

Tensor Contraction Layers for Parsimonious Deep Nets

Jean Kossaifi
Amazon AI

Imperial College London

jean.kossaifi@imperial.ac.uk

Aran Khanna
Amazon AI

arankhan@amazon.com

Zachary Lipton
Amazon AI

University of California, San Diego

zlipton@cs.ucsd.edu

Tommaso Furlanello
Amazon AI

University of Southern California

furlanel@usc.edu

Anima Anandkumar
Amazon AI

California Institute of Technology

anima@amazon.com

Abstract

Tensors offer a natural representation for many kinds of data frequently encountered in machine learning. Images, for example, are naturally represented as third order tensors, where the modes correspond to height, width, and channels. In particular, tensor decompositions are noted for their ability to discover multi-dimensional dependencies and produce compact low-rank approximations of data. In this paper, we explore the use of tensor contractions as neural network layers and investigate several ways to apply them to activation tensors. Specifically, we propose the Tensor Contraction Layer (TCL), the first attempt to incorporate tensor contractions as end-to-end trainable neural network layers. Applied to existing networks, TCLs reduce the dimensionality of the activation tensors and thus the number of model parameters. We evaluate the TCL on the task of image recognition, augmenting popular networks (AlexNet, VGG). The resulting models are trainable end-to-end. We evaluate TCL's performance on the task of image recognition, using the CIFAR100 and ImageNet datasets, studying the effect of parameter reduction via tensor contraction on performance. We demonstrate significant model compression without significant impact on the accuracy and, in some cases, improved performance.

1. Introduction

Following their successful application to computer vision, speech recognition, and natural language processing, deep neural networks have become ubiquitous in the machine learning community. And yet many questions remain unanswered: Why do deep neural networks work? How many parameters are really necessary to achieve state of the art performance?

Recently, tensor methods have been used in attempts to better understand the success of deep neural networks [4, 6]. One class of broadly useful techniques within tensor methods are tensor decompositions. While the properties of tensors have long been studied, in the past decade they have come to prominence in machine learning in such varied applications as learning latent variable models [1], and developing recommender systems [10]. Several recent papers apply tensor learning and tensor decomposition to deep neural networks for the purpose of devising neural network learning algorithms with theoretical guarantees of convergence [17, 9].

Other lines of research have investigated practical applications of tensor decomposition to deep neural networks with aims including speeding up convolutional neural networks [15], multi-task learning [20], and sharing residual units [3]. However, to our knowledge, no attempt has been made to apply tensor decompositions as a generic layer directly on the activations or weights of a deep neural network and to train the resulting network end-to-end. Most current work applies decompositions only for either initialization [20] or post-training [16]. These techniques then often require additional fine-tuning to compensate for the loss of information [11].

In deep convolutional neural networks, the output of each layer is a tensor. We posit that tensor algebraic techniques can allow us to exploit multidimensional dependencies in the activation tensors. We propose to leverage that structure by incorporating Tensor Contraction Layers (TCLs) into neural networks. Specifically, in our experiments, we apply TCLs directly to the third-order activation tensors produced by the final convolutional layer of an image recognition network. Traditional networks flatten this activation tensor, passing it to subsequent fully-connected layers. However, the flattening process loses information

about the multidimensional structure of the tensor. Our experiments show that incorporating TCLs into several popular deep convolutional networks can improve their performance, despite reducing the number of parameters. Moreover, inference on TCL-equipped networks, which contain less parameters, requires considerably fewer floating point operations.

We organize the rest of this paper as follows: Section 1.1 introduces prerequisite concepts needed to understand the TCL; Section 2 explains the TCL in detail; Section 3 experimentally evaluates the TCL.

1.1. Tensor Contraction

Notation: We define tensors as multidimensional arrays, denoting first-order tensors \mathbf{v} as *vectors*, second-order tensors \mathbf{M} as *matrices* and by $\tilde{\mathcal{X}}$, refer to tensors of order 3 or greater. \mathbf{M}^\top denotes the transpose of \mathbf{M} .

