

Factorized Higher-Order CNNs with an Application to Spatio-Temporal Emotion Estimation

Jean Kossaifi^{*†}
NVIDIA

Antoine Toisoul^{*}
Samsung AI Center

Adrian Bulat
Samsung AI Center

Yannis Panagakis[†]
University of Athens

Timothy Hospedales
Samsung AI Center

Maja Pantic[†]
Imperial College London

Abstract

Training deep neural networks with spatio-temporal (i.e., 3D) or multidimensional convolutions of higher-order is computationally challenging due to millions of unknown parameters across dozens of layers. To alleviate this, one approach is to apply low-rank tensor decompositions to convolution kernels in order to compress the network and reduce its number of parameters. Alternatively, new convolutional blocks, such as MobileNet, can be directly designed for efficiency. In this paper, we unify these two approaches by proposing a tensor factorization framework for efficient multidimensional (separable) convolutions of higher-order. Interestingly, the proposed framework enables a novel higher-order transduction, allowing to train a network on a given domain (e.g., 2D images or N -dimensional data in general) and using transduction to generalize to higher-order data such as videos (or $(N+K)$ -dimensional data in general), capturing for instance temporal dynamics while preserving the learnt spatial information.

We apply the proposed methodology, coined CP-Higher-Order Convolution (HO-CPConv), to spatio-temporal facial emotion analysis. Most existing facial affect models focus on static imagery and discard all temporal information. This is due to the above-mentioned burden of training 3D convolutional nets and the lack of large bodies of video data annotated by experts. We address both issues with our proposed framework. Initial training is first done on static imagery before using transduction to generalize to the temporal domain. We demonstrate superior performance on three challenging large scale affect estimation datasets, AffectNet, SEWA, and AFEW-VA.

^{*}Joint first authors.

[†]Jean Kossaifi, Yannis Panagakis and Maja Pantic were with Samsung AI Center, Cambridge and Imperial College London.

1. Introduction

With the unprecedented success of deep convolutional neural networks came the quest for training always deeper networks. However, while deeper neural networks give better performance when trained appropriately, that depth also translates into memory and computationally heavy models, typically with tens of millions of parameters. This is especially true when training higher-order convolutional nets –e.g. third order (3D) on videos. However, such models are necessary to perform predictions in the spatio-temporal domain and are crucial in many applications, including action recognition and emotion recognition.

In this paper, we depart from conventional approaches and propose a novel factorized multidimensional convolutional block that achieves superior performance through efficient use of the structure in the data. In addition, our model can first be trained on the image domain and extended seamlessly to transfer performance to the temporal domain. This novel transduction method is made possible by the structure of our proposed block. In addition, it allows one to drastically decrease the number of parameters, while improving performance and computational efficiency and it can be applied to already existing spatio-temporal network architectures such as a ResNet3D [47]. Our method leverages a CP tensor decomposition, in order to separately learn the disentangled spatial and temporal information of 3D convolutions. This improves accuracy by a large margin while reducing the number of parameters of spatio-temporal architectures and greatly facilitating training on video data.

In summary, we make the following contributions:

- We show that many of deep nets architectural improvements, such as *MobileNet* or *ResNet*'s *Bottleneck* blocks, are in fact drawn from the same larger family of tensor decomposition methods (Section 3). We propose a general framework unifying tensor decompositions and efficient architectures, showing how these ef-

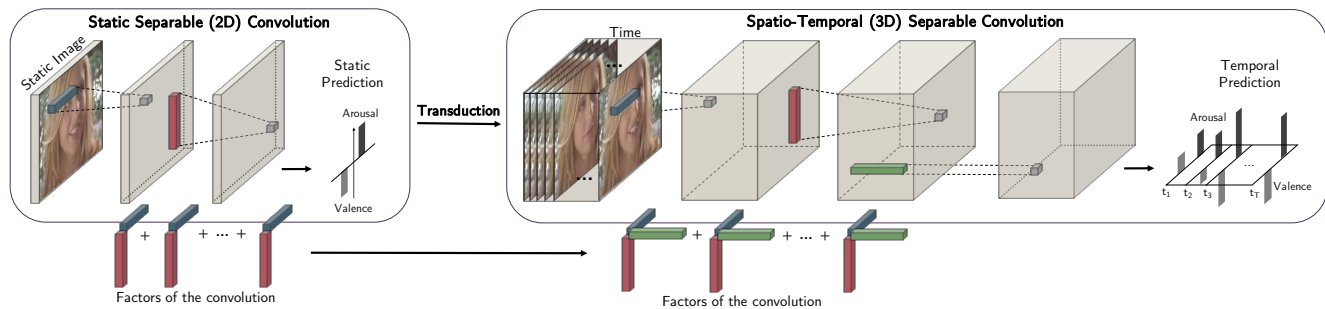


Figure 1: **Overview of our method**, here represented for a single channel of a single input. We start by training a 2D CNN with our proposed factorized convolutional block on static images (**left**). We then apply transduction to extend the model from the static to the spatio-temporal domain (**right**). The pretrained spatial factors (blue and red) are first kept fixed, before jointly fine-tuning all the parameters once the temporal factors (green) have been trained.

