
Wider and Deeper, Cheaper and Faster: Tensorized LSTMs for Sequence Learning

Zhen He^{1,2}, Shaobing Gao³, Liang Xiao², Daxue Liu², Hangen He², and David Barber^{1,4*}

¹University College London, ²National University of Defense Technology, ³Sichuan University,
⁴Alan Turing Institute

Abstract

Long Short-Term Memory (LSTM) is a popular approach to boosting the ability of Recurrent Neural Networks to store longer term temporal information. The *capacity* of an LSTM network can be increased by widening and adding layers. However, usually the former introduces additional parameters, while the latter increases the runtime. As an alternative we propose the *Tensorized LSTM* in which the hidden states are represented by *tensors* and updated via a *cross-layer convolution*. By increasing the tensor size, the network can be widened efficiently without additional parameters since the parameters are shared across different locations in the tensor; by delaying the output, the network can be deepened implicitly with little additional runtime since deep computations for each timestep are merged into temporal computations of the sequence. Experiments conducted on five challenging sequence learning tasks show the potential of the proposed model.

1 Introduction

We consider the time-series prediction task of producing a desired output \mathbf{y}_t at each timestep $t \in \{1, \dots, T\}$ given an observed input sequence $\mathbf{x}_{1:t} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t\}$, where $\mathbf{x}_t \in \mathbb{R}^R$ and $\mathbf{y}_t \in \mathbb{R}^S$ are vectors¹. The Recurrent Neural Network (RNN) [17, 43] is a powerful model that learns how to use a hidden state vector $\mathbf{h}_t \in \mathbb{R}^M$ to encapsulate the relevant features of the entire input history $\mathbf{x}_{1:t}$ up to timestep t . Let $\mathbf{h}_{t-1}^{cat} \in \mathbb{R}^{R+M}$ be the concatenation of the current input \mathbf{x}_t and the previous hidden state \mathbf{h}_{t-1} :

$$\mathbf{h}_{t-1}^{cat} = [\mathbf{x}_t, \mathbf{h}_{t-1}] \quad (1)$$

The update of the hidden state \mathbf{h}_t is defined as:

$$\mathbf{a}_t = \mathbf{h}_{t-1}^{cat} \mathbf{W}^h + \mathbf{b}^h \quad (2)$$

$$\mathbf{h}_t = \phi(\mathbf{a}_t) \quad (3)$$

where $\mathbf{W}^h \in \mathbb{R}^{(R+M) \times M}$ is the weight, $\mathbf{b}^h \in \mathbb{R}^M$ the bias, $\mathbf{a}_t \in \mathbb{R}^M$ the hidden activation, and $\phi(\cdot)$ the element-wise tanh function. Finally, the output \mathbf{y}_t at timestep t is generated by:

$$\mathbf{y}_t = \varphi(\mathbf{h}_t \mathbf{W}^y + \mathbf{b}^y) \quad (4)$$

where $\mathbf{W}^y \in \mathbb{R}^{M \times S}$ and $\mathbf{b}^y \in \mathbb{R}^S$, and $\varphi(\cdot)$ can be any differentiable function, depending on the task.

However, this vanilla RNN has difficulties in modeling long-range dependencies due to the vanishing/exploding gradient problem [4]. Long Short-Term Memories (LSTMs) [19, 24] alleviate

*Corresponding authors: Shaobing Gao <gaoshaobing@scu.edu.cn> and Zhen He <hezhen.cs@gmail.com>.

¹Vectors are assumed to be in row form throughout this paper.

these problems by employing memory cells to preserve information for longer, and adopting gating mechanisms to modulate the information flow. Given the success of the LSTM in sequence modeling, it is natural to consider how to increase the complexity of the model and thereby increase the set of tasks for which the LSTM can be profitably applied.

We consider the *capacity* of a network to consist of two components: the *width* (the amount of information handled in parallel) and the *depth* (the number of computation steps) [5]. A naive way to widen the LSTM is to increase the number of units in a hidden layer; however, the parameter number scales quadratically with the number of units. To deepen the LSTM, the popular Stacked LSTM (sLSTM) stacks multiple LSTM layers [20]; however, runtime is proportional to the number of layers and information from the input is potentially lost (due to gradient vanishing/explosion) as it propagates vertically through the layers.

In this paper, we introduce a way to both widen and deepen the LSTM whilst keeping the parameter number and runtime largely unchanged. In summary, we make the following contributions:

- (a) We tensorize RNN hidden state vectors into higher-dimensional tensors which allow more flexible parameter sharing and can be widened more efficiently without additional parameters.
- (b) Based on (a), we merge RNN deep computations into its temporal computations so that the network can be deepened with little additional runtime, resulting in a *Tensorized RNN (tRNN)*.
- (c) We extend the tRNN to an LSTM, namely the *Tensorized LSTM (tLSTM)*, which integrates a novel memory cell convolution to help to prevent the vanishing/exploding gradients.

2 Method

2.1 Tensorizing Hidden States

It can be seen from (2) that in an RNN, the parameter number scales quadratically with the size of the hidden state. A popular way to limit the parameter number when widening the network is to organize parameters as higher-dimensional tensors which can be factorized into lower-rank sub-tensors that contain significantly fewer elements [6, 15, 18, 26, 32, 39, 46, 47, 51], which is known as tensor factorization. This implicitly widens the network since the hidden state vectors are in fact broadcast to interact with the tensorized parameters. Another common way to reduce the parameter number is to share a small set of parameters across different locations in the hidden state, similar to Convolutional Neural Networks (CNNs) [34, 35].