Tensor unfolding: Given a tensor, $\tilde{\mathcal{X}} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, the mode- n unfolding of $\tilde{\mathcal{X}}$ is a matrix $\mathbf{X}_{[n]} \in \mathbb{R}^{I_n \times I_M}$, with $M = \prod_{k=1, k \neq n}^N I_k$ and is defined by the mapping from element (i_1, i_2, \dots, i_N) to (i_n, j) , with $j = \sum_{k=1, k \neq n}^N i_k \times \prod_{m=k+1}^N I_m$.

n-mode product: For a tensor $\tilde{\mathcal{X}} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ and a matrix $\mathbf{M} \in \mathbb{R}^{R \times I_n}$, the n -mode product of $\tilde{\mathcal{X}}$ by \mathbf{M} is a tensor of size $(I_1 \times \dots \times I_{n-1} \times R \times I_{n+1} \times \dots \times I_N)$ and can be expressed using the unfolding of $\tilde{\mathcal{X}}$ and the classical matrix multiplication as:

$$\tilde{\mathcal{X}} \times_n \mathbf{M} = \mathbf{M} \tilde{\mathbf{X}}_{[n]} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times R \times I_{n+1} \times \dots \times I_N} \quad (1)$$

Tensor contraction: Given a tensor $\tilde{\mathcal{X}} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, we can decompose it into a low-dimensional core tensor $\tilde{\mathcal{G}} \in \mathbb{R}^{R_1 \times R_2 \times \dots \times R_N}$ through projection along each of its modes by projection factors $(\mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)})$, with $\mathbf{U}^{(k)} \in \mathbb{R}^{R_k \times I_k}, k \in (1, \dots, N)$. In other words, we can write:

$$\tilde{\mathcal{G}} = \tilde{\mathcal{X}} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times \dots \times_N \mathbf{U}^{(N)} \quad (2)$$

or, in short:

$$\tilde{\mathcal{G}} = \llbracket \tilde{\mathcal{X}}; \mathbf{U}^{(1)}, \dots, \mathbf{U}^{(N)} \rrbracket \quad (3)$$

Typically, the factors of the contraction are obtained by solving a least squares problem. In particular, closed form solutions can be obtained for the factor by considering the n -mode unfolding of $\tilde{\mathcal{X}}$ that can be expressed as:

$$\mathbf{G}_{[n]} = \mathbf{U}^{(n)} \mathbf{X}_{[n]} \left(\mathbf{U}^{(1)} \otimes \dots \otimes \mathbf{U}^{(n-1)} \otimes \mathbf{U}^{(n+1)} \otimes \dots \otimes \mathbf{U}^{(N)} \right)^T \quad (4)$$

We refer the interested reader to the seminal work of Kolda and Bader [12].




Figure 1. A representation of the Tensor Contraction Layer (TCL) applied on a tensor of order 3. The input tensor $\tilde{\mathcal{X}}$ is contracted into a low-dimensionality core $\tilde{\mathcal{G}}$.

1.2. Networks with Large fully connected layers

Many popular convolutional neural networks for computer vision, e.g. AlexNet, ResNet, and Inception, require hundreds of millions of parameters to achieve the reported results. This can be problematic when running these networks for inference on resource-constrained devices, where it may not be easy to run hundreds of millions of calculations just to classify a single image.

While these widely used architectures exhibit considerable variety, they also exhibit some commonalities. Often, they consist of blocks containing convolution, activation and pooling layers followed by fully-connected layers before the final classification layer. Both the popular networks AlexNet [14] and VGG [19] follow this meta-architecture, with both containing two fully-connected layers of 4096 hidden units each. In both networks, these fully connected layers hold over 80 percent of the parameters. In VGG, the hidden units contain 119,545,856 of the 138,357,544 total parameters, and in AlexNet the hidden units contain 54,534,144 out of the 62,378,344 total parameters.

Given the enormous computational costs for both training and running inference in these networks, we desire techniques that preserve high accuracy while reducing the number of parameters in the network. Notable work in this direction includes approaches to induce and exploit sparsity in the parameters during training [7].

2. Tensor Contraction Layer

In this paper, we propose to incorporate the tensor contraction into convolutional neural networks as an end-to-end trainable layer, applying it to the third order activation tensor output by the final convolutional layer.

In particular, given an activation tensor $\tilde{\mathcal{X}}$ of size (D_1, \dots, D_N) , we seek a low dimensional core $\tilde{\mathcal{G}}$ of smaller size (R_1, \dots, R_N) such that:

$$\tilde{\mathcal{G}} = \tilde{\mathcal{X}} \times_1 \mathbf{V}^{(1)} \times_2 \mathbf{V}^{(2)} \times \dots \times_N \mathbf{V}^{(N)} \quad (5)$$

with $\mathbf{V}^{(k)} \in \mathbb{R}^{R_k \times I_k}, k \in (1, \dots, N)$.

