrating tree structure into a language model is also important for different reasons:

1. to obtain a hierarchical representation with increasing levels of abstraction, a key feature of deep neural networks (Bengio et al., 2009; LeCun et al., 2015; Schmidhuber, 2015);
2. to capture complex linguistic phenomena, like the long-term dependency problem (Tai et al., 2015) and the compositional effects (Socher et al., 2013);
3. to provide shortcut for gradient back-propagation (Chung et al., 2016).

Developing deep neural networks that can leverage syntactic knowledge, or at least some tree structure (Williams et al., 2018; Shi et al., 2018), to form better semantic representations have received a great deal of attention in recent years (Shen et al., 2017; Jacob et al., 2018; Bowman et al., 2016; Choi et al., 2018; Yogatama et al., 2016).

One straightforward way of obtaining the tree structure is through a supervised syntactic parser. Trees produced by these parsers have been used to guide the composition of word semantics into sentence semantics (Socher et al., 2013; Bowman et al., 2015), or even to help next word prediction given previous words (Wu et al., 2017). However, supervised parsers are limiting for several reasons: 1) few languages have comprehensive annotated data for supervised parser training; 2) in available language data, syntax rules tend to be broken (e.g. in tweets); and 3) languages change over time with use, so syntax rules may evolve.

On the other hand, learning the tree structure in an unsupervised manner from available data remains an open problem. Many such attempts suffer from inducing trivial structure (e.g., a left-branching or right-branching tree structure (Williams et al., 2018)), or the difficulty in training caused by resort to RL (Yogatama et al., 2016). Further, some methods are relatively complex to implement and train, like the PRPN from Shen et al. (2017).

Recurrent neural networks (RNNs) have proven highly effective at the task of language modeling (Merity et al., 2017; Melis et al., 2017). RNNs implicitly impose a chain structure on the data. This chain structure may seem at odds with the latent non-sequential structure of language and poses several difficulties for the processing of natural language data with deep learning methods, such as capturing long-term dependencies (Bengio et al., 2009), achieving good generalization (Bowman et al., 2015), handling negation (Socher et al., 2013), etc. Meanwhile, some evidence exists that an RNN with sufficient capacity has the potential to encode such a tree structure implicitly (Kuncoro et al., 2018). But, the question remains: Would imposing a tree-structure inductive prior on the model architecture result in better models of language?

In this work, we introduce a new inductive bias for recurrent neural networks: *Ordered Neurons*. This inductive bias enforces a dependency between the neurons that reflects the life cycle of information stored inside each neuron. In other words, some high-ranking neurons store long-term information, while low-ranking neurons store short-term information. To avoid a fixed division between high-ranking and low-ranking neurons, we further propose a new activation function `cumax()` to actively allocate neurons to store long/short-term information. Based on the `cumax()` and the LSTM architecture, we have designed a new model, ON-LSTM, that enables RNN models to perform tree-like compositions without breaking its sequential form. Our model achieve good performance on four tasks: language modeling, unsupervised constituency parsing, targeted syntactic evaluation (Marvin & Linzen, 2018) and logical inference (Bowman et al., 2015). The result on unsupervised constituency parsing task suggests that the proposed inductive bias aligns with the syntax principles proposed by human experts. The experiments also show that ON-LSTM performs better than standard LSTM models in terms of long-term dependency and longer sequence generalization.

2 RELATED WORK

There has been prior work leveraging tree structures for natural language tasks in the literature. Socher et al. (2010); Alvarez-Melis & Jaakkola (2016); Zhou et al. (2017); Zhang et al. (2015) uses labeled data from a treebank to perform supervised learning for inferring parse trees. Socher et al. (2013); Tai et al. (2015) explicitly models the tree-structure using parsing information from an external parser. Later, Bowman et al. (2016) used supervised signals from a parser (Klein & Manning, 2003) to train a stack-augmented neural network.

Theoretically, RNNs and LSTMs can model data produced by context-free grammars and context-sensitive grammars (Gers & Schmidhuber, 2001). However, recent results suggest that introducing structure information into an LSTM model is beneficial. Kuncoro et al. (2018) showed that RNNs (Dyer et al., 2016), which have an explicit bias to model the syntactic structures, outperform LSTMs on the subject-verb agreement task (Linzen et al., 2016). In our paper, we run a more extensive suite of grammatical tests provided by Marvin & Linzen (2018). Bowman et al. (2014; 2015) also demonstrate that these recursive structures work better for downstream, predictive tasks if the data was generated with such a structure. Interestingly, Shi et al. (2018) suggests that the prescribed grammar tree may not be ideal, but some sort of hierarchical structure, perhaps task dependent, might help. However, the problem of efficiently learning such structures from data remains an open question.