We adopt parameter sharing to cutdown the parameter number for RNNs, since compared with factorization, it has the following advantages: (i) *scalability*, i.e., the number of shared parameters can be set independent of the hidden state size, and (ii) *separability*, i.e., the information flow can be carefully managed by controlling the *receptive field*, allowing one to shift RNN deep computations to the temporal domain (see Sec. 2.2). We also explicitly tensorize the RNN hidden state vectors, since compared with vectors, tensors have a better: (i) *flexibility*, i.e., one can specify which dimensions to share parameters and then can just increase the size of those dimensions without introducing additional parameters, and (ii) *efficiency*, i.e., with higher-dimensional tensors, the network can be widened faster w.r.t. its depth when fixing the parameter number (see Sec. 2.3).

For ease of exposition, we first consider 2D tensors (matrices): we tensorize the hidden state $\mathbf{h}_t \in \mathbb{R}^M$ to become $\mathbf{H}_t \in \mathbb{R}^{P \times M}$, where P is the *tensor size*, and M the *channel size*. We locally-connect the first dimension of \mathbf{H}_t in order to share parameters, and fully-connect the second dimension of \mathbf{H}_t to allow global interactions. This is analogous to the CNN which fully-connects one dimension (e.g., the RGB channel for input images) to globally fuse different feature planes. Also, if one compares \mathbf{H}_t to the hidden state of a Stacked RNN (sRNN) (see Fig. 1(a)), then P is akin to the number of stacked hidden layers, and M the size of each hidden layer. We start to describe our model based on 2D tensors, and finally show how to strengthen the model with higher-dimensional tensors.

2.2 Merging Deep Computations

Since an RNN is already *deep* in its temporal direction, we can deepen an input-to-output computation by associating the input \mathbf{x}_t with a (delayed) future output. In doing this, we need to ensure that the output \mathbf{y}_t is *separable*, i.e., not influenced by any future input $\mathbf{x}_{t'}$ ($t' > t$). Thus, we concatenate the projection of \mathbf{x}_t to the *top* of the previous hidden state \mathbf{H}_{t-1} , then gradually shift the input




Figure 1: Examples of sRNN, tRNNs and tLSTMs. (a) A 3-layer sRNN. (b) A 2D tRNN without (–) feedback (F) connections, which can be thought as a *skewed* version of (a). (c) A 2D tRNN. (d) A 2D tLSTM without (–) memory (M) cell convolutions. (e) A 2D tLSTM. In each model, the blank circles in column 1 to 4 denote the hidden state at timestep $t-1$ to $t+2$, respectively, and the blue region denotes the receptive field of the current output \mathbf{y}_t . In (b)-(e), the outputs are delayed by $L-1=2$ timesteps, where $L=3$ is the depth.

information down when the temporal computation proceeds, and finally generate \mathbf{y}_t from the *bottom* of \mathbf{H}_{t+L-1} , where $L-1$ is the number of delayed timesteps for computations of *depth* L . An example with $L=3$ is shown in Fig. 1(b). This is in fact a *skewed* sRNN as used in [1] (also similar to [48]). However, our method does not need to change the network structure and also allows different kinds of interactions as long as the output is separable, e.g. one can increase the local connections and use feedback (see Fig. 1(c)), which can be beneficial for sRNNs [10]. In order to share parameters, we update \mathbf{H}_t using a convolution with a learnable kernel. In this manner we increase the complexity of the input-to-output mapping (by delaying outputs) and limit parameter growth (by sharing transition parameters using convolutions).

To describe the resulting tRNN model, let $\mathbf{H}_{t-1}^{cat} \in \mathbb{R}^{(P+1) \times M}$ be the concatenated hidden state, and $p \in \mathbb{Z}_+$ the *location* at a tensor. The channel vector $\mathbf{h}_{t-1,p}^{cat} \in \mathbb{R}^M$ at location p of \mathbf{H}_{t-1}^{cat} is defined as:

$$\mathbf{h}_{t-1,p}^{cat} = \begin{cases} \mathbf{x}_t \mathbf{W}^x + \mathbf{b}^x & \text{if } p = 1 \\ \mathbf{h}_{t-1,p-1} & \text{if } p > 1 \end{cases} \quad (5)$$

where $\mathbf{W}^x \in \mathbb{R}^{R \times M}$ and $\mathbf{b}^x \in \mathbb{R}^M$. Then, the update of tensor \mathbf{H}_t is implemented via a convolution:

$$\mathbf{A}_t = \mathbf{H}_{t-1}^{cat} \circledast \{\mathbf{W}^h, \mathbf{b}^h\} \quad (6)$$

$$\mathbf{H}_t = \phi(\mathbf{A}_t) \quad (7)$$

where $\mathbf{W}^h \in \mathbb{R}^{K \times M^i \times M^o}$ is the *kernel weight* of size K , with $M^i = M$ input channels and $M^o = M$ output channels, $\mathbf{b}^h \in \mathbb{R}^{M^o}$ is the *kernel bias*, $\mathbf{A}_t \in \mathbb{R}^{P \times M^o}$ is the hidden activation, and \circledast is the convolution operator (see Appendix A.1 for a more detailed definition). Since the kernel convolves across different hidden layers, we call it the *cross-layer convolution*. The kernel enables interaction, both bottom-up and top-down across layers. Finally, we generate \mathbf{y}_t from the channel vector $\mathbf{h}_{t+L-1,p} \in \mathbb{R}^M$ which is located at the *bottom* of \mathbf{H}_{t+L-1} :

$$\mathbf{y}_t = \varphi(\mathbf{h}_{t+L-1,p} \mathbf{W}^y + \mathbf{b}^y) \quad (8)$$

where $\mathbf{W}^y \in \mathbb{R}^{M \times S}$ and $\mathbf{b}^y \in \mathbb{R}^S$. To guarantee that the *receptive field* of \mathbf{y}_t only covers the current and previous inputs $\mathbf{x}_{1:t}$ (see Fig. 1(c)), L , P , and K should satisfy the constraint:

$$L = \left\lceil \frac{2P}{K - K \bmod 2} \right\rceil \quad (9)$$

where $\lceil \cdot \rceil$ is the ceil operation. For the derivation of (9), please see Appendix B.