We leverage this formulation and define a new layer that takes the activation tensor $\tilde{\mathcal{X}}$ obtained from a previous layer

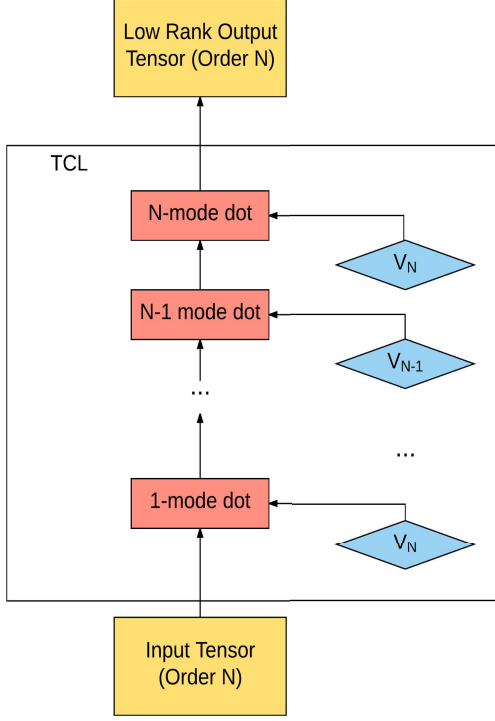


Figure 2. A representation of the symbolic graph of the Tensor Contraction Layer.

and applies such a projection to it (Figure. 1). We optimize the projection factors $(\mathbf{V}^{(k)})_{k \in [1, \dots, N]}$ to obtain a low dimensional projection of the activation tensor as the output of the layer. We learn the projection factors by backpropagation jointly with the rest of the network’s parameters. We coin this new layer the *Tensor Contraction Layer (TCL)* and denote by $size-(R_1, \dots, R_N)$ *TCL*, or *TCL*– (R_1, \dots, R_N) a tensor contraction layer producing a contracted output of size (R_1, \dots, R_N) .

The gradients with respect to each of the factors can be derived easily from 4. Specifically, for each $k \in 1, \dots, N$, we use the following equivalences:

$$\begin{aligned} \frac{\partial \tilde{g}}{\partial \mathbf{V}^{(k)}} &= \frac{\partial \tilde{\mathcal{X}} \times_1 \mathbf{V}^{(1)} \times_2 \mathbf{V}^{(2)} \times \dots \times_N \mathbf{V}^{(N)}}{\partial \mathbf{V}^{(k)}} = \frac{\partial \tilde{g}_{[k]}}{\partial \mathbf{V}^{(k)}} \\ &= \frac{\partial \mathbf{V}^{(k)} \mathbf{x}_{[k]} \left(\mathbf{V}^{(1)} \otimes \dots \otimes \mathbf{V}^{(k-1)} \otimes \mathbf{V}^{(k+1)} \otimes \dots \otimes \mathbf{V}^{(N)} \right)^T}{\partial \mathbf{V}^{(k)}} \end{aligned}$$

In practice, with minibatch training, we might think of the first mode of an activation tensor as corresponding to the batch-size. Technically, it is possible to applying a transformation along this dimension too, but we leave this consideration for future work. It is trivial to address this case by either starting the n –mode products at the second mode or by setting the first factor to be the Identity and not optimize

over it. Therefore, in the remainder of the paper, we consider the activation tensor for a single sample for clarity, without loss of generality.

Figure. 2 presents the symbolic graph of the tensor contraction layer. Note that when taking the n –mode product over different modes, the order in which the n –mode products are computed does not matter.

2.1. Complexity of the TCL

In this section we detail the number of parameters and complexity of the tensor contraction layer.

Number of parameters Let $\tilde{\mathcal{X}}$ be an activation tensor of size (D_1, \dots, D_N) which we pass through a $size-(R_1, \dots, R_N)$ Tensor Contraction Layer.

This TCL has a total of $\sum_{k=1}^N D_k \times R_k$ parameters (corresponding to the factors of the N n –mode products) and produces as input a tensor of size (R_1, \dots, R_N) .