One possible solution would be to develop models with varying time-scales of recurrence as a way of emulating this hierarchy. There has been precedence for such models: El Hihi & Bengio (1996); Schmidhuber (1991); Lin et al. (1998) describe models that model data at different, pre-determined time-scales. More recently, Koutnik et al. (2014) segments an RNN hidden state with different time-scales for updating called the Clockwork RNN. These approaches typically make a strong assumption about the regularity of the hierarchy involved in modelling the data. Chung et al. (2016) proposed a method that, unlike the Clockwork RNN, would learn the multi-scale hierarchical recurrence. However, the model still has a pre-determined depth to the hierarchy, depending on the number of layers it was parameterised with.

In models developed specifically for language modelling, there has been precedent for incorporating syntactic structure for the task (Roark, 2001; Charniak, 2001; Chelba & Jelinek, 2000). More recently, Yogatama et al. (2018) implicitly learned structure by using a stack-like memory. While they did not perform analysis on its ability to induce a parse tree, the authors perform the Linzen et al. (2016) test on their model. Shen et al. (2017) introduced the Parsing-Reading-Predict Networks (PRPN) model, which attempts to perform parsing with only a language modelling signal. The model uses self-attention to compose previous states. They introduced a new value, syntactic distance, to control the range of attention. This value is then found to correspond to the depth of the parse tree. However, the added complexity in using the PRPN model makes it unwieldy in practice.

3 ORDERED NEURONS




Figure 1: The relationship between a constituency parse tree and an ON-LSTM. Given a sequence of tokens (x_1, x_2, x_3) , their constituency-based parse tree is illustrated in (a). (b) provides a block view of the tree structure, where S and VP node strides across more than one time step. The representation for high-ranking nodes should be relatively consistent across multiple time steps. (c) visualization of the ratio of updated neurons for each group of neurons at each time step. At each time step, given the input word, darker grey blocks are completely updated, lighter grey blocks are partially updated. The three groups of neurons have different update frequencies. Higher groups update less frequently and lower groups update more frequently.

Given a sequence of tokens x_1, \dots, x_T governed by a latent tree structure as shown in Figure 1(a), our goal is to infer the unobserved structure from observed tokens and compute a hidden state h_t for each time step t . One ideal interpretation for h_t is that it represents all nodes on the path between current leaf node x_t to the root node S. As shown in Figure 1(c), h_t contains representations for all constituents that include the current token x_t , even when the respective constituent is only partially observed. We can also further assume that different nodes are represented by different chunks of adjacent neurons in the hidden states. However, while the dimension of hidden states is fixed, the numbers of nodes on the path are different across different time steps and sentences. Thus, allowing the model to actively allocate different numbers of neurons to each node would allow more flexibility.

In our model, high-ranking nodes contain long-term/global information that will last anywhere from several time steps to the entire sentence, while low-ranking nodes contain only short-term/local information that only last one or a few time steps, as shown in Figure 1(b). It is also therefore important to allow the model to actively control the updating frequency of neurons to differentiate long/short-term information.

Given these requirements, we introduce a new inductive bias: *ordered neurons* to enable dynamic allocation of neurons to represent different time-scale dependencies by controlling the update frequency of neurons. The *ordered neurons* make the assumption that:

- A order should exist between neurons: the high-ranking neurons store long-term information, while the low-ranking neurons store short-term information. To erase (or update) high-ranking neurons, the model should first erase (or update) all lower-ranking neurons.
- This ordering is independent of the data, thus we can enforce it on hidden states as an inductive bias.

In other words, some neurons always update more (or less) frequently than the others, and that order is pre-determined as part of the model architecture.

4 ON-LSTM

In this section, we introduce a new RNN unit ON-LSTM, as an implementation of ordered neurons. The new model shares a similar architecture with the standard LSTM model:

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f) \quad (1)$$

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i) \quad (2)$$

$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \quad (3)$$

$$\hat{c}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c) \quad (4)$$

$$h_t = o_t \circ \tanh(c_t) \quad (5)$$

The only difference with the standard LSTM is that we exclude the update function for cell state c_t and replace it with a new update rule that will be explained in the following sections. The forget gates f_t and input gates i_t are used to control the erasing and writing operation on cell states c_t , as before. Since the gates in the standard LSTM do not impose a topology on the individual units in the gates, in general, the behavior of the individual cells does not reflect an ordering.