We call the model defined in (5)-(8) the *Tensorized RNN (tRNN)*. The model can be widened by increasing the tensor size P , whilst the parameter number remains fixed (thanks to the convolution). Also, unlike the sRNN of runtime complexity $O(TL)$, tRNN breaks down the runtime complexity to $O(T+L)$, which means either increasing the sequence length T or the network depth L would not significantly increase the runtime.

2.3 Extending to LSTMs

To allow the t RNN to capture long-range temporal dependencies, one can straightforwardly extend it to an LSTM by replacing the t RNN tensor update equations of (6)-(7) as follows:

$$[\mathbf{A}_t^g, \mathbf{A}_t^i, \mathbf{A}_t^f, \mathbf{A}_t^o] = \mathbf{H}_{t-1}^{cat} \circledast \{\mathbf{W}^h, \mathbf{b}^h\} \quad (10)$$

$$[\mathbf{G}_t, \mathbf{I}_t, \mathbf{F}_t, \mathbf{O}_t] = [\phi(\mathbf{A}_t^g), \sigma(\mathbf{A}_t^i), \sigma(\mathbf{A}_t^f), \sigma(\mathbf{A}_t^o)] \quad (11)$$

$$\mathbf{C}_t = \mathbf{G}_t \odot \mathbf{I}_t + \mathbf{C}_{t-1} \odot \mathbf{F}_t \quad (12)$$

$$\mathbf{H}_t = \phi(\mathbf{C}_t) \odot \mathbf{O}_t \quad (13)$$

where the kernel $\{\mathbf{W}^h, \mathbf{b}^h\}$ is of size K , with $M^i=M$ input channels and $M^o=4M$ output channels, $\mathbf{A}_t^g, \mathbf{A}_t^i, \mathbf{A}_t^f, \mathbf{A}_t^o \in \mathbb{R}^{P \times M}$ are activations for the new content \mathbf{G}_t , input gate \mathbf{I}_t , forget gate \mathbf{F}_t , and output gate \mathbf{O}_t , respectively, $\sigma(\cdot)$ is the element-wise sigmoid function, \odot is the element-wise multiplication, and $\mathbf{C}_t \in \mathbb{R}^{P \times M}$ is the memory cell. However, since in (12) the previous memory cell \mathbf{C}_{t-1} is only gated along the temporal direction (see Fig. 1(d)), long-range dependencies from the input to output might be lost when the tensor size P becomes large.

Memory Cell Convolution. To capture long-range dependencies from multiple directions, we additionally introduce a novel *memory cell convolution*, by which the memory cells can have a larger receptive field (see Fig. 1(e)). We also dynamically generate this convolution kernel so that it is both time- and location-dependent, allowing for flexible control over long-range dependencies from different directions. This results in our t LSTM tensor update equations:

$$[\mathbf{A}_t^g, \mathbf{A}_t^i, \mathbf{A}_t^f, \mathbf{A}_t^o, \mathbf{A}_t^q] = \mathbf{H}_{t-1}^{cat} \circledast \{\mathbf{W}^h, \mathbf{b}^h\} \quad (14)$$

$$[\mathbf{G}_t, \mathbf{I}_t, \mathbf{F}_t, \mathbf{O}_t, \mathbf{Q}_t] = [\phi(\mathbf{A}_t^g), \sigma(\mathbf{A}_t^i), \sigma(\mathbf{A}_t^f), \sigma(\mathbf{A}_t^o), \varsigma(\mathbf{A}_t^q)] \quad (15)$$

$$\mathbf{W}_t^c(p) = \text{reshape}(\mathbf{q}_{t,p}, [K, 1, 1]) \quad (16)$$

$$\mathbf{C}_{t-1}^{conv} = \mathbf{C}_{t-1} \circledast \mathbf{W}_t^c(p) \quad (17)$$

$$\mathbf{C}_t = \mathbf{G}_t \odot \mathbf{I}_t + \mathbf{C}_{t-1}^{conv} \odot \mathbf{F}_t \quad (18)$$

