

Learning Compact Recurrent Neural Networks with Block-Term Tensor Decomposition

Jinmian Ye¹, Linnan Wang², Guangxi Li¹, Di Chen¹, Shandian Zhe³, Xinqi Chu⁴ and Zenglin Xu¹

¹{jinmian.y, gxli2017, chendi1995425, zenglin}@gmail.com, SMILE Lab, University of Electronic Science and Technology of China

²wangnan318@gmail.com, Department of Computer Science, Brown University

³zhe@cs.utah.edu, School of Computing, University of Utah

⁴ethan@xjeralabs.com, Xjera Labs, Pte.Ltd

Abstract

Recurrent Neural Networks (RNNs) are powerful sequence modeling tools. However, when dealing with high dimensional inputs, the training of RNNs becomes computational expensive due to the large number of model parameters. This hinders RNNs from solving many important computer vision tasks, such as Action Recognition in Videos and Image Captioning. To overcome this problem, we propose a compact and flexible structure, namely Block-Term tensor decomposition, which greatly reduces the parameters of RNNs and improves their training efficiency. Compared with alternative low-rank approximations, such as tensor-train RNN (TT-RNN), our method, Block-Term RNN (BT-RNN), is not only more concise (when using the same rank), but also able to attain a better approximation to the original RNNs with much fewer parameters. On three challenging tasks, including Action Recognition in Videos, Image Captioning and Image Generation, BT-RNN outperforms TT-RNN and the standard RNN in terms of both prediction accuracy and convergence rate. Specifically, BT-LSTM utilizes 17,388 times fewer parameters than the standard LSTM to achieve an accuracy improvement over 15.6% in the Action Recognition task on the UCF11 dataset.

1. Introduction

Best known for the sequence-to-sequence learning, the Recurrent Neural Networks (RNNs) belong to a class of neural architectures designed to capture the dynamic temporal behaviors of data. The vanilla fully connected RNN utilizes a feedback loop to memorize previous information, while it is inept to handle long sequences as the gradient exponentially vanishes along the time [13, 2]. Unlike the vanilla RNNs passing information between layers with direct matrix-vector multiplications, the Long Short-Term

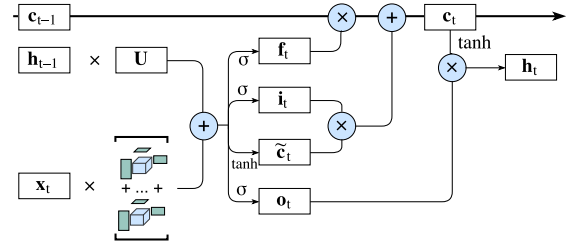


Figure 1: Architecture of BT-LSTM. The redundant dense connections between input and hidden state is replaced by low-rank BT representation.

Memory (LSTM) introduces a number of gates and passes information with element-wise operations [14]. This improvement drastically alleviates the gradient vanishing issue; therefore LSTM and its variants, e.g. Gated Recurrent Unit (GRU) [5], are widely used in various Computer Vision (CV) tasks [3, 22, 37] to model the long-term correlations in sequences.

The current formulation of LSTM, however, suffers from an excess of parameters, making it notoriously difficult to train and susceptible to overfitting. The formulation of LSTM can be described by the following equations:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \cdot \mathbf{x}_t + \mathbf{U}_f \cdot \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (1)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \cdot \mathbf{x}_t + \mathbf{U}_i \cdot \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (2)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \cdot \mathbf{x}_t + \mathbf{U}_o \cdot \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (3)$$

$$\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_c \cdot \mathbf{x}_t + \mathbf{U}_c \cdot \mathbf{h}_{t-1} + \mathbf{b}_c) \quad (4)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t \quad (5)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t), \quad (6)$$

where \odot denotes the element-wise product, $\sigma(\cdot)$ denotes the sigmoid function and $\tanh(\cdot)$ is the hyperbolic tangent function. The weight matrices \mathbf{W}_* and \mathbf{U}_* transform the input \mathbf{x}_t and the hidden state \mathbf{h}_{t-1} , respectively, to cell up-

date \tilde{c}_t and three gates f_t , i_t , and o_t . Please note that given an image feature vector \mathbf{x}_t fetch from a Convolutional Neural Network (CNN) network, the shape of \mathbf{x}_t will raise to $I = 4096$ and $I = 14 \times 14 \times 512$ w.r.t vgg16 [33] and Inception v4 [35]. If the number of hidden states is $J = 256$, the total number of parameters in calculating the four \mathbf{W}_* is $4 \times I \times J$, which can up to 4.1×10^6 and 1.0×10^8 , respectively. Therefore, the giant matrix-vector multiplication, i.e., $\mathbf{W}_* \cdot \mathbf{x}_t$, leads to the major inefficiency – the current parameter-intensive design not only subjects the model difficult to train, but also lead to high computation complexity and memory usage.

In addition, each $\mathbf{W}_* \cdot \mathbf{x}_t$ essentially represents a fully connected operation that transforms the input vector \mathbf{x}_t into the hidden state vector. However, extensive research on CNNs has proven that the dense connection is significantly inefficient at extracting the spatially latent local structures and local correlations naturally exhibited in the image [20, 10]. Recent leading CNNs architectures, e.g., DenseNet [15], ResNet [11] and Inception v4 [35], also try to circumvent one huge cumbersome dense layer [36]. But the discussions of improving the dense connections in RNNs are still quite limited [26, 30]. It is imperative to seek a more efficient design to replace $\mathbf{W}_* \cdot \mathbf{x}_t$.

In this work, we propose to design a sparsely connected tensor representation, i.e., the Block-Term decomposition (BTD) [7], to replace the redundant and densely connected operation in LSTM¹. The Block-Term decomposition is a low-rank approximation method that decomposes a high-order tensor into a sum of multiple Tucker decomposition models [39, 44, 45, 21]. In detail, we represent the four weight matrices (i.e., \mathbf{W}_*) and the input data \mathbf{x}_t into a various order of tensor. In the process of RNNs training, the BTD layer automatically learns inter-parameter correlations to implicitly prune redundant dense connections rendered by $\mathbf{W} \cdot \mathbf{x}$. By plugging the new BTD layer into current RNNs formulations, we present a new BT-RNN model with a similar representation power but several orders of fewer parameters. The refined LSTM model with the Block-term representation is illustrated in Fig. 1.