4.1 ACTIVATION FUNCTION: cumax()

To enforce an order to the update frequency, we introduce a new activation function:

$$\hat{g} = \text{cumax}(\dots) = \text{cumsum}(\text{softmax}(\dots)) \quad (6)$$

The vector \hat{g} can be seen as the expectation of a binary gate $g = (0, \dots, 0, 1, \dots, 1)$. This binary gate split the cell state into two segments: the 0-segment and the 1-segment. Thus, the model can apply different update rules on the two segments to differentiate long/short-term information. The index for the first 1 in g is parametrised as:

$$p(d) = \text{softmax}(\dots) \quad (7)$$

This discrete variable d represents the split point between the two segments. We can further compute the probability of the k -th value being 1, by evaluating the probability of the disjunction of any of the values before the k -th being the split point: $d \leq k = (d = 0) \vee (d = 1) \vee \dots \vee (d = k)$. Since the categories are mutually exclusive, we can do this by computing the cumulative distribution function,

$$p(g_k = 1) = p(d \leq k) = \sum_{i \leq k} p(d = i) \quad (8)$$

Ideally, g should take the form of discrete values. Unfortunately, computing gradient through a discrete value is not trivial, so in practice we use a relaxation in the form of computing the quantity $p(d \leq k)$ by computing a cumulative sum of the softmax. As g_k is binary, this is equivalent to computing $\mathbb{E}[g_k]$. Hence, $\hat{g} = \mathbb{E}[g]$.

4.2 STRUCTURED GATING MECHANISM

Based on the cumax() function, we introduce a master forget gate \tilde{f}_t and a master input gate \tilde{i}_t :

$$\tilde{f}_t = \text{cumax}(W_{\tilde{f}} x_t + U_{\tilde{f}} h_{t-1} + b_{\tilde{f}}) \quad (9)$$

$$\tilde{i}_t = 1 - \text{cumax}(W_{\tilde{i}} x_t + U_{\tilde{i}} h_{t-1} + b_{\tilde{i}}) \quad (10)$$

where the values in master forget gate are constrained to monotonously increase from 0 to 1, and those in master input gate monotonously decrease from 1 to 0. These gates serve as a high-level control unit for the update operations of cell states. Using the master gates, we define a new update rule,

$$\omega_t = \tilde{f}_t \circ \tilde{i}_t \quad (11)$$

$$\hat{f}_t = f_t \circ \omega_t + (\tilde{f}_t - \omega_t) = \tilde{f}_t \circ (f_t \circ \tilde{i}_t + 1 - \tilde{i}_t) \quad (12)$$

$$\hat{i}_t = i_t \circ \omega_t + (\tilde{i}_t - \omega_t) = \tilde{i}_t \circ (i_t \circ \tilde{f}_t + 1 - \tilde{f}_t) \quad (13)$$

$$c_t = \hat{f}_t \circ c_{t-1} + \hat{i}_t \circ \hat{c}_t \quad (14)$$

To explain the intuition behind the new update rule, we make the assumption that the master gates are binary.

- The master forget gate \tilde{f}_t controls the erasing behavior of the model. Suppose $\tilde{f}_t = (0, \dots, 0, 1, \dots, 1)$ and the split point is d_t^f . Given the Eq. (12) and (14), the information stored in the first d_t^f neurons of the previous cell states c_{t-1} will be completely erased. Assuming that the model learned the constituency parse as pictured in Figure 1(c), this has the effect of completing previous constituents. A large number of zeroed neurons, i.e. a large d_t^f , represents the end of a high-level constituent in a constituent-based parse tree, as most of the information will be discarded. Conversely, a small d_t^f conveys the end of a low-level constituent as high-level information is kept for further processing.
- The master input gate \tilde{i}_t is meant to control the writing behavior of model. Suppose $\tilde{i}_t = (1, \dots, 1, 0, \dots, 0)$ and the split point is d_t^i . Given Eq. (13) and (14), a large d_t^i means that the current input x_t contains long-term information that needs to be preserved for several time steps. Conversely, a small d_t^i means that the current input x_t just provides local information that could be erased by \tilde{f}_t in the next few time steps.
- The product of two master gates ω_t represents the overlap of \tilde{f}_t and \tilde{i}_t . When the overlap exists ($\exists k, \omega_{tk} > 0$), the segment is further controlled by the \tilde{f}_t and \tilde{i}_t in standard LSTM model to enable more fine-grained operations. This segment of neurons is related to the incomplete constituents that contain some previous words and the current input word x_t . For example, in figure 1, the word x_3 belongs to the constituents S and VP . At this time step, the overlap ω_3 would cover the related blocks of neurons, such that these neurons could be partial updated.

