

Compressing CNN Kernels for Videos Using Tucker Decompositions: Towards Lightweight CNN Applications

Tobias Engelhardt Rasmussen (tenra@dtu.dk)^{*1}, Line H Clemmensen (lkhc@dtu.dk)¹,
and Andreas Baum (andba@dtu.dk)¹

¹Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark

Abstract

Convolutional Neural Networks (CNN) are the state-of-the-art in the field of visual computing. However, a major problem with CNNs is the large number of floating point operations (FLOPs) required to perform convolutions for large inputs. When considering the application of CNNs to video data, convolutional filters become even more complex due to the extra temporal dimension. This leads to problems when respective applications are to be deployed on mobile devices, such as smart phones, tablets, micro-controllers or similar, indicating less computational power.

Kim et al. proposed using a Tucker-decomposition to compress the convolutional kernel of a pre-trained network for images in order to reduce the complexity of the network, i.e. the number of FLOPs [5]. In this paper, we generalize the aforementioned method for application to videos (and other 3D signals) and evaluate the proposed method on a modified version of the THETIS data set, which contains videos of individuals performing tennis shots. We show that the compressed network reaches comparable accuracy, while indicating a memory compression by a factor of 51. However, the actual computational speed-up (factor 1.4) does not meet our theoretically derived expectation (factor 6).

1 Introduction

Neural networks (NNs) are powerful machine learning tools and their use on mobile phones and de-

vices with limited computational power is increasing, thus a general method for making the NNs more efficient is needed. In visual computing the convolutional neural networks (CNNs) have for many years been the state of the art due to their intuitive nature and their ability to learn features in images or videos using filters. A problem with CNNs is that they are computationally heavy. In 2013, Denil et al. concluded that NNs are often over-parametrized and that many of the weights are redundant [1]. In the same year, Sironi et al. found that a number of convolutional filters can be computed as a linear combination of a smaller number of separable filters thus exploiting the between-filter redundancy [9].

There have been numerous attempts to take advantage of this using different tensor decomposition methods. The usual approach is to compress the weights of a pre-trained network, change the network architecture to fit the subspace projections of the compressed kernel, and fine-tune the new network using the new weights. Attempts to speed up the convolutional operation include Lebedev et al. (2015)[6] using Canonical decomposition (CP), Wang et al. (2016)[10] using block-term decomposition, Kim et al. (2016)[5] using Tucker-decomposition, and Jaderberg et al. (2014)[3] using two original low-rank expansion schemes.

The convolutional operation itself can be understood as the sums of the element-wise multiplications of the filter with the image at a given location. The output of the convolution operation is generated by moving the filter iteratively across the entire image. This makes each parameter responsible for numerous floating point operations (FLOPs), thus making the operation complex and slow. Es-

^{*}Corresponding Author

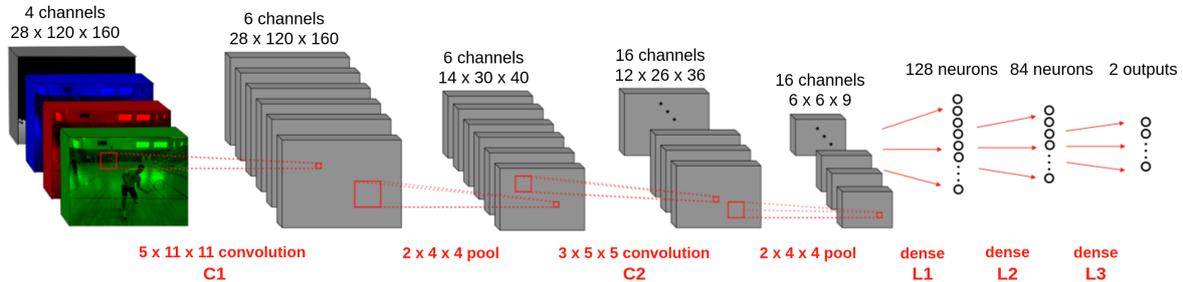


Figure 1: The original architecture used to learn the THETIS data set including layer names

pecially for videos, this becomes an increasing problem due to the additional, temporal dimension.

When classifying videos, rather than images, there are multiple ways of dealing with the added temporal information. The *early fusion* approach can be considered a 3-dimensional convolution because it also includes the temporal dimension of the video. This allows for detection of precise movement and speed because it is calculated *early* at the pixel-level [4]. Early fusion works well for classification of specific actions instead of general activities. The 3D convolutional approach intice us to extend the compression work by Kim et al.