The major merits of BT-RNN are shown as follows:

- The low-rank BTD can compress the dense connections in the input-to-hidden transformation, while still retaining the current design philosophy of LSTM. By reducing several orders of model parameters, BT-LSTM has better convergence rate than the traditional LSTM architecture, significantly enhancing the training speed.
- Each dimension in the input data can share weights with all the other dimensions as the existence of core tensors, thus BT representation has the strong connection between different dimensions, enhancing the

¹we focus on LSTM in this paper, but the proposed approach also applies for other variants such as GRU.

ability to capture sufficient local correlations. Empirical results show that, compared with the Tensor Train model [29], the BT model has a better representation power with the same amount of model parameters.

- The design of multiple Tucker models can significantly reduce the sensitivity to noisy input data and widen network, leading to a more robust RNN model. In contrary to the Tensor Train based tensor approaches [47, 28], the BT model does not suffer from the difficulty of ranks setting, releasing researchers from intolerable work in choosing hyper-parameters.

In order to demonstrate the performance of the BT-LSTM model, we design three challenging computer vision tasks – Action Recognition in Videos, Image Caption and Image Generation – to quantitatively and qualitatively evaluate the proposed BT-LSTM against the baseline LSTM and other low-rank variants such as the Tensor Train LSTM (TT-LSTM). Experimental results have demonstrated the promising performance of the BT-LSTM model.

2. Related Work

The poor image modeling efficiency of full connections in the perception architecture, i.e., $\mathbf{W} \cdot \mathbf{x}$ [41], has been widely recognized by the Computer Vision (CV) community. The most prominent example is the great success made by Convolutional Neural Networks (CNNs) for the general image recognition. Instead of using the dense connections in multi-layer perceptions, CNNs relies on sparsely connected convolutional kernels to extract the latent regional features in an image. Hence, going sparse on connections is the key to the success of CNNs [8, 16, 27, 12, 34]. Though extensive discussions toward the efficient CNNs design, the discussions of improving the dense connections in RNNs are still quite limited [26, 30].

Compared with aforementioned explicit structure changes, the low-rank method is one orthogonal approach to implicitly prune the dense connections. Low-rank tensor methods have been successfully applied to address the redundant dense connection problem in CNNs [28, 47, 1, 38, 18]. Since the key operation in one perception is $\mathbf{W} \cdot \mathbf{x}$, Sainath et al. [31] decompose \mathbf{W} with Singular Value Decomposition (SVD), reducing up to 30% parameters in \mathbf{W} , but also demonstrates up to 10% accuracy loss [46]. The accuracy loss majorly results from losing the high-order spatial information, as intermediate data after image convolutions are intrinsically in 4D.

In order to capture the high order spatial correlations, recently, tensor methods were introduced into Neural Networks to approximate $\mathbf{W} \cdot \mathbf{x}$. For example, Tensor Train (TT) method was employed to alleviate the large computation $\mathbf{W} \cdot \mathbf{x}$ and reduce the number of parameters [28, 47, 38]. Yu et al. [48] also used a tensor train representation to forecast long-term information. Since this approach targets in long historic states, it increases additional parameters, leading to a difficulty in training. Other tensor decomposition

methods also applied in Deep Neural Networks (DNNs) for various purposes [19, 49, 18].

Although TT decomposition has obtained a great success in addressing dense connections problem, there are some limitations which block TT method to achieve better performance: 1) The optimal setting of TT-ranks is that they are small in the border cores and large in middle cores, e.g., like an olive [50]. However, in most applications, TT-ranks are set equally, which will hinder TT's representation ability. 2) TT-ranks has a strong constraint that the rank in border tensors must set to 1 ($R_1 = R_{d+1} = 1$), leading to a seriously limited representation ability and flexibility [47, 50].

Instead of difficultly finding the optimal TT-ranks setting, BTM has these advantages: 1) Tucker decomposition introduces a core tensor to represent the correlations between different dimensions, achieving better weight sharing. 2) ranks in core tensor can be set to equal, avoiding unbalance weight sharing in different dimensions, leading to a robust model toward different permutations of input data. 3) BTM uses a sum of multiple Tucker models to approximate a high-order tensor, breaking a large Tucker decomposition to several smaller models, widening network and increasing representation ability. Meanwhile, multiple Tucker models also lead to a more robust RNN model to noisy input data.

3. Tensorizing Recurrent Neural Networks

The core concept of this work is to approximate $\mathbf{W} \cdot \mathbf{x}$ with much fewer parameters, while still preserving the memorization mechanism in existing RNN formulations. The technique we use for the approximation is Block Term Decomposition (BTM), which represents $\mathbf{W} \cdot \mathbf{x}$ as a series of light-weighted small tensor products. In the process of RNN training, the BTM layer automatically learns inter-parameter correlations to implicitly prune redundant dense connections rendered by $\mathbf{W} \cdot \mathbf{x}$. By plugging the new BTM layer into current RNN formulations, we present a new BT-RNN model with several orders of magnitude fewer parameters while maintaining the representation power.

This section elaborates the details of the proposed methodology. It starts with exploring the background of tensor representations and BTM, before delving into the transformation of a regular RNN model to the BT-RNN; then we present the back propagation procedures for the BT-RNN; finally, we analyze the time and memory complexity of the BT-RNN compared with the regular one.

3.1. Preliminaries and Background

Tensor Representation We use the boldface Euler script letter, e.g., \mathcal{X} , to denote a tensor. A d -order tensor represents a d dimensional multiway array; thereby a vector and a matrix is a 1-order tensor and a 2-order tensor, respectively. An element in a d -order tensor is denoted as $\mathcal{X}_{i_1, \dots, i_d}$.

Tensor Product and Contraction Two tensors can perform product on a k th order if their k th dimension matches.

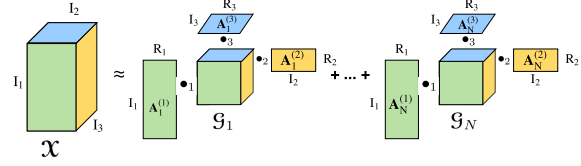


Figure 2: Block Term decomposition for a 3-order case tensor. A 3-order tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ can be approximated by N Tucker decompositions. We call the N the CP-rank, R_1, R_2, R_3 the Tucker-rank and d the Core-order.