As the master gates only focus on coarse-grained control, modeling them with the same dimensions as the hidden states is computationally expensive and unnecessary. In practice, we parameterize \tilde{f}_t and \tilde{i}_t to be $D_m = \frac{D}{C}$ dimension vectors, where D is the dimension of hidden state, and C is a chunk size factor. We repeat each dimension C times, before the element-wise multiplication with f_t and i_t . The downsizing significantly reduces the number of extra parameters that we add to standard LSTM. This behavior means that every unit within each C -sized chunk receives the same gating behavior from the master gates.

5 EXPERIMENTS

We evaluate the proposed model on four tasks: language modeling, unsupervised constituency parsing, targeted syntactic evaluation (Marvin & Linzen, 2018), and logical inference (Bowman et al., 2015).

5.1 LANGUAGE MODELING

Word-level language modeling is a macroscopic evaluation of the model’s ability to deal with various linguistic phenomena (e.g. co-occurrence, syntactic structure, verb-subject agreement, etc). We evaluate our model by measuring perplexity on the Penn TreeBank (PTB) (Marcus et al., 1993; Mikolov, 2012) task.

For fair comparison, we closely follow the model hyper-parameters, regularization and optimization techniques introduced in AWD-LSTM (Merity et al., 2017). Our model uses a three-layer ON-LSTM model with 1150 units in the hidden layer and an embedding of size 400. For master gates, the downsize factor $C = 10$. The total number of parameters was slightly increased from 24 millions to 25 millions with additional matrices for computing master gates. We manually searched some of the dropout values for ON-LSTM based on the validation performance. The values used for dropout on the word vectors, the output between LSTM layers, the output of the final LSTM layer, and embedding dropout were (0.5, 0.3, 0.45, 0.1) respectively. A weight-dropout of 0.45 was applied to the recurrent weight matrices.

As shown in table 1, our model performs better than the standard LSTM while sharing the same number of layers, embedding dimensions, and hidden states units. Recall that the master gates only

Model	Parameters	Validation	Test
Zaremba et al. (2014) - LSTM (large)	66M	82.2	78.4
Gal & Ghahramani (2016) - Variational LSTM (large, MC)	66M	—	73.4
Kim et al. (2016) - CharCNN	19M	—	78.9
Merity et al. (2016) - Pointer Sentinel-LSTM	21M	72.4	70.9
Grave et al. (2016) - LSTM	—	—	82.3
Grave et al. (2016) - LSTM + continuous cache pointer	—	—	72.1
Inan et al. (2016) - Variational LSTM (tied) + augmented loss	51M	71.1	68.5
Zilly et al. (2016) - Variational RHN (tied)	23M	67.9	65.4
Zoph & Le (2016) - NAS Cell (tied)	54M	—	62.4
Shen et al. (2017) - PRPN-LM	—	—	62.0
Melis et al. (2017) - 4-layer skip connection LSTM (tied)	24M	60.9	58.3
Merity et al. (2017) - AWD-LSTM - 3-layer LSTM (tied)	24M	60.0	57.3
ON-LSTM - 3-layer (tied)	25M	58.29 ± 0.10	56.17 ± 0.12
Yang et al. (2017) - AWD-LSTM-MoS*	22M	56.5	54.4

Table 1: Single model perplexity on validation and test sets for the Penn Treebank language modeling task. Models noting *tied* use weight tying on the embedding and softmax weights. Model noting * focus on improving the softmax component of RNN language model. Their contribution is orthogonal to ours.

controls how information is stored in different neurons. Therefore, it is interesting to note that we can improve the performance of RNN model without skip connections or a significant increase in the number of parameters.

5.2 UNSUPERVISED CONSTITUENCY PARSING

The unsupervised constituency parsing task compares the latent tree structure induced by the model with those annotated by human experts. Following the experiment settings proposed in Htut et al. (2018), we take our best model for the language modeling task, and test it on WSJ10 dataset and WSJ test set. WSJ10 has 7422 sentences, filtered from the WSJ dataset with the constraint of 10 words or less, after the removal of punctuation and null elements (Klein & Manning, 2002). The WSJ test set contains 2416 sentences with various lengths. It is worth noting that the WSJ10 test set contains sentences from the training, validation, and test set of the PTB dataset, while WSJ test uses the same set of sentences as the PTB test set.

To generate a tree structure from the trained model and a sentence, we initialise the hidden states with $\mathbf{0}$, then feed the sentence into the model as in language modeling task. For each time step, we compute an estimation of d_t^f :

$$\hat{d}_t = \mathbb{E} \left[d_t^f \right] = \sum_{k=1}^{D_m} k p(y_t = k) = \sum_{k=1}^{D_m} \sum_{i=1}^k p(y_t = k) = D_m - \sum_{k=1}^{D_m} \tilde{f}_{tk} \quad (15)$$

Given \hat{d}_t , we can use the parsing algorithm proposed in Shen et al. (2017) for unsupervised constituency parsing.