Kim et al. proposed using a Tucker decomposition to compress both the convolutional kernel and the linear layer in a 2D CNN, thus developing a method able to compress an entire network using the same method, namely the *one-shot whole network compression* algorithm [5]. This approach yielded promising results, however is not directly applicable for video (3D) convolutions. We propose to generalize the method developed by Kim et al. for images, to also cover video convolutions. We evaluate the method on a modified version of the THETIS data set, [2], which is described in detail below, and using the architecture given in Figure 1.

After a short introduction to the data modifications in section 2 the proposed method will be presented in section 3. The experiments and results will be presented in section 4 and discussed in section 5.

2 The THETIS Data Set

The THETIS (THree dimEnsional TennIs Shots) action data set consists of 1980 videos of individ-

uals performing tennis shots. Each of the 12 shot types¹ has been performed multiple times by 55 individuals (31 beginners and 24 experts). Each observation consists of an RGB video of a single shot, a depth video (gray-scale), a silhouette video (binary), and both 2D and 3D skeleton videos. The videos have a resolution of 480×640 , while they vary in length (approximately 3-7 s). Although the videos are fairly standardized, they still contain a significant amount of variability due to different locations (arena with background noise/changing room), gender (male/female), age (kid-adult), handedness (right/left), and skill set (amateurs/professionals).

2.1 Modified THETIS data set

In order to ease the training of the model, the data has been modified as described in the following. First, only two types of shots (forehand and backhand) were selected, making it a binary classification problem. All videos were scaled down by a factor of four in each dimension. The lengths of the videos have been standardized by extracting the same number of frames (14) on either side of a manually assigned time-point reflecting the perceived midway through the shot. Lastly, the RGB video and the black and white depth video have been concatenated, resulting in four input channels for each observation (see Figure 1). With these modifications the resulting data set contains 327 observations that are all 4-way tensors of size: 4 channels \times 28 frames \times 120 vertical \times 160 horizontal pixels.

¹The types include: forehand, backhand, smash, service, volley, etc.[2]

3 Proposed Method

The proposed method follows the same course of action as in the work done by Kim et al.[5], however, the derivation of the compression of the convolutional kernel was described for image input data and will, therefore, be expanded in order to generalize to input data containing 4 dimensions, i.e. an extra video dimension.

3.1 Tucker-2 Decomposition of the 4D Convolutional Kernel

We will denote matrices using bold capitals \mathbf{X} with elements $\mathbf{X}(i, j) = x_{ij}$ and tensors with N modes using calligraphic letters $\mathcal{X}^{d_1 \times d_2 \times \dots \times d_N}$ with elements $\mathcal{X}(i_1, i_2, \dots, i_N) = x_{i_1 i_2 \dots i_N}$.

As described by Kim et al. we will use the Tucker-2 decomposition to compress the input and output channels (S and T) in order to exploit the cross-filter redundancy. The spatial and temporal dimensions of the kernel typically do not require compression due to their small size. The convolution of an input tensor \mathcal{X} of size $F \times H \times W \times S$ into an output tensor \mathcal{Y} of size $F' \times H' \times W' \times T$ using a video kernel is given by the linear mapping:

$$\mathcal{Y}(f', h', w', t) = \sum_{i=1}^{D_F} \sum_{j=1}^{D_H} \sum_{l=1}^{D_W} \sum_{s=1}^S \mathcal{K}(i, j, l, s, t) \mathcal{X}(f_i, h_j, w_l, s) \quad (1)$$

where:

$$f_i = (f' - 1) \Delta_F + i - P_F \quad (2)$$

$$h_j = (h' - 1) \Delta_H + j - P_H \quad (3)$$

$$w_l = (w' - 1) \Delta_W + l - P_W \quad (4)$$

Here Δ_\bullet is the stride, P_\bullet is the padding, and D_\bullet is the filter width for the given dimension F , H , W or S . (f', h', w', t) is the position in the output tensor \mathcal{Y} and \mathcal{K} is the 5-dimensional convolutional kernel, i.e. the stack of T 4-dimensional filters. The tucker-2 decomposition [7] of \mathcal{K} with respect to the input and output channel dimensions is given by:

$$\mathcal{K}(i, j, l, s, t) = \sum_{r_s=1}^{R_s} \sum_{r_t=1}^{R_t} \mathcal{C}(i, j, l, r_s, r_t) \mathbf{U}^{(s)}(s, r_s) \mathbf{U}^{(t)}(t, r_t) \quad (5)$$