By comparison, a fully connected layer producing an output of the same size, i.e. with $H = \prod_{k=1}^N R_k$ hidden units, and taking the same (flattened) tensor as input would have a total of $\prod_{k=1}^N D_k \times \prod_{k=1}^N R_k$ parameters.

Complexity As previously exposed, one way to look at the TCL is as a series of matrix multiplications between the factors of the contraction and the unfolded activation tensor. Let’s place ourselves in the setting previously detailed with an activation tensor $\tilde{\mathcal{X}}$ of size (D_1, \dots, D_N) and a *TCL*– (R_1, \dots, R_N) of complexity $O(C_{\text{TCL}})$. We can write $C_{\text{TCL}} = \sum_{k=1}^N C_k$ where C_k is the complexity of the k^{th} n –mode product. Note that the order in which the products are taken doesn’t matter due to the commutativity of the n –mode product over disjoint modes (e.g. it is commutative for $\tilde{\mathcal{X}} \times_i \mathbf{U}^{(i)} \times_j \mathbf{U}^{(j)}$ as long as $i \neq j$). However, for illustrative purposes, we consider them to be done in order, from the first mode to the N^{th} . We then have:

$$C_k = R_k \times D_k \prod_{i=1}^{k-1} R_i \prod_{j=k+1}^N D_j \quad (6)$$

It follows that the overall complexity of the TCL is:

$$C_{\text{TCL}} = \sum_{k=1}^N \prod_{i=1}^k R_i \prod_{j=k}^N D_j \quad (7)$$

Comparison with a fully-connected layer A fully-connected layer with H hidden units has complexity $O(C_{\text{FC}})$, with:

$$C_{\text{FC}} = H \prod_{i=1}^N D_i \quad (8)$$

Method	Added TCL	1 st fully connected	2 nd fully connected	Accuracy (%)	space savings (%)
Baseline	-	4096 hidden units	4096 hidden units	65.41	0
Added TCL	TCL-(256, 3, 3)	4096 hidden units	4096 hidden units	65.53	-0.25
Added TCL	TCL-(192, 3, 3)	3072 hidden units	3072 hidden units	65.92	43.28
Added TCL	TCL-(128, 3, 3)	2048 hidden units	2048 hidden units	66.57	74.49
1 TCL substitution	-	TCL-(256, 3, 3)	4096 hidden units	65.52	62.77
1 TCL substitution	-	TCL-(192, 3, 3)	3072 hidden units	65.95	78.72
1 TCL substitution	-	TCL-(128, 3, 3)	2048 hidden units	64.95	90.25
2 TCL substitutions	-	TCL-(256, 3, 3)	TCL-(256, 3, 3)	62.98	98.64
2 TCL substitutions	-	TCL-(192, 3, 3)	TCL-(144, 3, 3)	62.06	99.22

Table 1. Results with AlexNet on CIFAR100. The first column presents the method, the second specifies whether a tensor contraction was added and when this is the case, the size of the TCL. Columns 3 and 4 specify the number of hidden units in the fully connected layers or the size of the TCL used instead when relevant. Column 5 presents the top-1 accuracy on the validation. Finally, the last column presents the reduction factor in the number of parameters in the fully connected layers (which represent, as previously mentioned, more than 80% of the total number of parameters of the networks) where the reference is the original network without any modification (*Baseline*).

Consider a TCL that maintains the size of its input, i.e., for any k in $[1 \dots N]$, $R_k = D_k$. In other words, $C_k = D_k \prod_{i=1}^N D_i$. Therefore,

$$C_{\text{TCL}} = \sum_{k=1}^N D_k \prod_{i=1}^N D_i \quad (9)$$

By comparison, a fully-connected layer that also maintains the size of its input, i.e. $H = \prod_{k=1}^N D_k$, would have a complexity of:

$$C_{\text{FC}} = \left(\prod_{i=1}^N D_i \right)^2 \quad (10)$$

Notice the product in the fully-connected case versus a sum for the TCL case.

2.2. Incorporating TCL in a network

We see several straightforward ways to incorporate the tensor contraction layer into existing neural network architectures.

TCL as An Additional Layer First, we can insert a tensor contraction layer following the last pooling layer, reducing the dimensionality of the activation tensor before feeding it to the subsequent two fully connected layers and softmax output of the network. In general, flattening induces a loss of information. By applying tensor contraction we reduce dimensionality efficiently by leveraging the multi dimensional dependencies in the activation tensor.