Let's denote \bullet_k as the tensor-tensor product on k th order [17]. Given two d -order tensor $\mathcal{A} \in \mathbb{R}^{I_1 \times \dots \times I_d}$ and $\mathcal{B} \in \mathbb{R}^{J_1 \times \dots \times J_d}$, the tensor product on k th order is:

$$(\mathcal{A} \bullet_k \mathcal{B})_{i_k^-, i_k^+, j_k^-, j_k^+} = \sum_{p=1}^{I_k} \mathcal{A}_{i_k^-, p, i_k^+} \mathcal{B}_{j_k^-, p, j_k^+}. \quad (7)$$

To simplify, we use i_k^- denotes indices (i_1, \dots, i_{k-1}) , while i_k^+ denotes (i_{k+1}, \dots, i_d) . The whole indices can be denoted as $\vec{i} := (i_k^-, i_k, i_k^+)$. As we can see that each tensor product will be calculated along I_k dimension, which is consistent with matrix product.

Contraction is an extension of tensor product [6]; it conducts tensor products on multiple orders at the same time. For example, if $I_k = J_k, I_{k+1} = J_{k+1}$, we can conduct a tensor product according the k th and $(k+1)$ th order:

$$(\mathcal{A} \bullet_{k, k+1} \mathcal{B})_{i_k^-, i_{k+1}^+, j_k^-, j_{k+1}^+} = \sum_{p=1}^{I_k} \sum_{q=1}^{I_{k+1}} \mathcal{A}_{i_k^-, p, q, i_{k+1}^+} \mathcal{B}_{j_k^-, p, q, j_{k+1}^+}. \quad (8)$$

Block Term Decomposition (BTM) Block Term decomposition is a combination of CP decomposition [4] and Tucker decomposition [39]. Given a d -order tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times \dots \times I_d}$, BTM decomposes it into N block terms; And each term conducts \bullet_k between a core tensor $\mathcal{G}_n \in \mathbb{R}^{R_1 \times \dots \times R_d}$ and d factor matrices $\mathcal{A}_n^{(k)} \in \mathbb{R}^{I_k \times R_k}$ on \mathcal{G}_n 's k th dimension, where $n \in [1, N]$ and $k \in [1, d]$ [7]. The formulation of BTM is as follows:

$$\mathcal{X} = \sum_{n=1}^N \mathcal{G}_n \bullet_1 \mathcal{A}_n^{(1)} \bullet_2 \mathcal{A}_n^{(2)} \bullet_3 \dots \bullet_d \mathcal{A}_n^{(d)}. \quad (9)$$

We call the N the CP-rank, R_1, R_2, R_3 the Tucker-rank and d the Core-order. Fig. 2 demonstrates an example of how 3-order tensor \mathcal{X} being decomposed into N block terms.

3.2. BT-RNN model

This section demonstrates the core steps of BT-RNN model. 1) We transform \mathbf{W} and \mathbf{x} into tensor representations, \mathcal{W} and \mathcal{X} ; 2) then we decompose \mathcal{W} into several low-rank core tensors \mathcal{G}_n and their corresponding factor

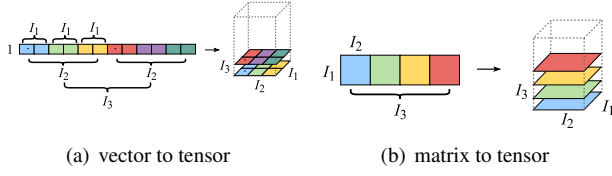


Figure 3: Tensorization operation in a case of 3-order tensors. (a) Tensorizing a vector with shape $I = I_1 \cdot I_2 \cdot I_3$ to a tensor with shape $I_1 \times I_2 \times I_3$; (b) Tensorizing a matrix with shape $I_1 \times (I_2 \cdot I_3)$ to a tensor with shape $I_1 \times I_2 \times I_3$.

tensors $\mathcal{A}_n^{(d)}$ using BTD; 3) subsequently, the original product $\mathbf{W} \cdot \mathbf{x}$ is approximated by the tensor contraction between decomposed weight tensor \mathcal{W} and input tensor \mathcal{X} ; 4) finally, we present the gradient calculations amid Back Propagation Through Time (BPTT) [43, 42] to demonstrate the learning procedures of BT-RNN model.

Tensorizing \mathbf{W} and \mathbf{x} we tensorize the input vector \mathbf{x} to a high-order tensor \mathcal{X} to capture spatial information of the input data, while we tensorize the weight matrix \mathbf{W} to decomposed weight tensor \mathcal{W} with BTD.

Formally, given an input vector $\mathbf{x}_t \in \mathbb{R}^I$, we define the notation φ to denote the *tensorization* operation. It can be either a *stack* operation or a *reshape* operation. We use reshape operation for tensorization as it does not need to duplicate the element of the data. Essentially reshaping is regrouping the data. Fig. 3 outlines how we reshape a vector and a matrix into 3-order tensors.

Decomposing \mathbf{W} with BTD Given a 2 dimensions weight matrix $\mathbf{W} \in \mathbb{R}^{J \times I}$, we can tensorize it as a $2d$ dimensions tensor $\mathcal{W} \in \mathbb{R}^{J_1 \times I_1 \times J_2 \times \dots \times J_d \times I_d}$, where $I = I_1 I_2 \dots I_d$ and $J = J_1 J_2 \dots J_d$. Following BTD in Eq. (9), we can decomposes \mathcal{W} into:

$$BTD(\mathcal{W}) = \sum_{n=1}^N \mathcal{G}_n \bullet_1 \mathcal{A}_n^{(1)} \bullet_2 \dots \bullet_d \mathcal{A}_n^{(d)}, \quad (10)$$

where $\mathcal{G}_n \in \mathbb{R}^{R_1 \times \dots \times R_d}$ denotes the core tensor, $\mathcal{A}_n^{(d)} \in \mathbb{R}^{I_d \times J_d \times R_d}$ denotes the factor tensor, N is the CP-rank and d is the Core-order. From the mathematical property of BT's ranks [17], we have $R_k \leq I_k$ (and J_k), $k = 1, \dots, d$. If $R_k > I_k$ (or J_k), it is difficult for the model to obtain bonus in performance. What's more, to obtain a robust model, in practice, we set each Tucker-rank to be equal, e.g., $R_i = R$, $i \in [1, d]$, to avoid unbalanced weight sharing in different dimensions and to alleviate the difficulty in hyper-parameters setting.