Where \mathcal{C} is the core of the decomposition of size $D_F \times D_H \times D_W \times R_s \times R_t$ and \mathbf{U}^\bullet is the loading matrix along the input (s) or output (t) dimension, respectively, using rank R_s or R_t . Substituting this expression into (1) gives:

$$\mathcal{Y}(f', h', w', t) = \sum_{i=1}^{D_F} \sum_{j=1}^{D_H} \sum_{l=1}^{D_W} \sum_{s=1}^S \sum_{r_s=1}^{R_s} \sum_{r_t=1}^{R_t}$$

$$\mathcal{C}(i, j, l, r_s, r_t) \mathbf{U}^{(s)}(s, r_s) \mathbf{U}^{(t)}(t, r_t) \mathcal{X}(f_i, h_j, w_l, s) \quad (6)$$

By rearranging the sums we obtain:

$$\mathcal{Y}(f', h', w', t) = \sum_{r_t=1}^{R_t} \mathbf{U}^{(t)}(t, r_t) \sum_{i=1}^{D_F} \sum_{j=1}^{D_H} \sum_{l=1}^{D_W} \sum_{r_s=1}^{R_s}$$

$$\mathcal{C}(i, j, l, r_s, r_t) \underbrace{\sum_{s=1}^S \mathbf{U}^{(s)}(s, r_s) \mathcal{X}(f_i, h_j, w_l, s)}_{\mathcal{Q}(f_i, h_j, w_l, r_s)} \quad (7)$$

$$\underbrace{\dots}_{\mathcal{Q}'(f', h', w', r_t)}$$

Initially summing out s yields an intermediate tensor \mathcal{Q} of size $F \times H \times W \times R_s$. Subsequently another intermediate tensor \mathcal{Q}' of size $F' \times H' \times W' \times R_t$ is obtained by summing out everything except t . Using the intermediate tensors, the three linear mappings can be described as in equations 8, 9 and 10.

$$\mathcal{Q}(f, h, w, r_s) = \sum_{s=1}^S \mathbf{U}^{(s)}(s, r_s) \mathcal{X}(f, h, w, s) \quad (8)$$

$$\mathcal{Q}'(f', h', w', r_t) = \sum_{i=1}^{D_F} \sum_{j=1}^{D_H} \sum_{l=1}^{D_W} \sum_{r_s=1}^{R_s} \mathcal{C}(i, j, l, r_s, r_t) \mathcal{Q}(f_i, h_j, w_l, r_s) \quad (9)$$

$$\mathcal{Y}(f', h', w', t) = \sum_{r_t=1}^{R_t} \mathbf{U}^{(t)}(t, r_t) \mathcal{Q}'(f', h', w', r_t) \quad (10)$$

These equations correspond to the sequence of convolutions visualized in Figure 2. In detail we have:

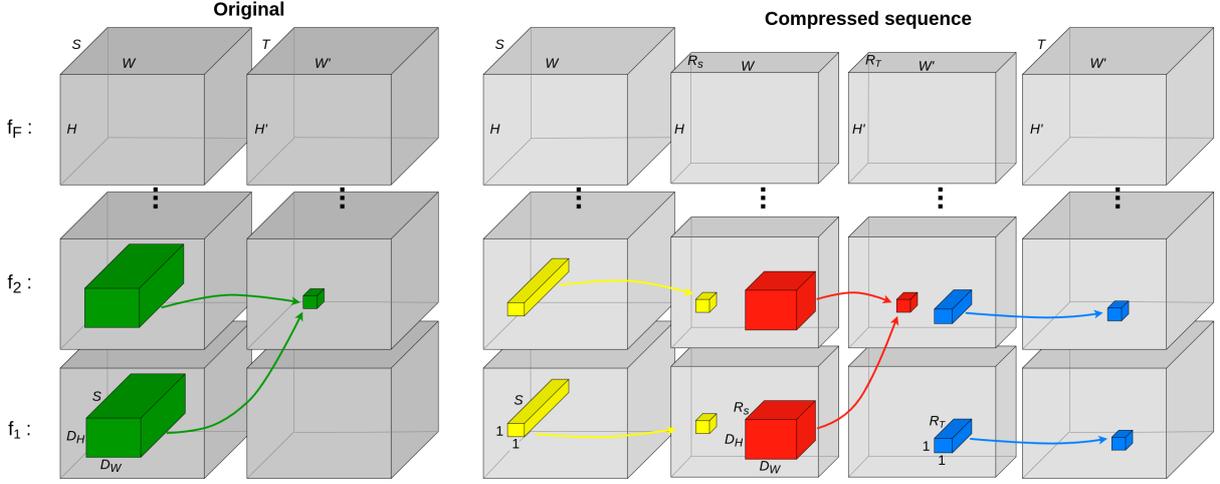


Figure 2: Visualization of the compression of a convolution on a video into a sequence of smaller convolutions based on [5]. For this visualisation the frame dimension is assumed to have adequate padding such that it does not change. Noticeably, D_F is the time dimension of the kernel \mathcal{K} , i.e. how many frames it stretches over (two in this illustration).

Table 1: The number of multiplications and number of parameters needed for the original convolutional operation and the convolutional operation using the compressed kernel, respectively.

	Original	Compressed
# multiplications	$S \cdot T \cdot \Gamma \cdot \Lambda$	$S \cdot R_s \cdot \Gamma + R_s \cdot R_t \cdot \Lambda \cdot \Gamma' + R_t \cdot T \cdot \Gamma'$
# parameters	$S \cdot T \cdot \Gamma$	$S \cdot R_s + R_s \cdot R_t \cdot \Gamma + R_t \cdot T$

- Yellow (eq. 8) - $1 \times 1 \times 1$ convolution with S input channels and R_s output channels (the rank of the decomposition).
- Red (eq. 9) - $D_F \times D_H \times D_W$ convolution like the original (green), but with R_s input channels and R_t output channels
- Blue (eq. 10) - $1 \times 1 \times 1$ convolution with R_t input channels and T output channels

The key to the theoretical speed-up is to evaluate the convolution on a smaller tensor using less filters. The first $1 \times 1 \times 1$ convolution reduces the number of input channels to R_s , hence, the convolution is carried out on R_s input channels and using only R_t filters, and in the end another $1 \times 1 \times 1$ convolution is bringing the number of channels from R_t to T .

The ranks R_t and R_s are estimated using Variational Bayesian Matrix Factorization (VBMF). They are assumed to be small under the assumption that filter parameter redundancy exists. Low ranks

are desirable in order to decrease the number of parameters and corresponding number of FLOPs. A complexity analysis is presented below illustrating how ranks can be chosen appropriately.

3.1.1 Complexity Analysis

In this section we discuss the number of parameters and multiplications (proportional to FLOPs) required to perform a convolution operation on an input video of size $F \times H \times W$ into an output video of size $F' \times H' \times W'$ using a kernel of size $D_F \times D_H \times D_W \times S \times T$. In the following we define $\Gamma = F \cdot H \cdot W$, $\Gamma' = F' \cdot H' \cdot W'$ and $\Lambda = D_F \cdot D_H \cdot D_W$ (number of pixels per channel in the input image, output image, and kernel respectively). The number of multiplications and parameters needed to perform the original vs. compressed mapping are given in Table 1. The ratios between the original and the compressed version for, both, the number of multiplications and the

Table 2: Overview of the layers of the decomposed network resulting after applying the *one-shot whole network compression* scheme to the architecture given in Figure 1 including number of weights and FLOPs for each layer. The number of FLOPs have been used to calculate the theoretical speed up, and the weights have been used to calculate the storage improvement, where $\times x$ means an improvement of x times. The observed time is given as the mean and standard deviation of 500 evaluations. In the bracket, the time for each sub-layer is given.

Layer	Comp.	S/R_{in}	T/R_{out}	Weights	FLOPs	CPU time (ms)
C1		4	6	14.5K	15609.2M	1231.60 \pm 67.06
Comp	Tucker2	2	2	2.5K	2618.6M (= 7.5 + 2600.9 + 10.2)	877.46 \pm 29.27 (= 1.43 + 872.63 + 3.40)
Impr				$\times 5.94$	$\times 5.96$	$\times 1.40$
C2		6	16	7.2K	161.6M	0.82 \pm 0.04
Comp	Tucker2	2	3	526	11.4M (= 0.4 + 10.8 + 0.9)	0.81 \pm 2.09 (= 0.17 + 0.39 + 0.25)
Impr				$\times 13.72$	$\times 14.2$	$\times 1.02$
L1		16	128	663.7K	1.3M	0.70 \pm 0.05
Comp	Tucker2	4	7	10.1K	60.0K (= 40.2 + 18.1 + 1.7)	0.36 \pm 0.55 (= 0.11 + 0.17 + 0.08)
Impr				$\times 65.32$	$\times 22.12$	$\times 1.97$
L2		128	84	10.8K	21.5K	0.08 \pm 0.51
Comp	Tucker1	-	1	296	423 (= 255 + 168)	0.07 \pm 0.04 (= 0.03 + 0.03)
Impr				$\times 36.61$	$\times 50.84$	$\times 1.15$
L3		84	2	170	336	0.03 \pm 0.002
Comp	-	-	-	170	336	0.03 \pm 0.002
Impr				$\times 1.00$	$\times 1.00$	$\times 1.03$
Total				696.4K	15772.1M	1238.56 \pm 67.33
Comp				13.6K	2630.1M	884.02 \pm 29.91
Impr				$\times 51.22$	$\times 6.0$	$\times 1.40$