TCL as Replacement of a Fully Connected Layer We can also incorporate the TCL into existing architectures by completely replacing fully-connected layers. This has the advantage of significantly reducing the number of parameters in our model. Concretely, consider an activation tensor of size $(256, 7, 7)$ that is fed to either a fully-connected

layer (after having been flattened) or to a TCL. A fully-connected layer with 4096 hidden units has $256 \times 7 \times 7 \times 4096 = 51,380,224$ parameters. A TCL that preserves the size of its input, on the other hand, only has $256^2 + 7^2 + 7^2 = 1,712,622$ parameters. The TCL has 30 times fewer parameters than the fully-connected layer. Similarly, a TCL-(128, 5, 5) (approximately half size) will have only $256 \times 128 + 7 \times 5 + 7 \times 5 = 32,838$ parameters, or 1,500 times fewer parameters than a fully-connected layer.

3. Experiments

Our experiments investigate the representational power of the tensor contraction layer, demonstrating results on the CIFAR100 dataset [13]. Subsequently, we offer some preliminary results on the ImageNet 1k dataset [5]. We hypothesize that a tensor contraction layer can efficiently represent an activation tensor for processing by subsequent layers of the network, allowing for a large reduction in parameters without a reduction in accuracy.

We conduct our investigation on CIFAR100 using the AlexNet [14] and VGG [19] architectures, each modified to take 32×32 images as inputs. We also present results with a traditional AlexNet on ImageNet. In all cases we report the accuracy (top-1) as well as the space saved, which we quantify as:

$$\text{space saving} = 1 - \frac{n_{\text{original}}}{n_{\text{TCL}}}$$

where n_{original} is the sum of the number of parameters in the fully connected layers of the standard network and n_{TCL} is the sum of the number of parameters in the fully-connected layers of the network modified to include the TCL.

To avoid vanishing or exploding gradients, and to make the TCL more robust to changes in the initialization of the factors, we added a batch normalization layer [8] before and after the TCL.

Method	Added TCL	1 st fully connected	2 nd fully connected	Accuracy (%)	space savings (%)
Baseline	-	4096 hidden units	4096 hidden units	69.98	0
Added TCL	TCL-(512, 3, 3)	4096 hidden units	4096 hidden units	70.07	-0.73
Added TCL	TCL-(384, 3, 3)	3072 hidden units	3072 hidden units	68.56	42.99
Added TCL	TCL-(256, 3, 3)	2048 hidden units	2048 hidden units	67.57	74.35
1 TCL substitution	-	TCL-(512, 3, 3)	4096 hidden units	69.71	45.8
1 TCL substitution	-	TCL-(384, 3, 3)	3072 hidden units	68.83	69.16
1 TCL substitution	-	TCL-(256, 3, 3)	2048 hidden units	68.51	85.98
2 TCL substitutions	-	TCL-(512, 3, 3)	TCL-(512, 3, 3)	67.20	97.27
2 TCL substitutions	-	TCL-(384, 3, 3)	TCL-(288, 3, 3)	67.38	98.43

Table 2. Results obtained on CIFAR100 using a VGG-19 network architecture with different variations of the Tensor Contraction Layer. In all cases we report Top-1 Accuracy and space savings with respect to the baseline. As observed with the AlexNet, TCL allows for large space savings with minimal impact on performance and even improvement in some cases.

3.1. Results on CIFAR100

The CIFAR100 dataset is composed of 100 classes containing 600 32×32 images each, with 500 training images and 100 testing images per class. In all cases, we report performance on the testing set in term of accuracy (Top-1). We implemented all models using the MXNet library [2] and ran all experiments training with data parallelism across multiple GPUs on Amazon Web Services, with two NVIDIA k80 GPUs.

Because both the original AlexNet and VGG architectures were defined for the ImageNet data set, which has a larger input image size, to adapt them for CIFAR100 by adjusting the stride size on the input convolution layer of both networks so that they would take 32×32 input images. We investigate two sets of experiments, described below.

Added TCL In the first experiments, we add TCL as additional layer after the last pooling layer and perform the contraction along the two spacial modes of the image, leaving the modes corresponding to the channel and the batch size untouched. We gradually reduced the number of hidden units in these last two layers with and without the TCL included and retrain the nets until convergence to demonstrate how the tensor contraction layer can learn more compact representations without compromising accuracy.