$$\mathbf{H}_t = \phi(\mathbf{C}_t) \odot \mathbf{O}_t \quad (19)$$

where, in contrast to (10)-(13), the kernel $\{\mathbf{W}^h, \mathbf{b}^h\}$ has additional $\langle K \rangle$ output channels² to generate the activation $\mathbf{A}_t^q \in \mathbb{R}^{P \times \langle K \rangle}$ for the dynamic kernel bank $\mathbf{Q}_t \in \mathbb{R}^{P \times \langle K \rangle}$, $\mathbf{q}_{t,p} \in \mathbb{R}^{\langle K \rangle}$ is the vectorized adaptive kernel at the location p of \mathbf{Q}_t , and $\mathbf{W}_t^c(p) \in \mathbb{R}^{K \times 1 \times 1}$ is the dynamic kernel of size K with a single input/output channel, which is reshaped from $\mathbf{q}_{t,p}$ (see Fig. 2(a) for an illustration). In (17), *each channel* of the previous memory cell \mathbf{C}_{t-1} is convolved with $\mathbf{W}_t^c(p)$ whose values vary with p , forming a *memory cell convolution* (see Appendix A.2 for a more detailed definition), which produces a convolved memory cell $\mathbf{C}_{t-1}^{conv} \in \mathbb{R}^{P \times M}$. Note that in (15) we employ a softmax function $\varsigma(\cdot)$ to normalize the channel dimension of \mathbf{Q}_t , which, similar to [37], can stabilize the value of memory cells and help to prevent the vanishing/exploding gradients (see Appendix C for details).

The idea of dynamically generating network weights has been used in many works [6, 14, 15, 23, 44, 46], where in [14] location-dependent convolutional kernels are also dynamically generated to improve CNNs. In contrast to these works, we focus on broadening the receptive field of t LSTM memory cells. Whilst the flexibility is retained, fewer parameters are required to generate the kernel since the kernel is shared by different memory cell channels.

Channel Normalization. To improve training, we adapt Layer Normalization (LN) [3] to our t LSTM. Similar to the observation in [3] that LN does not work well in CNNs where channel vectors at different locations have very different statistics, we find that LN is also unsuitable for t LSTM where lower level information is near the input while higher level information is near the output. We




Figure 2: Illustration of generating the memory cell convolution kernel, where (a) is for 2D tensors and (b) for 3D tensors.

²The operator $\langle \cdot \rangle$ returns the cumulative product of all elements in the input variable.

therefore normalize the channel vectors at different locations with their own statistics, forming a *Channel Normalization (CN)*, with its operator $\text{CN}(\cdot)$:

$$\text{CN}(\mathbf{Z}; \Gamma, \mathbf{B}) = \widehat{\mathbf{Z}} \odot \Gamma + \mathbf{B} \quad (20)$$

where $\mathbf{Z}, \widehat{\mathbf{Z}}, \Gamma, \mathbf{B} \in \mathbb{R}^{P \times M^z}$ are the original tensor, normalized tensor, *gain* parameter, and *bias* parameter, respectively. The m^z -th channel of \mathbf{Z} , i.e. $z_{m^z} \in \mathbb{R}^P$, is normalized element-wisely:

$$\widehat{z}_{m^z} = (z_{m^z} - z^\mu) / z^\sigma \quad (21)$$

where $z^\mu, z^\sigma \in \mathbb{R}^P$ are the *mean* and *standard deviation* along the channel dimension of \mathbf{Z} , respectively, and $\widehat{z}_{m^z} \in \mathbb{R}^P$ is the m^z -th channel of $\widehat{\mathbf{Z}}$. Note that the number of parameters introduced by CN/LN can be neglected as it is very small compared to the number of other parameters in the model.

Using Higher-Dimensional Tensors. One can observe from (9) that when fixing the kernel size K , the tensor size P of a 2D t LSTM grows linearly w.r.t. its depth L . How can we expand the tensor volume more rapidly so that the network can be widened more efficiently? We can achieve this goal by leveraging higher-dimensional tensors. Based on previous definitions for 2D t LSTMs, we replace the 2D tensors with D -dimensional ($D > 2$) tensors, obtaining $\mathbf{H}_t, \mathbf{C}_t \in \mathbb{R}^{P_1 \times P_2 \times \dots \times P_{D-1} \times M}$ with the tensor size $\mathbf{P} = [P_1, P_2, \dots, P_{D-1}]$. Since the hidden states are no longer matrices, we concatenate the projection of \mathbf{x}_t to one *corner* of \mathbf{H}_{t-1} , and thus (5) is extended as:

$$\mathbf{h}_{t-1, \mathbf{p}}^{\text{cat}} = \begin{cases} \mathbf{x}_t \mathbf{W}^x + \mathbf{b}^x & \text{if } p_d = 1 \text{ for } d = 1, 2, \dots, D-1 \\ \mathbf{h}_{t-1, p-1} & \text{if } p_d > 1 \text{ for } d = 1, 2, \dots, D-1 \\ \mathbf{0} & \text{otherwise} \end{cases} \quad (22)$$

where $\mathbf{h}_{t-1, \mathbf{p}}^{\text{cat}} \in \mathbb{R}^M$ is the channel vector at location $\mathbf{p} \in \mathbb{Z}_+^{D-1}$ of the concatenated hidden state $\mathbf{H}_{t-1}^{\text{cat}} \in \mathbb{R}^{(P_1+1) \times (P_2+1) \times \dots \times (P_{D-1}+1) \times M}$. For the tensor update, the convolution kernel \mathbf{W}^h and $\mathbf{W}_t^c(\cdot)$ also increase their dimensionality with kernel size $\mathbf{K} = [K_1, K_2, \dots, K_{D-1}]$. Note that $\mathbf{W}_t^c(\cdot)$ is reshaped from the vector, as illustrated in Fig. 2(b). Correspondingly, we generate the output \mathbf{y}_t from the opposite *corner* of \mathbf{H}_{t+L-1} , and therefore (8) is modified as:

$$\mathbf{y}_t = \varphi(\mathbf{h}_{t+L-1, \mathbf{p}} \mathbf{W}^y + \mathbf{b}^y) \quad (23)$$

For convenience, we set $P_d = P$ and $K_d = K$ for $d = 1, 2, \dots, D-1$ so that all dimensions of \mathbf{P} and \mathbf{K} can satisfy (9) with the same depth L . In addition, CN still normalizes the channel dimension of tensors.