number of parameters are bound by the ratio between the product of the original in- and output dimensions; and the product of the ranks of the compressed kernel, i.e. $\frac{S \cdot T}{R_s \cdot R_t}$.

3.2 One-Shot Network Compression

The *one-shot whole network compression* algorithm reported by Kim et al. [5] is used with minor modifications. VBMF is applied for rank selection in a given mode. Prior to this, the kernel is matrixized accordingly [7]. For the compression part it is noteworthy to mention that the Tucker-2 decomposition is used on all convolutional layers but the first as well as on the first linear layer (i.e. $C2 \dots CN_C$ and $L1$)². Tucker-1 is used on the remaining layers (i.e. $C1$ and $L2 \dots LN_L$). To allow for Tucker-2 decomposition of the first linear layer, it is compressed treating it as a convolutional layer. The full algorithm is provided in algorithm 1.

² C_{\bullet} and L_{\bullet} are convolutional respectively linear layers in the architecture with N_C convolutional layers followed by

4 Experiments and Results

The network architecture used to train the video data is illustrated in Figure 1. A hyper-parameter search was conducted using 5-fold cross validation due to the limited number of observations. After applying the *one-shot whole network compression* scheme, the compressed network was fine-tuned reaching the same accuracy (90.9%) as the original network; in very few epochs. The resulting network is shown in Table 2 along with the theoretical speed-up and storage improvements calculated from the number of FLOPs and parameters, respectively. The observed speed-up computed using the `profiler`-module in `torch` [8] is also given.

The total elapsed time also includes delays caused by non-linearities, transformations, and poolings, hence is not a direct sum of the individual layers. These extra computations are the same for both models, and therefore give a more truthful value for the speed-up of the entire model.

N_L linear layers

Algorithm 1: One-Shot Tucker Compression of a Video CNN

```
1 net  $\leftarrow$  Define appropriate CNN with convolutional layers  $C1, C2, \dots, CN_C$  and linear layers  $L1, L2, \dots, LN_L$ ;  
2 for layer in  $\{C2, \dots, CN_C, L1\}$  do  
3    $\mathcal{K}_{\text{layer}} \leftarrow$  Take out weight kernel;  
4    $R_s, R_t \leftarrow$  Choose ranks using VBMF;  
5    $\mathcal{G}, \mathbf{U}^{(s)}, \mathbf{U}^{(t)} \leftarrow$  Compress  $K_{\text{layer}}$  using Tucker-2 on in/out dimensions;  
6   comp1, comp2, comp3  $\leftarrow$  Define new layers;  
7    $\mathcal{K}_{\text{comp1}} \leftarrow \mathbf{U}^{(s)}$ ; //  $1 \times 1 \times 1$  convolution  
8    $\mathcal{K}_{\text{comp2}} \leftarrow \mathcal{G}$ ; //  $D_F \times D_H \times D_W$  convolution  
9    $\mathcal{K}_{\text{comp3}} \leftarrow \mathbf{U}^{(t)}$ ; //  $1 \times 1 \times 1$  convolution  
10   $\mathbf{b}_{\text{comp3}} \leftarrow \mathbf{b}_{\text{layer}}$ ; // Add bias to last layer  
11  layer  $\leftarrow$  Sequential(comp1, comp2, comp3)  
12 for layer in  $\{C1, L2, \dots, LN_L\}$  do  
13   $\mathcal{K}_{\text{layer}} \leftarrow$  Take out weight kernel;  
14   $R \leftarrow$  Choose rank using global analytical VBMF;  
15   $\mathcal{G}, \mathbf{U} \leftarrow$  Compress  $K_{\text{layer}}$  using Tucker-1 on out dimensions;  
16  comp1, comp2  $\leftarrow$  Define new layers;  
17   $\mathcal{K}_{\text{comp1}} \leftarrow \mathcal{G}$ ; //  $D_F \times D_H \times D_W$  convolution  
18   $\mathcal{K}_{\text{comp2}} \leftarrow \mathbf{U}$ ; //  $1 \times 1 \times 1$  convolution  
19   $\mathbf{b}_{\text{comp2}} \leftarrow \mathbf{b}_{\text{layer}}$ ; // Add bias to last layer  
20  layer  $\leftarrow$  Sequential(comp1, comp2)  
21 train layer on training data; // fine-tune the compressed network
```