Computation between \mathbf{W} and \mathbf{x} After substituting the matrix-vector product by BT representation and tensorized input vector, we replace the input-to-hidden matrix-vector product $\mathbf{W} \cdot \mathbf{x}_t$ with the following form:

$$\phi(\mathbf{W}, \mathbf{x}_t) = BTD(\mathcal{W}) \bullet_{1,2,\dots,d} \mathcal{X}_t, \quad (11)$$

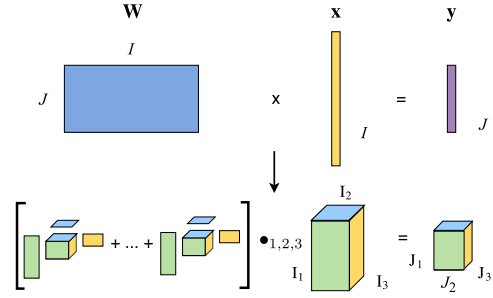


Figure 4: Diagrams of BT representation for matrix-vector product $\mathbf{y} = \mathbf{W}\mathbf{x}$, $\mathbf{W} \in \mathbb{R}^{J \times I}$. We substitute the weight matrix \mathbf{W} by the BT representation, then tensorize the input vector \mathbf{x} to a tensor with shape $I_1 \times I_2 \times I_3$. After operating the tensor contraction between BT representation and input tensor, we get the result tensor in shape $J_1 \times J_2 \times J_3$. With the reverse tensorize operation, we get the output vector $\mathbf{y} \in \mathbb{R}^{J_1 \cdot J_2 \cdot J_3}$.

where the tensor contraction operation $\bullet_{1,2,\dots,d}$ will be computed along all I_k dimensions in \mathcal{W} and \mathcal{X} , yielding the same size in the element-wise form as the original one. Fig. 4 demonstrates the substitution intuitively.

Training BT-RNN The gradient of RNN is computed by Back Propagation Through Time (BPTT) [43]. We derive the gradients amid the framework of BPTT for the proposed BT-RNN model.

Following the regular LSTM back-propagation procedure, the gradient $\frac{\partial L}{\partial \mathbf{y}}$ can be computed by the original BPTT algorithm, where $\mathbf{y} = \mathbf{W}\mathbf{x}_t$. Using the *tensorization* operation same to \mathbf{y} , we can obtain the tensorized gradient $\frac{\partial L}{\partial \mathcal{Y}}$. For a more intuitive understanding, we rewrite Eq. (11) in element-wise case:

$$\mathbf{y}_{\vec{j}} = \sum_{n=1}^N \sum_{\vec{r}}^{R_1, \dots, R_d} \sum_{\vec{i}}^{I_1, \dots, I_d} \prod_{k=1}^d \mathcal{X}_{t, \vec{i}} \mathcal{A}_{n, i_k, j_k, r_k}^{(k)} \mathcal{G}_{n, \vec{r}}. \quad (12)$$

Here, for simplified writing, we use \vec{i} , \vec{j} and \vec{r} to denote the indices (i_1, \dots, i_d) , (j_1, \dots, j_d) and (r_1, \dots, r_d) , respectively. Since the right hand side of Eq. (12) is a scalar, the element-wise gradient for parameters in BT-RNN is as follows:

$$\frac{\partial L}{\partial \mathcal{A}_{n, i_k, j_k, r_k}^{(k)}} = \sum_{\vec{r}' \neq \vec{r}_k} \sum_{\vec{i}' \neq \vec{i}_k} \sum_{\vec{j}' \neq \vec{j}_k} \prod_{k' \neq k} \mathcal{G}_{n, \vec{r}'} \mathcal{A}_{n, i_{k'}, j_{k'}, r_{k'}} \mathcal{X}_{t, \vec{i}'} \frac{\partial L}{\partial \mathbf{y}_{\vec{j}'}} \quad (13)$$

$$\frac{\partial L}{\partial \mathcal{G}_{n, \vec{r}}} = \sum_{\vec{i}} \sum_{\vec{j}} \prod_{k=1}^d \mathcal{A}_{n, i_k, j_k, r_k} \mathcal{X}_{t, \vec{i}} \frac{\partial L}{\partial \mathbf{y}_{\vec{j}}}. \quad (14)$$

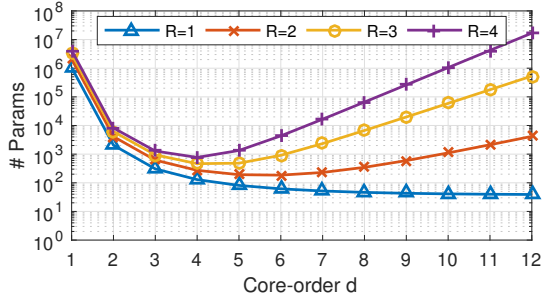


Figure 5: The number of parameters w.r.t Core-order d and Tucker-rank R , in the setting of $I = 4096, J = 256, N = 1$. While the vanilla RNN contains $I \times J = 1048576$ parameters. Refer to Eq. (15), when d is small, the first part $\sum_1^d I_k J_k R$ does the main contribution to parameters. While d is large, the second part R^d does. So we can see the number of parameters will go down sharply at first, but rise up gradually as d grows up (except for the case of $R = 1$).

3.3. Hyper-Parameters and Complexity Analysis

3.3.1 Hyper-Parameters Analysis

Total #Params BTD decomposes \mathcal{W} into N block terms and each block term is a tucker representation [18, 17], therefore the total amount of parameters is as follows:

$$P_{BTD} = N \left(\sum_{k=1}^d I_k J_k R + R^d \right). \quad (15)$$

By comparison, the original weight matrix \mathbf{W} contains $P_{RNN} = I \times J = \prod_{k=1}^d I_k J_k$ parameters, which is several orders of magnitude larger than it in the BTD representation.

#Params w.r.t Core-order (d) Core-order d is the most significant factor affecting the total amount of parameters as R^d term in Eq. (15). It determines the total dimensions of core tensors, the number of factor tensors, and the total dimensions of input and output tensors. If we set $d = 1$, the model degenerates to the original matrix-vector product with the largest number of parameters and the highest complexity. Fig. 5 demonstrates how total amount of parameters vary w.r.t different Core-order d . If the Tucker-rank $R > 1$, the total amount of parameters first decreases with d increasing until reaches the minimum, then starts increasing afterwards. This mainly results from the non-linear characteristic of d in Eq. (15).