The performance is shown in Table 2. The 2nd-layer of ON-LSTM model achieves state-of-the-art unsupervised constituency parsing results on the WSJ test set, while the 1st and 3rd layer of ON-LSTM do not perform as good. One possible interpretation is that only the first and last layers focus on fine-tuning the input and output vectors with local information, thus do not need to learn the entire tree structure. Since the WSJ test set contains sentence of various lengths which they are unobserved during training, we find that ON-LSTM provides better generalization and robustness toward longer sentences than previous models. We also see that ON-LSTM model can provide strong results for phrase detection, including ADJP (adjective phrases), PP (prepositional phrases), and NP (noun phrases). This feature could benefit many downstream tasks, like question-answering, named entity recognition, co-reference detection, etc.

5.3 TARGETED SYNTACTIC EVALUATION

Targeted syntactic evaluation is proposed in Marvin & Linzen (2018). The task evaluates language models along three different structure-sensitive linguistic phenomenon: subject-verb agreement, re-

Model	Training Data	Training Object	Vocab Size	Parsing F1				Depth WSJ	Accuracy on WSJ by Tag			
				WSJ10		WSJ			ADJP	NP	PP	INTJ
				μ (σ)	max	μ (σ)	max					
PRPN-UP	AIINLI Train	LM	76k	66.3 (0.8)	68.5	38.3 (0.5)	39.8	5.8	28.7	65.5	32.7	0.0
PRPN-LM	AIINLI Train	LM	76k	52.4 (4.9)	58.1	35.0 (5.4)	42.8	6.1	37.8	59.7	61.5	100.0
PRPN-UP	WSJ Train	LM	15.8k	62.2 (3.9)	70.3	26.0 (2.3)	32.8	5.8	24.8	54.4	17.8	0.0
PRPN-LM	WSJ Train	LM	10k	70.5 (0.4)	71.3	37.4 (0.3)	38.1	5.9	26.2	63.9	24.4	0.0
ON-LSTM 1st-layer	WSJ Train	LM	10k	35.2(4.1)	42.8	20.0(2.8)	24.0	5.6	38.1	23.8	18.3	100.0
ON-LSTM 2nd-layer	WSJ Train	LM	10k	65.1(1.7)	66.8	47.7(1.5)	49.4	5.6	46.2	61.4	55.4	0.0
ON-LSTM 3rd-layer	WSJ Train	LM	10k	54.0(3.9)	57.6	36.6(3.3)	40.4	5.3	44.8	57.5	47.2	0.0
300D ST-Gumbel	AIINLI Train	NLI	-	-	-	<i>19.0 (1.0)</i>	<i>20.1</i>	-	<i>15.6</i>	<i>18.8</i>	<i>9.9</i>	<i>59.4</i>
w/o Leaf GRU	AIINLI Train	NLI	-	-	-	22.8 (1.6)	25.0	-	18.9	24.1	14.2	51.8
300D RL-SPINN	AIINLI Train	NLI	-	-	-	<i>13.2 (0.0)</i>	<i>13.2</i>	-	<i>1.7</i>	<i>10.8</i>	<i>4.6</i>	<i>50.6</i>
w/o Leaf GRU	AIINLI Train	NLI	-	-	-	<i>13.1 (0.1)</i>	<i>13.2</i>	-	<i>1.6</i>	<i>10.9</i>	<i>4.6</i>	<i>50.0</i>
CCM	WSJ10 Full	-	-	-	71.9	-	-	-	-	-	-	-
DMV+CCM	WSJ10 Full	-	-	-	77.6	-	-	-	-	-	-	-
UML-DOP	WSJ10 Full	-	-	-	82.9	-	-	-	-	-	-	-
Random Trees	-	-	-	-	34.7	21.3 (0.0)	21.4	5.3	17.4	22.3	16.0	40.4
Balanced Trees	-	-	-	-	-	21.3 (0.0)	21.3	4.6	22.1	20.2	9.3	55.9
Left Branching	-	-	-	28.7	28.7	13.1 (0.0)	13.1	12.4	-	-	-	-
Right Branching	-	-	-	61.7	61.7	16.5 (0.0)	16.5	12.4	-	-	-	-