3 Experiments

We evaluate t LSTM on five challenging sequence learning tasks under different configurations:

- (a) **sLSTM (baseline)**: our implementation of sLSTM [21] with parameters shared across all layers.
- (b) **2D t LSTM**: the standard 2D t LSTM, as defined in (14)-(19).
- (c) **2D t LSTM-M**: removing (-) memory (M) cell convolutions from (b), as defined in (10)-(13).
- (d) **2D t LSTM-F**: removing (-) feedback (F) connections from (b).
- (e) **3D t LSTM**: tensorizing (b) into 3D t LSTM.
- (f) **3D t LSTM+LN**: applying (+) LN [3] to (e).
- (g) **3D t LSTM+CN**: applying (+) CN to (e), as defined in (20).

To compare different configurations, we also use L to denote the number of layers of a sLSTM, and M to denote the hidden size of each sLSTM layer. We set the kernel size K to 2 for 2D t LSTM-F and 3 for other t LSTMs, in which case we have $L = P$ according to (9).

For each configuration, we fix the parameter number and increase the tensor size to see if the performance of t LSTM can be boosted without increasing the parameter number. We also investigate how the runtime is affected by the depth, where the runtime is measured by the average GPU milliseconds spent by a *forward and backward pass over one timestep of a single example*. Next, we compare t LSTM against the state-of-the-art methods to evaluate its ability. Finally, we visualize the internal working mechanism of t LSTM. Please see Appendix D for training details.

3.1 Wikipedia Language Modeling

The Hutter Prize Wikipedia dataset [25] consists of 100 million characters taken from 205 different characters including alphabets, XML markups and special symbols. We model the dataset at the character-level, and try to predict the next character of the input sequence.

We fix the parameter number to 10M, corresponding to channel sizes M of 1120 for s LSTM and 2D t LSTM-F, 901 for other 2D t LSTMs, and 522 for 3D t LSTMs. All configurations are evaluated with depths $L = 1, 2, 3, 4$. We use Bits-per-character (BPC) to measure the model performance.

Results are shown in Fig. 3. When $L \leq 2$, s LSTM and 2D t LSTM-F outperform other models because of a larger M . With L increasing, the performances of s LSTM and 2D t LSTM-M improve but become saturated when $L \geq 3$, while t LSTMs with memory cell convolutions improve with increasing L and finally outperform both s LSTM and 2D t LSTM-M. When $L = 4$, 2D t LSTM-F is surpassed by 2D t LSTM, which is in turn surpassed by 3D t LSTM. The performance of 3D t LSTM+LN benefits from LN only when $L \leq 2$. However, 3D t LSTM+CN consistently improves 3D t LSTM with different L .

Whilst the runtime of s LSTM is almost proportional to L , it is nearly constant in each t LSTM configuration and largely independent of L .

We compare a larger model, i.e. a 3D t LSTM+CN with $L = 6$ and $M = 1200$, to the state-of-the-art methods on the test set, as reported in Table 1. Our model achieves 1.264 BPC with 50.1M parameters, and is competitive to the best performing methods [38, 54] with similar parameter numbers.

3.2 Algorithmic Tasks

(a) **Addition:** The task is to sum two 15-digit integers. The network first reads two integers with one digit per timestep, and then predicts the summation. We follow the processing of [30], where a symbol ‘-’ is used to delimit the integers as well as pad the input/target sequence. A 3-digit integer addition task is of the form:

Input: - 1 2 3 - 9 0 0 - - - - -
 Target: - - - - - 1 0 2 3 -

(b) **Memorization:** The goal of this task is to memorize a sequence of 20 random symbols. Similar to the addition task, we use 65 different




Figure 3: Performance and runtime of different configurations on Wikipedia.

Table 1: Test BPC on Wikipedia.

	BPC	# Param.
MI-LSTM [51]	1.44	≈17M
mLSTM [32]	1.42	≈20M
HyperLSTM+LN [23]	1.34	26.5M
HM-LSTM+LN [11]	1.32	≈35M
Large RHN [54]	1.27	≈46M
Large FS-LSTM-4 [38]	1.245	≈47M
2 × Large FS-LSTM-4 [38]	1.198	≈94M
3D t LSTM+CN ($L = 6, M = 1200$)	1.264	50.1M




Figure 4: Performance and runtime of different configurations on the addition (left) and memorization (right) tasks.

symbols. A 5-symbol memorization task is of the form:

Input: - a b c c b - - - - -
 Target: - - - - - a b c c b -

We evaluate all configurations with $L=1, 4, 7, 10$ on both tasks, where M is 400 for *addition* and 100 for *memorization*. The performance is measured by the symbol prediction accuracy.

Fig. 4 show the results. In both tasks, large L degrades the performances of *sLSTM* and 2D *tLSTM*-M. In contrast, the performance of 2D *tLSTM*-F steadily improves with L increasing, and is further enhanced by using feedback connections, higher-dimensional tensors, and CN, while LN helps only when $L=1$. Note that in both tasks, the correct solution can be found (when 100% test accuracy is achieved) due to the repetitive nature of the task. In our experiment, we also observe that for the addition task, 3D *tLSTM*+CN with $L=7$ outperforms other configurations and finds the solution with only 298K training samples, while for the memorization task, 3D *tLSTM*+CN with $L=10$ beats others configurations and achieves perfect memorization after seeing 54K training samples. Also, unlike in *sLSTM*, the runtime of all *tLSTMs* is largely unaffected by L .