TCL substitution In this case, we completely replace one or both of the fully-connected layers by a tensor contraction layer. We reduce the number of hidden units in the subsequent layers proportionally to the reduction in the size of the activation tensor.

Network architectures We experimented with an AlexNet, with an adjusted stride and filter size in the final convolutional layer. From the last convolutional layer, we get an activation tensor of size $(batch_size, 256, 3, 3)$. Similarly, in the case of the VGG network, we obtain activation tensors of size $(batch_size, 512, 3, 3)$. We experiment with

several variations of the tensor contraction layer. First, we consider the case where we project the activations to a tensor of identical shape. Additionally, we evaluate the effect of reducing the dimensionality of the activation tensor by 25% and by 50%. For AlexNet, because the spatial modes already compact are already, we preserve the spatial dimensions, and reduce dimensionality along the channel.

3.1.1 Results

Table 1 summarizes our results on CIFAR100 using the AlexNet, while results with VGG are presented in Table 2. The first column presents the method, the second specifies whether a tensor contraction was added and when this is the case, the size of the contracted core. Columns 3 and 4 specify the number of hidden units in the fully connected layers or the size of the TCL used instead when relevant. Column 5 presents the top-1 accuracy on the validation. Finally, the last column presents the reduction factor in the number of parameters in the fully connected layers (which represent, as previously mentioned, more than 80% of the total number of parameters of the networks) where the reference is the original network without any modification (*Baseline*).

A first observation is that adding a tensor contraction layer (*Added TCL* in Tables 1 and 2) consistently increases performance while having minimal impact on the overall number of parameters. Replacing the first fully-connected layer (*1 TCL substitution* in the Tables) allows us to reduce the number of parameters in the fully connected layers by a factor of more than 3, while observing the same performance as the original network. By replacing both fully connected layers (*2 TCL substitutions* in the Tables) we can obtain a reduction of more than $92\times$, with only a 2.5% decrease in performance.

3.2. Results on ImageNet

In this section, we present preliminary experiments using the larger ILSVRC 2012 (ImageNet) dataset [5], using

Method	Additional TCL	1 st fully connected	2 nd fully connected	Accuracy (in %)	space savings (%)
Baseline	-	4096 hidden units	4096 hidden units	56.29	0
Added TCL	TCL-(256, 5, 5)	4096 hidden units	4096 hidden units	57.54	-0.11
Added TCL	TCL-(200, 5, 5)	3276 hidden units	3276 hidden units	56.11	35.36
TCL substitution	-	TCL-(256, 5, 5)	4096 hidden units	56.57	35.49

Table 3. Results obtained with AlexNet on ImageNet, for a standard AlexNet (*baseline*), with an added Tensor Contraction Layer (*Added TCL*) and by replacing the first fully-connected layer with a TCL (*TCL substitution*). Simply adding the TCL results in a higher performance while having a minimal impact on the number of parameters in the fully connected layers. By reducing the size of the TCL or using a TCL to replace a fully connected layer, we can obtain a space savings of more than 35% with virtually no deterioration in performance.

the AlexNet architecture. ImageNet is composed of 1.2 million images for testing and 50,000 for validation and comprises 1,000 labeled classes.

For these experiments, we trained each network simultaneously on 4 NVIDIA k80 GPUs using data parallelism and report preliminary results. We report Top-1 accuracy on the validation set, across all 1000 classes. All experiments were run using the same setting.

Network architecture We use a standard AlexNet [14]. From the last convolutional layer, we get an activation tensor of size (**batch_size**, 256, 5, 5). As in the CIFAR100 case, we experiment with several variations of the tensor contraction layer. We first insert a TCL before the fully connected layers, either a size-preserving TCL (i.e. projecting to a tensor of the same size) or with a smaller size TCL and a proportionally smaller number of hidden units in the subsequent fully connected layers. We then experiment with replacing completely the first fully connected layer by a TCL.

3.2.1 Results

In Table 3 we summarize the results from a classical AlexNet (*Baseline*, first row), with an added tensor contraction layer (*Added TCL*) that preserves the dimensionality of its input (row 2) or reduces it (last row). We also report result for substituting the first fully connected layer with a TCL (*1 TCL substitution*, last row). Simply adding the TCL improves performance while the increase in number of parameters in the fully connected layers is negligible. We can obtain similar performance by first adding a TCL to reduce the dimensionality of the activation tensor and reducing the number of hidden units in the fully connected layers, leading to a large space saving with virtually no decrease in performance. Replacing the first fully connected layer with a size-preserving TCL results in a similar space savings while maintaining the same performance as the standard network.