Table 2: Unlabeled parsing F1 results evaluated on full WSJ10 and WSJ test set. Our language model has three layers, each of them provides a sequence of \hat{d}_t . We provide the parsing performance for all layers. Results with RL-SPINN and ST-Gumbel are evaluated on the full WSJ (Williams et al., 2017). PRPN models are evaluated on WSJ test set (Htut et al., 2018). We run the model with 5 different random seeds to calculate the average F1. The *Accuracy* columns represent the fraction of ground truth constituents of a given type that corresponds to constituents in the model parses. We use the model with the best F1 score to report ADJP, NP, PP, and INTJ. WSJ10 baselines are from Klein & Manning (2002, CCM), Klein & Manning (2005, DMV+CCM), and Bod (2006, UML-DOP). As the WSJ10 baselines are trained using additional information such as POS tags and dependency parser, they are not strictly comparable with the latent tree learning results. Italics mark results that are worse than the random baseline.

flexive anaphora and negative polarity items. Given a large number of minimally different pairs of English sentences, each consisting of a grammatical and an ungrammatical sentence, a language model should assign a higher probability to a grammatical sentence than an ungrammatical one.

Using the released codebase¹ and the same settings proposed in Marvin & Linzen (2018), we train both the ON-LSTM and LSTM language models on a 90 million word subset of Wikipedia. The RNN LMs has two layers of 650 units, a batch size of 128, a dropout rate of 0.2, a learning rate of 20.0, and was trained for 40 epochs. The input embedding was 200 dimensions and the output embedding was 650 dimensions.

Table 3 shows that ON-LSTM perform better on long-term dependency cases, while LSTM is better on short-term ones. This is possibly due to the relatively small number of units in the hidden states, which is insufficient to take into account both long and short-term information. We also notice that the results for NPI test cases have unusually high variance across different hyper-parameters. This result maybe due to the non-syntactic cues discussed in Marvin & Linzen (2018). Despite this, ON-LSTM actually achieves better perplexity on the validation.

5.4 LOGICAL INFERENCE

We also analyze the model’s performance on the logical inference task described in Bowman et al. (2015). This task is based on a language that has a vocabulary of six words and three logical operations, *or*, *and*, *not*. There are seven mutually exclusive logical relations that describe the relationship between two sentences: two types of entailment, equivalence, exhaustive and non-exhaustive con-

¹https://github.com/BeckyMarvin/LM_syneval. We notice that the test set generated from the code is different from the one used in their paper. Hence, our results are not comparable with the results in Marvin & Linzen (2018).

	ON-LSTM	LSTM
Short-Term Dependency		
SUBJECT-VERB AGREEMENT:		
Simple	0.99	1.00
In a sentential complement	0.95	0.98
Short VP coordination	0.89	0.92
In an object relative clause	0.84	0.88
In an object relative (no <i>that</i>)	0.78	0.81
REFLEXIVE ANAPHORA:		
Simple	0.89	0.82
In a sentential complement	0.86	0.80
NEGATIVE POLARITY ITEMS:		
Simple (grammatical vs. intrusive)	0.18	1.00
Simple (intrusive vs. ungrammatical)	0.50	0.01
Simple (grammatical vs. ungrammatical)	0.07	0.63
Long-Term Dependency		
SUBJECT-VERB AGREEMENT:		
Long VP coordination	0.74	0.74
Across a prepositional phrase	0.67	0.68
Across a subject relative clause	0.66	0.60
Across an object relative clause	0.57	0.52
Across an object relative (no <i>that</i>)	0.54	0.51
REFLEXIVE ANAPHORA:		
Across a relative clause	0.57	0.58
NEGATIVE POLARITY ITEMS:		
Across a relative clause (grammatical vs. intrusive)	0.59	0.95
Across a relative clause (intrusive vs. ungrammatical)	0.20	0.00
Across a relative clause (grammatical vs. ungrammatical)	0.11	0.04

Table 3: Overall accuracy for the ON-LSTM and LSTM on each test case. “Long-term dependency” means that an unrelated phrase (or a clause) exist between the targeted pair of words, while “short-term dependency” means there is no such distraction.

tradition, and two types of semantic independence. Similar to the natural language inference task, this logical inference task requires the model to predict the correct label for given pair of sentences. The train/test split is as described in the original codebase², and 10% of training set is set aside as the validation set.




Figure 2: Test accuracy of the models, trained on short sequences (≤ 6) in logic data. The horizontal axis indicates the length of the sequence, and the vertical axis indicates the accuracy of models performance on the corresponding test set.

We evaluate the ON-LSTM and the standard LSTM on this dataset. Given a pair of sentences (s_1, s_2), we feed both sentences into an RNN encoder, taking the last hidden state (h_1, h_2) as the sentence embedding. The concatenation of ($h_1, h_2, h_1 \circ h_2, \text{abs}(h_1 - h_2)$) is used as input to a multi-layer classifier, which gives a probability distribution over seven labels. In our experiment, the RNN models were parameterised with 400 units in one hidden layer, and the input embedding size was 128. A dropout of 0.2 was applied between different layers. Both models are trained on sequences with 6 or less logical operations and tested on sequences with at most 12 operations.