From Table 2 it is clear that the improvements in terms of weights and FLOPs are substantial with an overall storage improvement of more than 51 times, and a theoretical speed-up of 6 times. This is not the case for the actual computation time that, even though it has decreased by a factor 1.4, does not meet that of the expected in any case.

5 Discussion

From the results, it is clear that the observed speed-up does not meet the expected. The key operation in the given scheme is the $1 \times 1 \times 1$ convolution that is used to bring down the dimensions for the actual convolution allowing it to run much faster. In the work by Kim et al. [5], they conclude that the lack of observed speed-up is caused by the cache-inefficiency of the 1×1 convolution which can be confirmed by the results achieved here. For example for the second convolution, the $1 \times 1 \times 1$ convolutions take up 31% and 22% of the time while they only correspond to 4% and 8% of the FLOPs, respectively. For the first, bigger convolution the observed speed-up is greater hence it seems that

the proposed scheme works better for bigger layers. Kim et al. also report greater observed speed-ups when running on smaller devices with less computational power, which makes it seem that the inefficiency of the $1 \times 1 \times 1$ convolution is more pronounced with more computational power, hence the proposed scheme is more appropriate for low-power applications.

PyTorch is well optimised in terms of evaluation due to its ability to make use of all computational power accessible [8]. This means that it is excellent at performing large operations quickly, and will likely be prone to penalties caused by additional yet smaller layers. This fact implies that caution should be taken when choosing optimal ranks as ranks too big will result in too many FLOPs according to section 3.1.1, but choosing ranks too small will increasingly result in penalties towards evaluation optimisation performed by PyTorch.

We believe that the theoretical speed-up could be approached, but that it would require the hardware to have a more linear relation between execution time and the number of operations required. For future work, this could be investigated by use of small devices with limited computational power

(that have no optimization of the calculations by e.g. parallelization) or by defining the calculations outside of the PyTorch framework and potentially translating it into C.

References

- [1] M. Denil, B. Shakibi, L. Dinh, M. A. Ranzato, and N. de Freitas. Predicting parameters in deep learning. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL <https://proceedings.neurips.cc/paper/2013/file/7fec306d1e665bc9c748b5d2b99a6e97-Paper.pdf>.
- [2] S. Gourgari, G. Goudelis, K. Karpouzis, and S. Kollias. Thetis: Three dimensional tennis shots a human action dataset. In *2013 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 676–681, 2013. URL <https://doi.org/10.1109/CVPRW.2013.102>.
- [3] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 2014. URL <http://dx.doi.org/10.5244/C.28.88>.
- [4] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. URL <https://doi.org/10.1109/CVPR.2014.223>.
- [5] Y. D. 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. *4th International Conference on Learning Representations, Iclr 2016 - Conference Track Proceedings*, 2016. URL <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85083951289&partnerID=40&md5=8115bfcad3c1b4338ff76ada9045ae94>.
- [6] V. Lebedev, Y. Ganin, M. Rakhuba, I. Oseledets, and V. Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015. URL <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85083952441&partnerID=40&md5=689b44a127d12ad8be206459e1de7187>.
- [7] M. Mørup. Applications of tensor (multiway array) factorizations and decompositions in data mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1):24–40, 2011. ISSN 1942-4787. URL <https://doi.org/10.1002/widm.1>.
- [8] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [9] R. Rigamonti, A. Sironi, V. Lepetit, and P. Fua. Learning separable filters. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2754–2761, 2013. URL <https://doi.org/10.1109/ICCV.2013.355>.
- [10] P. Wang and J. Cheng. Accelerating convolutional neural networks for mobile applications. In *Proceedings of the 24th ACM International Conference on Multimedia*, MM '16, page 541–545, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450336031. URL <https://doi.org/10.1145/2964284.2967280>.