4. Discussion

We introduced a new neural network layer that performs a tensor contraction on an activation tensor to yield a low

dimensional representation of it. By exploiting the natural multi-linear structure of the data in the activation tensor, where each mode of the tensor corresponds to a distinct modality of the data (i.e. the dimensions of the image and the channels) we are able to decrease the size of the data representation passed to later layers in the network without compromising accuracy on image recognition tasks.

The biggest practical contribution of the tensor contraction layer is the drastic reduction in the number of parameters with little to no performance penalty. This also allows neural networks to perform faster inference with fewer parameters by increasing their representational power. We demonstrated this via the performance of our Tensor Contraction Layer on the widely used CIFAR100 dataset with two established architectures, namely AlexNet and VGG. We also show results with AlexNet on the ImageNet dataset. Our proposed tensor contraction layer seems to be able to capture the underlying structure in the activation tensor and improve performance when added to an existing network. When we replace fully-connected layers with TCLs, we significantly reduce the number of parameters and nevertheless maintain (or in some cases even improve) performance.

Going forward, we plan to extend our work to more network architectures, especially in settings where raw data or learned representations exhibit natural multi-modal structure that we might capture via high-order tensors. We also endeavor to advance our experimental study of TCLs for large-scale, high-resolutions vision datasets. Given the time required to train a large network on such datasets we are investigating ways to reduce the dimension of the tensor contractions of an already trained model and simply fine tune. In addition, recent work [18] has shown that new extended BLAS primitives can avoid transpositions needed to compute the tensor contractions. This will further speed up the computations and we plan to implement it in future. Furthermore, we will look into methods to induce and exploit sparsity in the TCL, to understand the parameter reductions this method can yield over existing state of the art pruning methods. Finally, we are working on an extension to the TCL: a tensor regression layer to replace both the fully connected and final classification, potentially yielding increased accuracy with even greater parameter reductions.

References

- [1] A. Anandkumar, R. Ge, D. J. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832, 2014. [1](#)
- [2] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, and Z. Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *CoRR*, abs/1512.01274, 2015. [5](#)
- [3] Y. Chen, X. Jin, B. Kang, J. Feng, and S. Yan. Sharing residual units through collective tensor factorization in deep neural networks. 2017. [1](#)
- [4] N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor analysis. *CoRR*, abs/1509.05009, 2015. [1](#)
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. [4](#), [5](#)
- [6] B. D. Haeffele and R. Vidal. Global optimality in tensor factorization, deep learning, and beyond. *CoRR*, abs/1506.07540, 2015. [1](#)
- [7] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)*, 2016. [2](#)
- [8] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. [4](#)
- [9] M. Janzamin, H. Sedghi, and A. Anandkumar. Generalization bounds for neural networks through tensor factorization. *CoRR*, abs/1506.08473, 2015. [1](#)
- [10] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 79–86. ACM, 2010. [1](#)
- [11] Y. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. *CoRR*, abs/1511.06530, 2015. [1](#)
- [12] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM REVIEW*, 51(3):455–500, 2009. [2](#)
- [13] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009. [4](#)
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012. [2](#), [4](#), [6](#)
- [15] V. Lebedev, Y. Ganin, M. Rakhuba, I. V. Oseledets, and V. S. Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *CoRR*, abs/1412.6553, 2014. [1](#)
- [16] A. Novikov, D. Podoprikin, A. Osokin, and D. Vetrov. Tensorizing neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, NIPS’15, pages 442–450, 2015. [1](#)
- [17] H. Sedghi and A. Anandkumar. Training input-output recurrent neural networks through spectral methods. *CoRR*, abs/1603.00954, 2016. [1](#)
- [18] Y. Shi, U. N. Niranjan, A. Anandkumar, and C. Cecka. Tensor contractions with extended blas kernels on cpu and gpu. In *2016 IEEE 23rd International Conference on High Performance Computing (HiPC)*, pages 193–202, Dec 2016. [6](#)
- [19] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. [2](#), [4](#)
- [20] Y. Yang and T. M. Hospedales. Deep multi-task representation learning: A tensor factorisation approach. *CoRR*, abs/1605.06391, 2016. [1](#)