Figure 2 shows the performance of ON-LSTM and standard LSTM on the logical inference task. While both models achieve nearly 100% accuracy on short sequences (≤ 3), ON-LSTM attains better performance on sequences longer

²<https://github.com/sleepinyourhat/vector-entailment>

then 3. The performance gap continues to increase on longer sequences (≥ 7) that were not present during training. Hence, the ON-LSTM model shows better generalization while facing structured data with various lengths and comparing to the standard LSTM. However, a recursive neural network model can achieve stronger performance on this dataset (Bowman et al., 2015), since they have structure information as input. We also include the result of RRNet from Jacob et al. (2018), which can induce the latent tree structure from downstream tasks. However, the results may not be comparable, because the hyper-parameters for training were not provided. The repetitive composition using the same function is better suited for this synthetic task.

6 CONCLUSION

In this paper, we propose the ordered neuron inductive bias. This unifies modelling tree structures and RNNs, through separately allocating hidden state neurons with long and short-term information. Based on this idea, we propose a new RNN unit, the ON-LSTM, which includes a new gating mechanism and a new activation function $\text{cumax}(\cdot)$. The model’s results on unsupervised constituency parsing result shows that the ON-LSTM induces the latent structure of natural language in a way that is coherent with human expert annotation. The inductive bias also enables ON-LSTM to achieve good performance on language modeling, long-term dependency, and logical inference tasks.

REFERENCES

- David Alvarez-Melis and Tommi S Jaakkola. Tree-structured decoding with doubly-recurrent neural networks. 2016.
- Yoshua Bengio et al. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.
- Rens Bod. An all-subtrees approach to unsupervised parsing. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pp. 865–872. Association for Computational Linguistics, 2006.
- Samuel R Bowman, Christopher Potts, and Christopher D Manning. Recursive neural networks can learn logical semantics. *arXiv preprint arXiv:1406.1827*, 2014.
- Samuel R Bowman, Christopher D Manning, and Christopher Potts. Tree-structured composition in neural networks without tree-structured architectures. *arXiv preprint arXiv:1506.04834*, 2015.
- Samuel R Bowman, Jon Gauthier, Abhinav Rastogi, Raghav Gupta, Christopher D Manning, and Christopher Potts. A fast unified model for parsing and sentence understanding. *arXiv preprint arXiv:1603.06021*, 2016.
- Eugene Charniak. Immediate-head parsing for language models. In *Proceedings of the 39th Annual Meeting on Association for Computational Linguistics*, pp. 124–131. Association for Computational Linguistics, 2001.
- Ciprian Chelba and Frederick Jelinek. Structured language modeling. *Computer Speech & Language*, 14(4):283–332, 2000.
- Jihun Choi, Kang Min Yoo, and Sang-goo Lee. Learning to compose task-specific tree structures. In *Proceedings of the 2018 Association for the Advancement of Artificial Intelligence (AAAI), and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, 2018.
- Noam Chomsky and David W Lightfoot. *Syntactic structures*. Walter de Gruyter, 2002.
- Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. Hierarchical multiscale recurrent neural networks. *arXiv preprint arXiv:1609.01704*, 2016.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A Smith. Recurrent neural network grammars. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 199–209, 2016.