Hence, a proper choice of d is particularly important. Enlarging the parameter d is the simplest way to reduce the number of parameters. But due to the second term R^d in Eq. (15), enlarging d will also increase the amount of parameters in the core tensors, resulting in the high computational complexity and memory usage. With the Core-order d increasing, each dimensions of the input tensor decreases

logarithmically. However, this will result in the loss of important spatial information in an extremely high order BT model. In practice, Core-order $d \in [2, 5]$ is recommended.

#Params w.r.t Tucker-rank (R) The Tucker-rank R controls the complexity of Tucker decomposition. This hyper-parameter is conceptually similar to the number of singular values in Singular Value Decomposition (SVD). Eq. (15) and Fig. 5 also suggest the total amount of parameters is sensitive to R . Particularly, BT degenerates to a CP decomposition if we set it as $R = 1$. Since $R \leq I_k$ (and J_k), the choice of R is limited in a small value range, releasing researchers from heavily hyper-parameters setting.

#Params w.r.t CP-rank (N) The CP-rank N controls the number of block terms. If $N = 1$, BT degenerates to a Tucker decomposition. As we can see from Table 1 that N does not affect the memory usage in forward and backward passes, so if we need a more memory saving model, we can enlarge N while decreasing R and d at the same time.

3.3.2 Computational Complexity Analysis

Complexity in Forward Process Eq. (10) raises the computation peak, $\mathcal{O}(IJR)$, at the last tensor product \bullet_d , according to left-to-right computational order. However, we can reorder the computations to further reduce the total model complexity $\mathcal{O}(NdIJR)$. The reordering is:

$$\phi(\mathbf{W}, \mathbf{x}_t) = \sum_{n=1}^N \mathbf{x}_t \bullet_1 \mathcal{A}_n^{(1)} \bullet_2 \cdots \bullet_d \mathcal{A}_n^{(d)} \bullet_{1,2,\dots,d} \mathcal{G}_n. \quad (16)$$

The main difference is each tensor product will be first computed along all R dimensions in Eq. (11), while in Eq. (16) along all I_k dimensions. Since BT is a low-rank decomposition method, e.g., $R \leq J_k$ and $J_{max} R^{(d-1)} \leq J$, the new computation order can significantly reduce the complexity of the last tensor product from $\mathcal{O}(IJR)$ to $\mathcal{O}(IJ_{max} R^d)$, where $J = J_1 J_2 \cdots J_d$, $J_{max} = \max_k(J_k)$, $k \in [1, d]$. And then the total complexity of our model reduces from $\mathcal{O}(NdIJR)$ to $\mathcal{O}(NdIJ_{max} R^d)$. If we decrease Tucker-rank R , the computation complexity decreases logarithmically in Eq. (16) while linearly in Eq. (11).

Complexity in Backward Process To derive the computational complexity in the backward process, we present gradients in the tensor product form. The gradients of factor tensors and core tensors are:

$$\frac{\partial L}{\partial \mathcal{A}_n^{(k)}} = \frac{\partial L}{\partial \mathbf{y}} \bullet_k \mathbf{x}_t \bullet_1 \mathcal{A}_n^{(1)} \bullet_2 \cdots \quad (17)$$

$$\bullet_{k-1} \mathcal{A}_n^{(k-1)} \bullet_{k+1} \mathcal{A}_n^{(k+1)} \bullet_{k+2} \cdots$$

$$\bullet_d \mathcal{A}_n^{(d)} \bullet_{1,2,\dots,d} \mathcal{G}_n$$

$$\frac{\partial L}{\partial \mathcal{G}_n} = \mathbf{x}_t \bullet_1 \mathcal{A}_n^{(1)} \bullet_2 \cdots \bullet_d \mathcal{A}_n^{(d)} \bullet_{1,2,\dots,d} \frac{\partial L}{\partial \mathbf{y}} \quad (18)$$

Method	Time	Memory
RNN forward	$\mathcal{O}(IJ)$	$\mathcal{O}(IJ)$
RNN backward	$\mathcal{O}(IJ)$	$\mathcal{O}(IJ)$
TT-RNN forward	$\mathcal{O}(dIR^2J_{max})$	$\mathcal{O}(RI)$
TT-RNN backward	$\mathcal{O}(d^2IR^4J_{max})$	$\mathcal{O}(R^3I)$
BT-RNN forward	$\mathcal{O}(NdIR^dJ_{max})$	$\mathcal{O}(R^dI)$
BT-RNN backward	$\mathcal{O}(Nd^2IR^dJ_{max})$	$\mathcal{O}(R^dI)$

Table 1: Comparison of complexity and memory usage of vanilla RNN, Tensor-Train representation RNN (TT-RNN) [28, 47] and our BT representation RNN (BT-RNN). In this table, the weight matrix’s shape is $I \times J$. The input and hidden tensors’ shapes are $I = I_1 \times \dots \times I_d$ and $J = J_1 \times \dots \times J_d$, respectively. Here, $J_{max} = \max_k(J_k), k \in [1, d]$. Both TT-RNN and BT-RNN are set in same rank R .

Since Eq. (17) and Eq. (18) follow the same form of Eq. (11), the backward computational complexity is same as the forward pass $\mathcal{O}(dIJ_{max}R^d)$. Therefore, the $N \cdot d$ factor tensors demonstrate a total complexity of $\mathcal{O}(Nd^2IJ_{max}R^d)$.

Complexity Comparisons We analyze the time complexity and memory usage of RNN, Tensor Train RNN, and BT-RNN. The statistics are shown in Table 1. In our observation, both TT-RNN and BT-RNN hold lower computation complexity and memory usage than the vanilla RNN, since the extra hyper-parameters are several orders smaller than I or J . As we claim that the suggested choice of Core-order is $d \in [2, 5]$, the complexity of TT-RNN and BT-RNN should be comparable.

4. Experiments

RNN is a versatile and powerful modeling tool widely used in various computer vision tasks. We design three challenging computer vision tasks-Action Recognition in Videos, Image Caption and Image Generation-to quantitatively and qualitatively evaluate proposed BT-LSTM against baseline LSTM and other low-rank variants such as Tensor Train LSTM (TT-LSTM). Finally, we design a control experiment to elucidate the effects of different hyper-parameters.