efficient architectures can be derived by applying tensor decomposition to convolutional kernels (Section 3.4)

- Using this framework, we propose factorized higher-order convolutional neural networks, that leverage efficient general multi-dimensional convolutions. These achieve the same performance with a fraction of the parameters and floating point operations (Section 4).
- Finally, we propose a novel mechanism called higher-order transduction which can be employed to convert our model, trained in N dimensions, to $N + K$ dimensions.
- We show that our factorized higher-order networks outperform existing works on static affect estimation on the AffectNet, SEWA and AFEW-VA datasets.
- Using transduction on the static models, we also demonstrate state-of-the-art results for continuous facial emotion analysis from video on both SEWA and AFEW-VA datasets.

2. Background and related work

Multidimensional convolutions arise in several mathematical models across different fields. They are a cornerstone of Convolutional Neural Networks (CNNs) [26, 28] enabling them to effectively learn from high-dimensional data by mitigating the curse of dimensionality [34]. However, CNNs are computationally demanding, with the cost of convolutions being dominant during both training and inference. As a result, there is an increasing interest in improving the efficiency of multidimensional convolutions.

Several efficient implementations of convolutions have been proposed. For instance, 2D convolution can be efficiently implemented as matrix multiplication by converting the convolution kernel to a Toeplitz matrix. However, this procedure requires replicating the kernel values multiple

times across different matrix columns in the Toeplitz matrix, thus increasing the memory requirements. Implementing convolutions via the im2col approach is also memory intensive due to the space required for building the column matrix. These memory requirements may be prohibitive for mobile or embedded devices, hindering the deployment of CNNs in resource-limited platforms.

In general, most existing attempts at efficient convolutions are isolated and there currently is no unified framework to study them. In particular we are interested in two different branches of work, which we review next. Firstly, approaches that leverage tensor methods for efficient convolutions, either to compress or reformulate them for speed. Secondly, approaches that directly formulate efficient neural architecture, e.g., using separable convolutions.

Tensor methods for efficient deep networks The properties of tensor methods [14, 41, 17] make them a prime choice for deep learning. Beside theoretical study of the properties of deep neural networks [3], they have been especially studied in the context of reparametrizing existing layers [49]. One goal of such reparametrization is parameter space savings [7]. [31] for instance proposed to reshape the weight matrix of fully-connected layers into high-order tensors with a Tensor-Train (TT) [32] structure. In a follow-up work [4], the same strategy is also applied to convolutional layers. Fully connected layers and flattening layers can be removed altogether and replaced with tensor regression layers [21]. These express outputs through a low-rank multi-linear mapping from a high-order activation tensor to an output tensor of arbitrary order. Parameter space saving can also be obtained, while maintaining multi-linear structure, by applying tensor contraction [20].

Another advantage of tensor reparametrization is computational speed-up. In particular, a tensor decomposition is an efficient way of obtaining separable filters from convolutional kernels. These separable convolutions were proposed

in computer vision by [36] in the context of filter banks. [12] first applied this concept to deep learning and proposed leveraging redundancies across channels using separable convolutions. [1, 27] proposed to apply CP decomposition directly to the (4-dimensional) kernels of pretrained 2D convolutional layers, making them separable. The incurred loss in performance is compensated by fine-tuning.

Efficient rewriting of convolutions can also be obtained using Tucker decompositions instead of CP to decompose the convolutional layers of a pre-trained network [15]. This allows rewriting the convolution as a 1×1 convolution, followed by regular convolution with a smaller kernel and another 1×1 convolution. In this case, the spatial dimensions of the convolutional kernel are left untouched and only the input and output channels are compressed. Again, the loss in performance is compensated by fine-tuning the whole network. Finally, [46] propose to remove redundancy in convolutional layers and express these as the composition of two convolutional layers with less parameters. Each 2D filter is approximated by a sum of rank-1 matrices. Thanks to this restricted setting, a closed-form solution can be readily obtained with SVD. Here, we unify the above works and propose factorized higher-order (separable) convolutions.