- Salah El Hihi and Yoshua Bengio. Hierarchical recurrent neural networks for long-term dependencies. In *Advances in neural information processing systems*, pp. 493–499, 1996.
- Yarin Gal and Zoubin Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, pp. 1019–1027, 2016.
- Felix A Gers and E Schmidhuber. Lstm recurrent networks learn simple context-free and context-sensitive languages. *IEEE Transactions on Neural Networks*, 12(6):1333–1340, 2001.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. Improving neural language models with a continuous cache. *arXiv preprint arXiv:1612.04426*, 2016.
- Phu Mon Htut, Kyunghyun Cho, and Samuel R Bowman. Grammar induction with neural language models: An unusual replication. *arXiv preprint arXiv:1808.10000*, 2018.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. Tying word vectors and word classifiers: A loss framework for language modeling. *arXiv preprint arXiv:1611.01462*, 2016.
- Athul Paul Jacob, Zhouhan Lin, Alessandro Sordani, and Yoshua Bengio. 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, 2018.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M Rush. Character-aware neural language models. In *AAAI*, pp. 2741–2749, 2016.
- Dan Klein and Christopher D Manning. A generative constituent-context model for improved grammar induction. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pp. 128–135. Association for Computational Linguistics, 2002.
- Dan Klein and Christopher D Manning. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pp. 423–430. Association for Computational Linguistics, 2003.
- Dan Klein and Christopher D Manning. Natural language grammar induction with a generative constituent-context model. *Pattern recognition*, 38(9):1407–1419, 2005.
- Jan Koutnik, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber. A clockwork rnn. *arXiv preprint arXiv:1402.3511*, 2014.
- Adhiguna Kuncoro, Chris Dyer, John Hale, Dani Yogatama, Stephen Clark, and Phil Blunsom. Lstms can learn syntax-sensitive dependencies well, but modeling structure makes them better. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pp. 1426–1436, 2018.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Tsungnan Lin, Bill G Horne, Peter Tino, and C Lee Giles. Learning long-term dependencies is not as difficult with narx recurrent neural networks. Technical report, 1998.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. Assessing the ability of lstms to learn syntax-sensitive dependencies. *arXiv preprint arXiv:1611.01368*, 2016.
- Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. Building a large annotated corpus of english: The penn treebank. *Computational linguistics*, 19(2):313–330, 1993.
- Rebecca Marvin and Tal Linzen. Targeted syntactic evaluation of language models. *arXiv preprint arXiv:1808.09031*, 2018.
- Gábor Melis, Chris Dyer, and Phil Blunsom. On the state of the art of evaluation in neural language models. *arXiv preprint arXiv:1707.05589*, 2017.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.

- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and Optimizing LSTM Language Models. *arXiv preprint arXiv:1708.02182*, 2017.
- Tomáš Mikolov. Statistical language models based on neural networks. *Presentation at Google, Mountain View, 2nd April*, 2012.
- Brian Roark. Probabilistic top-down parsing and language modeling. *Computational linguistics*, 27(2):249–276, 2001.
- Dominiek Sandra and Marcus Taft. *Morphological Structure, Lexical Representation and Lexical Access (RLE Linguistics C: Applied Linguistics): A Special Issue of Language and Cognitive Processes*. Routledge, 2014.
- Jürgen Schmidhuber. Neural sequence chunkers. 1991.
- Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron Courville. Neural language modeling by jointly learning syntax and lexicon. *arXiv preprint arXiv:1711.02013*, 2017.
- Haoyue Shi, Hao Zhou, Jiaye Chen, and Lei Li. On tree-based neural sentence modeling. *arXiv preprint arXiv:1808.09644*, 2018.
- Richard Socher, Christopher D Manning, and Andrew Y Ng. Learning continuous phrase representations and syntactic parsing with recursive neural networks. In *Proceedings of the NIPS-2010 Deep Learning and Unsupervised Feature Learning Workshop*, volume 2010, pp. 1–9, 2010.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 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, 2013.
- Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*, 2017.
- Adina Williams, Andrew Drozdov*, and Samuel R Bowman. Do latent tree learning models identify meaningful structure in sentences? *Transactions of the Association of Computational Linguistics*, 6:253–267, 2018.
- Shuangzhi Wu, Dongdong Zhang, Nan Yang, Mu Li, and Ming Zhou. Sequence-to-dependency neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pp. 698–707, 2017.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W Cohen. Breaking the softmax bottleneck: A high-rank rnn language model. *arXiv preprint arXiv:1711.03953*, 2017.
- Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, and Wang Ling. Learning to compose words into sentences with reinforcement learning. *arXiv preprint arXiv:1611.09100*, 2016.
- Dani Yogatama, Yishu Miao, Gabor Melis, Wang Ling, Adhiguna Kuncoro, Chris Dyer, and Phil Blunsom. Memory architectures in recurrent neural network language models. 2018.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014.
- Xingxing Zhang, Liang Lu, and Mirella Lapata. Top-down tree long short-term memory networks. *arXiv preprint arXiv:1511.00060*, 2015.
- Ganbin Zhou, Ping Luo, Rongyu Cao, Yijun Xiao, Fen Lin, Bo Chen, and Qing He. Generative neural machine for tree structures. *CoRR*, 2017.

Julian Georg Zilly, Rupesh Kumar Srivastava, Jan Koutník, and Jürgen Schmidhuber. Recurrent highway networks. *arXiv preprint arXiv:1607.03474*, 2016.

Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.

A SAMPLE PARSES FROM THE MODEL WITH THE BEST PERPLEXITY




Figure A.1: *Left* parses are from the 2nd layer of the ON-LSTM model, *Right* parses are converted from human expert annotations (removing all punctuations).




Figure A.2: Expectations of d_t^f (blue bar) and d_t^i (orange bar) given by different layers of ON-LSTM