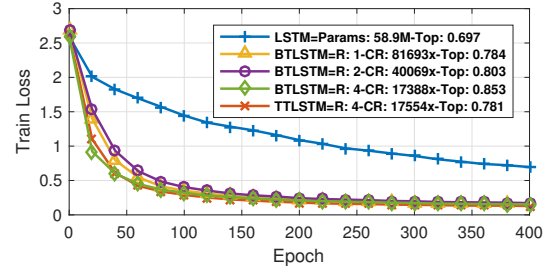
4.1. Implementations

Since operations in $\mathbf{f}_t, \mathbf{i}_t, \mathbf{o}_t$ and $\tilde{\mathbf{c}}_t$ follow the same computation pattern, we merge them together by concatenating $\mathbf{W}_f, \mathbf{W}_i, \mathbf{W}_o$ and \mathbf{W}_c into one giant \mathbf{W} , and so does \mathbf{U} . This observation leads to the following simplified LSTM formulations:

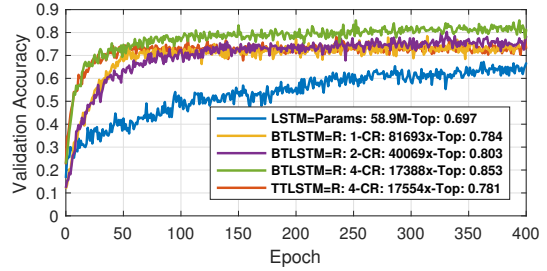
$$(\mathbf{f}'_t, \mathbf{i}'_t, \tilde{\mathbf{c}}'_t, \mathbf{o}'_t) = \mathbf{W} \cdot \mathbf{x}_t + \mathbf{U} \cdot \mathbf{h}_{t-1} + \mathbf{b}, \quad (19)$$

$$(\mathbf{f}_t, \mathbf{i}_t, \tilde{\mathbf{c}}_t, \mathbf{o}_t) = (\sigma(\mathbf{f}'_t), \sigma(\mathbf{i}'_t), \tanh(\tilde{\mathbf{c}}'_t), \sigma(\mathbf{o}'_t)). \quad (20)$$

We implemented BT-LSTM on the top of simplified LSTM formulation with Keras and TensorFlow. The initialization of baseline LSTM models use the default settings in Keras



(a) Training loss of baseline LSTM, TT-LSTM and BT-LSTM.



(b) Validation Accuracy of baseline LSTM, TT-LSTM and BT-LSTM.

Figure 6: Performance of different RNN models on the Action Recognition task trained with UCF11. CR stands for Compression Ratio; R is Tucker-rank, and Top is the highest validation accuracy observed in the training. Though BT-LSTM utilizes 17388 times less parameters than the vanilla LSTM (58.9 Millions), BT-LSTM demonstrates a 15.6% higher accuracy improvement than LSTM. BT-LSTM also demonstrates an extra 7.2% improvement over the TT-LSTM with comparable parameters.

	Method	Accuracy
Orthogonal Approaches	Original [25]	0.712
	Spatial-temporal [24]	0.761
	Visual Attention [32]	0.850
RNN Approaches	LSTM	0.697
	TT-LSTM [47]	0.796
	BT-LSTM	0.853

Table 2: State-of-the-art results on UCF11 dataset reported in literature, in comparison with our best model.

and TensorFlow, while we use Adam optimizer with the same learning rate (lr) across different tasks.

4.2. Quantitative Evaluations of BT-LSTM on the Task of Action Recognition in Videos

We use UCF11 YouTube Action dataset [25] for action recognition in videos. The dataset contains 1600 video clips, falling into 11 action categories. Each category contains 25 video groups, within each contains at least 4 clips. All video clips are converted to 29.97fps MPG². We scale down original frames from $320 \times 240 \times 3$ to $160 \times 120 \times 3$,

²http://crcv.ucf.edu/data/UCF11_updated_mpg.rar

then we sample 6 random frames in ascending order from each video clip as the input data. For more details on the preprocessing, please refer to [47].

We use a single LSTM cell as the model architecture to evaluate BT-LSTM against LSTM and TT-LSTM in Fig. 6. Please note there are other orthogonal approaches aiming at improving the model such as visual attention [32] and spatial-temporal [24]. Since our discussion is limited to a single LSTM cell, we can always replace the LSTM cells in those high-level models with BT-LSTM to acquire better accuracies. We set the hyper-parameters of BT-LSTM and TT-LSTM as follows: the factor tensor counts is $d = 4$; the shape of input tensor is $I_1 = 8, I_2 = 20, I_3 = 20, I_4 = 18$; and the hidden shape is $J_1 = J_2 = J_3 = J_4 = 4$; the rank of TT-LSTM is $R_1 = R_5 = 1, R_2 = R_3 = R_4 = 4$, while BT-LSTM is set to various Tucker-ranks.

Fig. 6 demonstrates the training loss and validation accuracy of BT-LSTM against LSTM and TT-LSTM under different settings. Table 2 demonstrates the top accuracies of different models. From these experiments, we claim that:

1) **8×10^4 times parameter reductions**: The vanilla LSTM has 58.9 millions parameters in \mathbf{W} , while BT-LSTM delivers better accuracies even with several orders of less parameters. The total parameters in BT-LSTM follows Eq. (15). At Tucker-rank 1, 2, 4, BT-LSTM uses 721, 1470, and 3387 parameters, demonstrating compression ratios of 81693x, 40069x and 17388x, respectively.

2) **faster convergence**: BT-LSTM demonstrates significant convergence improvement over the vanilla LSTM based on training losses and validation accuracies in Fig. 6(a) and Fig. 6(b). In terms of validation accuracies, BT-LSTM reaches 60% accuracies at epoch-16 while LSTM takes 230 epochs. The data demonstrates 14x convergence speedup. It is widely acknowledged that the model with few parameters is easier to train. Therefore, the convergence speedup majorly results from the drastic parameter reductions. At nearly same parameters, the training loss of BT-LSTM-4 also decreases faster than TTLSTM-4 (epochs[0, 50]), substantiating that BT model captures better spatial information than the Tensor Train model.

3) **better model efficiency**: Though several orders of parameter reductions, BT-LSTM demonstrates extra 15.6% accuracies than LSTM. In addition, BT-LSTM also demonstrates extra 7.2% accuracies than TT-LSTM with comparable parameters. In different Tucker-ranks, BT-LSTM converges to identical losses; but increasing Tucker ranks also improves the accuracy. This is consistent with the intuition since the high rank models capture additional relevant information.