We further compare the best performing configurations to the state-of-the-art methods for both tasks (see Table 2). Our models solve both tasks significantly faster (i.e., using fewer training samples) than other models, achieving the new state-of-the-art results.

Table 2: Test accuracies on two algorithmic tasks.

	Addition		Memorization	
	Acc.	# Samp.	Acc.	# Samp.
Stacked LSTM [21]	51%	5M	>50%	900K
Grid LSTM [30]	>99%	550K	>99%	150K
3D <i>tLSTM</i> +CN ($L=7$)	>99%	298K	>99%	115K
3D <i>tLSTM</i> +CN ($L=10$)	>99%	317K	>99%	54K

3.3 MNIST Image Classification

The MNIST dataset [35] consists of 50000/10000/10000 handwritten digit images of size 28×28 for training/validation/test. We have two tasks on this dataset:

(a) **Sequential MNIST**: The goal is to classify the digit after sequentially reading the pixels in a scan-line order [33]. It is therefore a 784 timestep sequence learning task where a single output is produced at the last timestep; the task requires very long range dependencies in the sequence.

(b) **Sequential Permuted MNIST**: We permute the original image pixels in a fixed random order as in [2], resulting in a permuted MNIST (*p*MNIST) problem that has even longer range dependencies across pixels and is harder.




Figure 5: Performance and runtime of different configurations on sequential MNIST (left) and sequential *p*MNIST (right).

In both tasks, all configurations are evaluated with $M=100$ and $L=1, 3, 5$. The model performance is measured by the classification accuracy.

Results are shown in Fig. 5. *sLSTM* and 2D *tLSTM*-M no longer benefit from the increased depth when $L=5$. Both increasing the depth and tensorization boost the performance of 2D *tLSTM*. However, removing feedback connections from 2D *tLSTM* seems not to affect the performance. On the other hand, CN enhances the 3D *tLSTM* and when $L \geq 3$ it outperforms LN. 3D *tLSTM*+CN with $L=5$ achieves the highest performances in both tasks, with a validation accuracy of 99.1% for MNIST and 95.6% for *p*MNIST. The runtime of *tLSTMs* is negligibly affected by L , and all *tLSTMs* become faster than *sLSTM* when $L=5$.




Figure 6: Visualization of the diagonal channel means of the t LSTM memory cells for each task. In each horizontal bar, the rows from top to bottom correspond to the diagonal locations from \mathbf{p}^{in} to \mathbf{p}^{out} , the columns from left to right correspond to different timesteps (from 1 to $T + L - 1$ for the full sequence, where $L - 1$ is the time delay), and the values are normalized to be in range $[0, 1]$ for better visualization. Both full sequences in (d) and (e) are zoomed out horizontally.

We also compare the configurations of the highest test accuracies to the state-of-the-art methods (see Table 3). For sequential MNIST, our 3D t LSTM+CN with $L = 3$ performs as well as the state-of-the-art Dilated GRU model [8], with a test accuracy of 99.2%. For the sequential p MNIST, our 3D t LSTM+CN with $L = 5$ has a test accuracy of 95.7%, which is close to the state-of-the-art of 96.7% produced by the Dilated CNN [40] in [8].

Table 3: Test accuracies (%) on sequential MNIST/ p MNIST.

	MNIST	p MNIST
i RNN [33]	97.0	82.0
LSTM [2]	98.2	88.0
u RNN [2]	95.1	91.4
Full-capacity u RNN [49]	96.9	94.1
s TANH [53]	98.1	94.0
BN-LSTM [13]	99.0	95.4
Dilated GRU [8]	99.2	94.6
Dilated CNN [40] in [8]	98.3	96.7
3D t LSTM+CN ($L = 3$)	99.2	94.9
3D t LSTM+CN ($L = 5$)	99.0	95.7

3.4 Analysis

The experimental results of different model configurations on different tasks suggest that the performance of t LSTMs can be improved by increasing the tensor size and network depth, requiring no additional parameters and little additional runtime. As the network gets wider and deeper, we found that the memory cell convolution mechanism is crucial to maintain improvement in performance. Also, we found that feedback connections are useful for tasks of sequential output (e.g., our Wikipedia and algorithmic tasks). Moreover, t LSTM can be further strengthened via tensorization or CN.

It is also intriguing to examine the internal working mechanism of t LSTM. Thus, we visualize the memory cell which gives insight into how information is routed. For each task, the best performing t LSTM is run on a random example. We record the channel mean (the mean over channels, e.g., it is of size $P \times P$ for 3D t LSTMs) of the memory cell at each timestep, and visualize the diagonal values of the channel mean from location $\mathbf{p}^{in} = [1, 1]$ (near the input) to $\mathbf{p}^{out} = [P, P]$ (near the output).

Visualization results in Fig. 6 reveal the distinct behaviors of t LSTM when dealing with different tasks: (i) Wikipedia: the input can be carried to the output location with less modification if it is sufficient to determine the next character, and vice versa; (ii) addition: the first integer is gradually encoded into memories and then interacts (performs addition) with the second integer, producing the sum; (iii) memorization: the network behaves like a shift register that continues to move the input symbol to the output location at the correct timestep; (iv) sequential MNIST: the network is more sensitive to the pixel value change (representing the contour, or topology of the digit) and can gradually accumulate evidence for the final prediction; (v) sequential p MNIST: the network is sensitive to high value pixels (representing the foreground digit), and we conjecture that this is because the permutation destroys the topology of the digit, making each high value pixel potentially important.