Efficient neural networks While concepts such as separable convolutions have been studied since the early successes of deep learning using tensor decompositions, they have only relatively recently been “rediscovered” and proposed as standalone end-to-end trainable efficient neural network architectures. The first attempts in the direction of neural network architecture optimization were proposed early in the ground-breaking VGG network [42] where the large convolutional kernels used in AlexNet [25] were replaced with a series of smaller ones that have an equivalent receptive field size: i.e. a convolution with a 5×5 kernel can be replaced by two consecutive convolutions of size 3×3 . In parallel, the idea of decomposing larger kernels into a series of smaller ones is explored in the multiple iterations of the Inception block [43, 44, 45] where a convolutional layer with a 7×7 kernel is approximated with two 7×1 and 1×7 kernels. [8] introduced the so-called *bottleneck* module that reduces the number of channels on which the convolutional layer with higher kernel size (3×3) operate on by projecting back and forth the features using two convolutional layers with 1×1 filters. [48] expands upon this by replacing the 3×3 convolution with a grouped convolutional layer that can further reduce the complexity of the model while increasing representational power at the same time. Recently, [10] introduced the *MobileNet* architecture where they proposed to replace the 3×3 convolutions with a depth-wise separable module: a depth-wise 3×3 convolution (the number of groups is equal to the number of channels) followed by a 1×1 convolutional layer that aggregates the information. These type of structures were shown to of-

fer a good balance between the performance offered and the computational cost they incur. [39] goes one step further and incorporates the idea of using separable convolutions in an *inverted bottleneck* module. The proposed module uses 1×1 layers to expand and then contract the channels (hence inverted bottleneck) while using separable convolutions for the 3×3 convolutional layer.

Facial affect analysis is the first step towards better human-computer interactions. Early research focused on detecting discrete emotions such as happiness and sadness, based on the hypothesis that these are *universal*. However, this categorization of human affect is limited and does not cover the wide emotional spectrum displayed by humans on a daily basis. Psychologists have since moved towards more fine grained dimensional measures of affect [35, 38]. The goal is to estimate *continuous* levels of valence –how positive or negative an emotional display is– and arousal –how exciting or calming is the emotional experience. This task is the subject of most of the recent research in affect estimation [37, 19, 23, 50] and is the focus of this paper.

Valence and arousal are dimensional measures of affect that vary in time. It is these changes that are important for accurate human-affect estimation. Capturing temporal dynamics of emotions is therefore crucial and requires video rather than static analysis. However, spatio-temporal models are difficult to train due to their large number of parameters, requiring very large amounts of annotated videos to be trained successfully. Unfortunately, the quality and quantity of available video data and annotation collected in naturalistic conditions is low [40]. As a result, most work in this area of affect estimation in-the-wild focuses on affect estimation from static imagery [23, 30]. Estimation from videos is then done on a frame-by-frame basis. Here, we tackle both issues and train spatio-temporal networks that outperform existing methods for affect estimation.

3. Convolutions in a tensor framework

In this section, we explore the relationship between tensor methods and deep neural networks’ convolutional layers. Without loss of generality, we omit the batch size in all the following formulas.

Mathematical background and notation We denote 1st-order tensors (vectors) as \mathbf{v} , 2nd-order tensor (matrices) as \mathbf{M} and tensors of order ≥ 3 as \mathcal{T} . We denote a regular convolution of \mathcal{X} with \mathbf{W} as $\mathcal{X} \star_n \mathbf{W}$. For 1-D convolutions, we write the convolution of a tensor $\mathcal{X} \in \mathbb{R}^{I_0, \dots, I_N}$ with a vector $\mathbf{v} \in \mathbb{R}^K$ along the n^{th} -mode as $\mathcal{X} \star_n \mathbf{v}$. In practice, as done in current deep learning frameworks [33], we use cross-correlation, which differs from a convolution by a flip of the kernel. This does not impact the results since the weights are learned end-to-end. In other words, $(\mathcal{X} \star_n \mathbf{v})_{i_0, \dots, i_N} = \sum_{k=1}^K \mathbf{v}_k \mathcal{X}_{i_0, \dots, i_{n-1}, i_n+k, i_{n+1}, \dots, i_N}$.

3.1. 1×1 convolutions and tensor contraction

We show that 1×1 convolutions are equivalent to a tensor contraction with the kernel of the convolution along the channels dimension. Let's consider a 1×1 convolution Φ , defined by kernel $\mathcal{W} \in \mathbb{R}^{T \times C \times 1 \times 1}$ and applied to an activation tensor $\mathcal{X} \in \mathbb{R}^{C \times H \times W}$. We denote the squeezed version of \mathcal{W} along the first mode as $\mathbf{W} \in \mathbb{R}^{T \times C}$.