4.3. Qualitative Evaluations of BT-LSTM on Tasks of Image Generation and Image Captioning

We also conduct experiments on Image Generation and Image Captioning to further substantiate the efficiency of BT-LSTM.

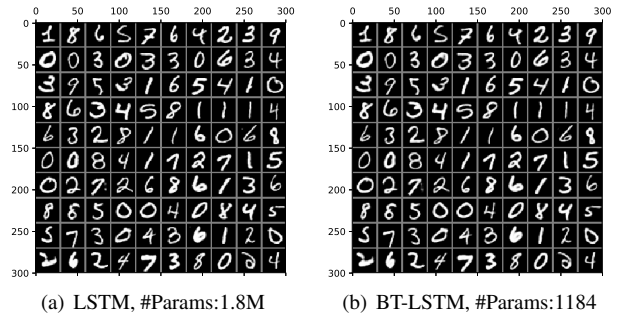


Figure 7: Image Generation: generating MNIST style digits with LSTM and BT-LSTM based model. The results are merely identical, while the parameters of BT-LSTM is 1577 times less.

Task 1: Image Generation Image generation intends to learn latent representation from images, then it tries to generate new image of same style from the learned model. The model for this task is Deep Recurrent Attentive Writer (DRAW) [9]. It uses an encoder RNN network to encode images into latent representations; then an decoder RNN network decodes the latent representations to construct an image. we substitute LSTM in encoder network with our BT-LSTM.

In this task, encoder network must capture sufficient correlations and visual features from raw images to generate high quality of feature vectors. As shown in Fig. 7, both LSTM and BT-LSTM model generate comparable images.

Task 2: Image Captioning Image Captioning intends to describe the content of an image. We use the model in Neural Image Caption[40] to evaluate the performance of BT-LSTM by replacing the LSTM cells.

The training dataset is MSCOCO [23], a large-scale dataset for the object detection, segmentation, and captioning. Each image is scaled to 224×224 in RGB channels and subtract the channel means as the input to a pretrained Inception-v3 model.

Fig. 8 demonstrates the image captions generated by BT-LSTM and LSTM. It is obvious that both BT-LSTM, TT-LSTM and LSTM can generate proper sentences to describe the content of an image, but with little improvement in BT-LSTM. Since the input data of BT model is a compact feature vector merged with the embedding images features from Inception-v3 and language features from a word embedding network, our model demonstrates the qualitative improvement in captioning. The results also demonstrate that BT-LSTM captures local correlations missed by traditional LSTM.

4.4. Sensitivity Analysis on Hyper-Parameters

There are 3 key hyper-parameters in BT-LSTM, which are core-order d , Tucker-rank R and CP-rank N . In order to scrutinize the impacts of these hyper-parameters, we design a control experiment illustrate their effects.



(a) **LSTM**: A train traveling down tracks next to a forest.
TT-LSTM: A train traveling down train tracks next to a forest.
BT-LSTM: A train traveling through a lush green forest.



(b) **LSTM**: A group of people standing next to each other.
TT-LSTM: A group of men standing next to each other.
BT-LSTM: A group of people posing for a photo.



(c) **LSTM**: A man and a dog are standing next to each other.
TT-LSTM: A man and a dog are in the snow.
BT-LSTM: A man and a dog playing with a frisbee.



(d) **LSTM**: A large elephant standing next to a baby elephant.
TT-LSTM: An elephant walking down a dirt road near trees.
BT-LSTM: A large elephant walking down a road with cars.

Figure 8: Results of image caption in MSCOCO dataset.

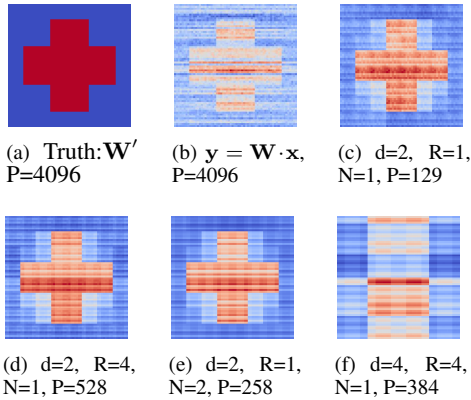


Figure 9: The trained W for different BT-LSTM settings. The closer to (a), the better W is.

We try to sample y from the distribution of $y = W' \cdot x$, where $x, y \in \mathbb{R}^{64}$. Each x is generated from a Gaussian distribution $\mathcal{N}(0, 0.5)$. We also add a small noise into x to avoid overfitting. y is generated by plugging x back to $y = W' \cdot x$. Given x and y , we randomly initialize W , and start training. Eventually, W should be similar to W' since x and y drawn from the distribution of $y = W' \cdot x$. Please note that the purpose of this experiment is to evaluate the impact of the BT model on different parameter settings, despite these are many other good methods such as L1 regularization and Lasso regularization, to recover the W weight matrix.

Core-order (d): Parameters goes down if d grows and $d < 5$. Parameters reduce about 1.3 times from Fig. 9(d) to Fig. 9(f); and d increase from 2 to 4. With less parameters, the reconstructed W deteriorates quickly. We claim that high Core-order d loses important spatial information, as tensor becomes too small to capture enough latent correlations. This result is consistent with our declaration.

Tucker-rank (R): the rank R take effectiveness exponentially to the parameters. By comparing Fig. 9(c) and Fig. 9(d), When R increases from 1 to 4, BT model has more parameters to capture sufficient information from input data, obtaining a more robust model.

CP-rank (N): CP-rank contributes to the number of parameters linearly, playing an important role when R is small. By comparing Fig. 9(c) and Fig. 9(e), we can see that the latter result has less noise in figure, showing that a proper CP-rank setting will lead to a more robust model, since we use multiple Tucker models to capture information from input data.

5. Conclusion

We proposed a Block-Term RNN architecture to address the redundancy problem in RNNs. By using a Block Term tensor decomposition to prune connections in the input-to-hidden weight matrix of RNNs, we provide a new RNN model with a less number of parameters and stronger correlation modeling between feature dimensions, leading to easy model training and improved performance. Experiment results on a video action recognition data set show that our BT-RNN architecture can not only consume several orders fewer parameters but also improve the model performance over standard traditional LSTM and the TT-LSTM. The next works are to 1) explore the sparsity in factor tensors and core tensors of BT model, further reducing the number of model parameters; 2) concatenate hidden states and input data for a period of time, respectively, extracting the temporal features via tensor methods; 3) quantify factor tensors and core tensors to reduce memory usage.