From Fig. 6 we can also observe common phenomena for all tasks: (i) at each timestep, the values at different tensor locations are markedly different, implying that wider (larger) tensors can encode more information, with less effort to compress it; (ii) from the input to the output, the values become increasingly distinct and are shifted by time, revealing that deep computations are indeed performed together with temporal computations, with long-range dependencies carried by memory cells.




Figure 7: Examples of models related to t LSTMs. (a) A single layer cLSTM [48] with vector array input. (b) A 3-layer sLSTM [21]. (c) A 3-layer Grid LSTM [30]. (d) A 3-layer RHN [54]. (e) A 3-layer QRNN [7] with kernel size 2, where costly computations are done by temporal convolution.

4 Related Work

Convolutional LSTMs. Convolutional LSTMs (cLSTMs) are proposed to parallelize the computation of LSTMs when the input at each timestep is *structured* (see Fig. 7(a)), e.g., a vector array [48], a vector matrix [41, 42, 50, 52], or a vector tensor [9, 45]. Unlike cLSTMs, t LSTM aims to increase the capacity of LSTMs when the input at each timestep is *non-structured*, i.e., a single vector, and is advantageous over cLSTMs in that: (i) it performs the convolution across different hidden layers whose structure is independent of the input structure, and integrates information bottom-up and top-down; while cLSTM performs the convolution within each hidden layer whose structure is coupled with the input structure, thus will fall back to the vanilla LSTM if the input at each timestep is a single vector; (ii) it can be widened efficiently without additional parameters by increasing the tensor size; while cLSTM can be widened by increasing the kernel size or kernel channel, which significantly increases the number of parameters; (iii) it can be deepened with little additional runtime by delaying the output; while cLSTM can be deepened by using more hidden layers, which significantly increases the runtime; (iv) it captures long-range dependencies from multiple directions through the memory cell convolution; while cLSTM struggles to capture long-range dependencies from multiple directions since memory cells are only gated along one direction.

Deep LSTMs. Deep LSTMs (d LSTMs) extend sLSTMs by making them deeper (see Fig. 7(b)-(d)). To keep the parameter number small and ease training, Graves [22], Kalchbrenner *et al.* [30], Mujika *et al.* [38], Zilly *et al.* [54] apply another RNN/LSTM along the *depth* direction of d LSTMs, which, however, multiplies the runtime. Though there are implementations to accelerate the deep computation [1, 16], they generally aim at simple architectures such as sLSTMs. Compared with d LSTMs, t LSTM performs the deep computation with little additional runtime, and employs a cross-layer convolution to enable the feedback mechanism. Moreover, the capacity of t LSTM can be increased more efficiently by using higher-dimensional tensors, whereas in d LSTM all hidden layers as a whole only equal to a 2D tensor (i.e., a stack of hidden vectors), the dimensionality of which is fixed.

Other Parallelization Methods. Some methods [7, 8, 28, 29, 36, 40] parallelize the temporal computation of the sequence (e.g., use the temporal convolution, as in Fig. 7(e)) during training, in which case full input/target sequences are accessible. However, during the online inference when the input presents sequentially, temporal computations can no longer be parallelized and will be blocked by deep computations of each timestep, making these methods potentially unsuitable for real-time applications that demand a high sampling/output frequency. Unlike these methods, t LSTM can speed up not only training but also online inference for many tasks since it performs the deep computation by the temporal computation, which is also human-like: we convert each signal to an action and *meanwhile* receive new signals in a non-blocking way. Note that for the online inference of tasks that use the previous output y_{t-1} for the current input x_t (e.g., autoregressive sequence generation), t LSTM cannot parallel the deep computation since it requires to delay $L-1$ timesteps to get y_{t-1} .

5 Conclusion

We introduced the Tensorized LSTM, which employs tensors to share parameters and utilizes the temporal computation to perform the deep computation for sequential tasks. We validated our model on a variety of tasks, showing its potential over other popular approaches.

Acknowledgements

This work is supported by the NSFC grant 91220301, the Alan Turing Institute under the EPSRC grant EP/N510129/1, and the China Scholarship Council.

References

- [1] Jeremy Appleyard, Tomas Kocisky, and Phil Blunsom. Optimizing performance of recurrent neural networks on gpus. *arXiv preprint arXiv:1604.01946*, 2016. 3, 9
- [2] Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *ICML*, 2016. 7, 8
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4, 5
- [4] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE TNN*, 5(2):157–166, 1994. 1
- [5] Yoshua Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2009. 2
- [6] Luca Bertinetto, João F Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In *NIPS*, 2016. 2, 4
- [7] James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. Quasi-recurrent neural networks. In *ICLR*, 2017. 9
- [8] Shiyu Chang, Yang Zhang, Wei Han, Mo Yu, Xiaoxiao Guo, Wei Tan, Xiaodong Cui, Michael Witbrock, Mark Hasegawa-Johnson, and Thomas Huang. Dilated recurrent neural networks. In *NIPS*, 2017. 8, 9
- [9] Jianxu Chen, Lin Yang, Yizhe Zhang, Mark Alber, and Danny Z Chen. Combining fully convolutional and recurrent neural networks for 3d biomedical image segmentation. In *NIPS*, 2016. 9
- [10] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Gated feedback recurrent neural networks. In *ICML*, 2015. 3, 13
- [11] Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. Hierarchical multiscale recurrent neural networks. In *ICLR*, 2017. 6
- [12] Ronan Collobert, Koray Kavukcuoglu, and Clément Farabet. Torch7: A matlab-like environment for machine learning. In *NIPS Workshop*, 2011. 13
- [13] Tim Cooijmans, Nicolas Ballas, César Laurent, and Aaron Courville. Recurrent batch normalization. In *ICLR*, 2017. 8
- [14] Bert De Brabandere, Xu Jia, Tinne Tuytelaars, and Luc Van Gool. Dynamic filter networks. In *NIPS*, 2016. 4
- [15] Misha Denil, Babak Shakibi, Laurent Dinh, Nando de Freitas, et al. Predicting parameters in deep learning. In *NIPS*, 2013. 2, 4
- [16] Greg Diamos, Shubho Sengupta, Bryan Catanzaro, Mike Chrzanowski, Adam Coates, Erich Elsen, Jesse Engel, Awni Hannun, and Sanjeev Satheesh. Persistent rnns: Stashing recurrent weights on-chip. In *ICML*, 2016. 9
- [17] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. 1
- [18] Timur Garipov, Dmitry Podoprikin, Alexander Novikov, and Dmitry Vetrov. Ultimate tensorization: compressing convolutional and fc layers alike. In *NIPS Workshop*, 2016. 2
- [19] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural computation*, 12(10):2451–2471, 2000. 1
- [20] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, 2013. 2
- [21] Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013. 5, 7, 9
- [22] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016. 9
- [23] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. In *ICLR*, 2017. 4, 6
- [24] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 1
- [25] Marcus Hutter. The human knowledge compression contest. *URL <http://prize.hutter1.net>*, 2012. 6
- [26] Ozan Irsoy and Claire Cardie. Modeling compositionality with multiplicative recurrent neural networks. In *ICLR*, 2015. 2

- [27] Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. An empirical exploration of recurrent network architectures. In *ICML*, 2015. 13
- [28] Łukasz Kaiser and Samy Bengio. Can active memory replace attention? In *NIPS*, 2016. 9
- [29] Łukasz Kaiser and Ilya Sutskever. Neural gpu learn algorithms. In *ICLR*, 2016. 9
- [30] Nal Kalchbrenner, Ivo Danihelka, and Alex Graves. Grid long short-term memory. In *ICLR*, 2016. 6, 7, 9, 13
- [31] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 13
- [32] Ben Krause, Liang Lu, Iain Murray, and Steve Renals. Multiplicative lstm for sequence modelling. In *ICLR Workshop*, 2017. 2, 6
- [33] Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015. 7, 8
- [34] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989. 2
- [35] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2, 7
- [36] Tao Lei and Yu Zhang. Training rnns as fast as cnns. *arXiv preprint arXiv:1709.02755*, 2017. 9
- [37] Gundram Leifert, Tobias Strauß, Tobias Grüning, Welf Wustlich, and Roger Labahn. Cells in multidimensional recurrent neural networks. *JMLR*, 17(1):3313–3349, 2016. 4, 13
- [38] Asier Mujika, Florian Meier, and Angelika Steger. Fast-slow recurrent neural networks. In *NIPS*, 2017. 6, 9
- [39] Alexander Novikov, Dmitrii Podoprikin, Anton Osokin, and Dmitry P Vetrov. Tensorizing neural networks. In *NIPS*, 2015. 2
- [40] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016. 8, 9
- [41] Viorica Patraucean, Ankur Handa, and Roberto Cipolla. Spatio-temporal video autoencoder with differentiable memory. In *ICLR Workshop*, 2016. 9
- [42] Bernardino Romera-Paredes and Philip Hilaire Sean Torr. Recurrent instance segmentation. In *ECCV*, 2016. 9
- [43] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. 1
- [44] Jürgen Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992. 4
- [45] Marijn F Stollenga, Wonmin Byeon, Marcus Liwicki, and Juergen Schmidhuber. Parallel multi-dimensional lstm, with application to fast biomedical volumetric image segmentation. In *NIPS*, 2015. 9
- [46] Ilya Sutskever, James Martens, and Geoffrey E Hinton. Generating text with recurrent neural networks. In *ICML*, 2011. 2, 4
- [47] Graham W Taylor and Geoffrey E Hinton. Factored conditional restricted boltzmann machines for modeling motion style. In *ICML*, 2009. 2
- [48] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016. 3, 9
- [49] Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. Full-capacity unitary recurrent neural networks. In *NIPS*, 2016. 8
- [50] Lin Wu, Chunhua Shen, and Anton van den Hengel. Deep recurrent convolutional networks for video-based person re-identification: An end-to-end approach. *arXiv preprint arXiv:1606.01609*, 2016. 9
- [51] Yuhuai Wu, Saizheng Zhang, Ying Zhang, Yoshua Bengio, and Ruslan Salakhutdinov. On multiplicative integration with recurrent neural networks. In *NIPS*, 2016. 2, 6
- [52] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *NIPS*, 2015. 9
- [53] Saizheng Zhang, Yuhuai Wu, Tong Che, Zhouhan Lin, Roland Memisevic, Ruslan R Salakhutdinov, and Yoshua Bengio. Architectural complexity measures of recurrent neural networks. In *NIPS*, 2016. 8
- [54] Julian Georg Zilly, Rupesh Kumar Srivastava, Jan Koutník, and Jürgen Schmidhuber. Recurrent highway networks. In *ICML*, 2017. 6, 9