Acknowledgment

This paper was in part supported by a grant from the Natural Science Foundation of China (No.61572111), 1000-Talent Program Startup Funding (A1098531023601041,G05QNQR004) and a Fundamental Research Fund for the Central Universities of China (No. A03017023701). Zenglin Xu is the major corresponding author.

References

- [1] M. Bai, B. Zhang, and J. Gao. Tensorial recurrent neural networks for longitudinal data analysis. *arXiv preprint arXiv:1708.00185*, 2017.

- [2] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
- [3] W. Byeon, T. M. Breuel, F. Raue, and M. Liwicki. Scene labeling with lstm recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3547–3555, 2015.
- [4] J. D. Carroll and J.-J. Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of eckart-young decomposition. *Psychometrika*, 35(3):283–319, 1970.
- [5] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [6] A. Cichocki. Era of big data processing: A new approach via tensor networks and tensor decompositions. *arXiv preprint arXiv:1403.2048*, 2014.
- [7] L. De Lathauwer. Decompositions of a higher-order tensor in block terms part ii: Definitions and uniqueness. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1033–1066, 2008.
- [8] M. Denil, B. Shakibi, L. Dinh, N. de Freitas, et al. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pages 2148–2156, 2013.
- [9] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.
- [10] S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, pages 1135–1143, 2015.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [12] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [13] S. Hochreiter. Untersuchungen zu dynamischen neuronalen netzen. *Diploma, Technische Universität München*, 91, 1991.
- [14] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [15] G. Huang, Z. Liu, K. Q. Weinberger, and L. van der Maaten. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2016.
- [16] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014.
- [17] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- [18] J. Kossaifi, Z. C. Lipton, A. Khanna, T. Furlanello, and A. Anandkumar. Tensor regression networks. *arXiv preprint arXiv:1707.08308*, 2017.
- [19] V. Lebedev, Y. Ganin, M. Rakhuba, I. Oseledets, and V. Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *arXiv preprint arXiv:1412.6553*, 2014.
- [20] Y. LeCun, Y. Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.
- [21] G. Li, J. Ye, H. Yang, D. Chen, S. Yan, and Z. Xu. Bt-nets: Simplifying deep neural networks via block term decomposition. *CoRR*, abs/1712.05689, 2017.
- [22] X. Liang, X. Shen, D. Xiang, J. Feng, L. Lin, and S. Yan. Semantic object parsing with local-global long short-term memory. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3185–3193, 2016.
- [23] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [24] D. Liu, M.-L. Shyu, and G. Zhao. Spatial-temporal motion information integration for action detection and recognition in non-static background. In *Information Reuse and Integration (IRI), 2013 IEEE 14th International Conference on*, pages 626–633. IEEE, 2013.
- [25] J. Liu, J. Luo, and M. Shah. Recognizing realistic actions from videos in the wild. In *Computer vision and pattern recognition, 2009. CVPR 2009. IEEE conference on*, pages 1996–2003. IEEE, 2009.
- [26] Z. Lu, V. Sindhvani, and T. N. Sainath. Learning compact recurrent neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 5960–5964. IEEE, 2016.
- [27] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz. Pruning convolutional neural networks for resource efficient inference. 2016.
- [28] A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov. Tensorizing neural networks. In *Advances in Neural Information Processing Systems*, pages 442–450, 2015.
- [29] I. V. Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.
- [30] R. Prabhavalkar, O. Alsharif, A. Bruguier, and L. McGraw. On the compression of recurrent neural networks with an application to lvsr acoustic modeling for embedded speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 5970–5974. IEEE, 2016.
- [31] T. N. Sainath, B. Kingsbury, V. Sindhvani, E. Arisoy, and B. Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 6655–6659. IEEE, 2013.
- [32] S. Sharma, R. Kiros, and R. Salakhutdinov. Action recognition using visual attention. *arXiv preprint arXiv:1511.04119*, 2015.
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [34] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–1958, 2014.
- [35] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, pages 4278–4284, 2017.
- [36] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the*

- IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [37] L. Theis and M. Bethge. Generative image modeling using spatial lstms. In *Advances in Neural Information Processing Systems*, pages 1927–1935, 2015.
 - [38] A. Tjandra, S. Sakti, and S. Nakamura. Compressing recurrent neural network with tensor train. *arXiv preprint arXiv:1705.08052*, 2017.
 - [39] L. R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966.
 - [40] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015.
 - [41] L. Wang, W. Wu, Z. Xu, J. Xiao, and Y. Yang. Blasx: A high performance level-3 blas library for heterogeneous multi-gpu computing. In *Proceedings of the 2016 International Conference on Supercomputing*, page 20. ACM, 2016.
 - [42] L. Wang, Y. Yang, R. Min, and S. Chakradhar. Accelerating deep neural network training with inconsistent stochastic gradient descent. *Neural Networks*, 2017.
 - [43] P. J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
 - [44] Z. Xu, F. Yan, and Y. A. Qi. Infinite tucker decomposition: Nonparametric bayesian models for multiway data analysis. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*, 2012.
 - [45] Z. Xu, F. Yan, and Y. A. Qi. Bayesian nonparametric models for multiway data analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(2):475–487, 2015.
 - [46] J. Xue, J. Li, and Y. Gong. Restructuring of deep neural network acoustic models with singular value decomposition. In *Interspeech*, pages 2365–2369, 2013.
 - [47] Y. Yang, D. Krompass, and V. Tresp. Tensor-train recurrent neural networks for video classification. *arXiv preprint arXiv:1707.01786*, 2017.
 - [48] R. Yu, S. Zheng, A. Anandkumar, and Y. Yue. Long-term forecasting using tensor-train rnns. *arXiv preprint arXiv:1711.00073*, 2017.
 - [49] C. Yunpeng, J. Xiaojie, K. Bingyi, F. Jiashi, and Y. Shuicheng. Sharing residual units through collective tensor factorization in deep neural networks. *arXiv preprint arXiv:1703.02180*, 2017.
 - [50] Q. Zhao, M. Sugiyama, and A. Cichocki. Learning efficient tensor representations with ring structure networks. *arXiv preprint arXiv:1705.08286*, 2017.