The tensor contraction of a tensor $\mathcal{T} \in \mathbb{R}^{I_0 \times I_1 \times \dots \times I_N}$ with matrix $\mathbf{M} \in \mathbb{R}^{J \times I_n}$, along the n^{th} -mode ($n \in [0 \dots N]$), known as n -mode product, is defined as $\mathcal{P} = \mathcal{T} \times_n \mathbf{M}$, with: $\mathcal{P}_{i_0, \dots, i_N} = \sum_{k=0}^{I_n} \mathcal{T}_{i_0, \dots, i_{n-1}, k, i_{n+1}, \dots, i_N} \mathbf{M}_{i_n, k}$

By plugging this into the expression of $\Phi(\mathcal{X})$, we readily observe that the 1×1 convolution is equivalent with an n -mode product between \mathcal{X} and the matrix \mathbf{W} :

$$\Phi(\mathcal{X})_{t,y,x} = \mathcal{X} \star \mathcal{W} = \sum_{k=0}^C \mathcal{W}_{t,k,y,x} \mathcal{X}_{k,y,x} = \mathcal{X} \times_0 \mathbf{W}$$

3.2. Kruskal convolutions

Here we show how separable convolutions can be obtained by applying CP decomposition to the kernel of a regular convolution [27]. We consider a convolution defined by its kernel weight tensor $\mathcal{W} \in \mathbb{R}^{T \times C \times K_H \times K_W}$, applied on an input of size $\mathbb{R}^{C \times H \times W}$. Let $\mathcal{X} \in \mathbb{R}^{C \times H \times W}$ be an arbitrary activation tensor. If we define the resulting feature map as $\mathcal{F} = \mathcal{X} \star \mathcal{W}$, we have:

$$\mathcal{F}_{t,y,x} = \sum_{k=1}^C \sum_{j=1}^H \sum_{i=1}^W \mathcal{W}(t, k, j, i) \mathcal{X}(k, j + y, i + x) \quad (1)$$

Assuming a low-rank Kruskal structure on the kernel \mathcal{W} (which can be readily obtained by applying CP decomposition), we can write:

$$\mathcal{W}_{t,s,j,i} = \sum_{r=0}^{R-1} \mathbf{U}_{t,r}^{(T)} \mathbf{U}_{s,r}^{(C)} \mathbf{U}_{j,r}^{(H)} \mathbf{U}_{i,r}^{(W)} \quad (2)$$

By plugging 2 into 1 and re-arranging the terms, we get:

$$\mathcal{F}_{t,y,x} = \sum_{r=0}^{R-1} \mathbf{U}_{t,r}^{(T)} \left(\sum_{i=1}^W \mathbf{U}_{i,r}^{(W)} \left(\sum_{j=1}^H \mathbf{U}_{j,r}^{(H)} \left[\sum_{k=1}^C \mathbf{U}_{k,r}^{(C)} \mathcal{X}(k, j + y, i + x) \right] \right) \right)$$

1 × 1 conv
depthwise conv
depthwise conv
1 × 1 convolution

This allows to replace the original convolution by a series of efficient depthwise separable convolutions [27], figure 2.

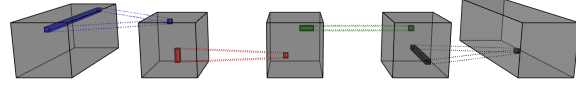


Figure 2: Illustration of a 2D Kruskal convolution.

3.3. Tucker convolutions

As previously, we consider the convolution $\mathcal{F} = \mathcal{X} \star \mathcal{W}$. However, instead of a Kruskal structure, we now assume a low-rank Tucker structure on the kernel \mathcal{W} (which can be readily obtained by applying Tucker decomposition) and yields an efficient formulation [15]. We can write:

$$\mathcal{W}(t, s, j, i) = \sum_{r_0=0}^{R_0-1} \sum_{r_1=0}^{R_1-1} \sum_{r_2=0}^{R_2-1} \sum_{r_3=0}^{R_3-1} \mathcal{G}_{r_0, r_1, r_2, r_3} \mathbf{U}_{t, r_0}^{(T)} \mathbf{U}_{s, r_1}^{(C)} \mathbf{U}_{j, r_2}^{(H)} \mathbf{U}_{i, r_3}^{(W)}$$

Plugging back into a convolution, we get:

$$\mathcal{F}_{t,y,x} = \sum_{k=1}^C \sum_{j=1}^H \sum_{i=1}^W \sum_{r_0=0}^{R_0-1} \sum_{r_1=0}^{R_1-1} \sum_{r_2=0}^{R_2-1} \sum_{r_3=0}^{R_3-1} \mathcal{G}_{r_0, r_1, r_2, r_3} \mathbf{U}_{t, r_0}^{(T)} \mathbf{U}_{k, r_1}^{(C)} \mathbf{U}_{j, r_2}^{(H)} \mathbf{U}_{i, r_3}^{(W)} \mathcal{X}_{k, j + y, i + x}$$

We can further absorb the factors along the spacial dimensions into the core by writing $\mathcal{H} = \mathcal{G} \times_2 \mathbf{U}_{j, r_2}^{(H)} \times_3 \mathbf{U}_{i, r_3}^{(W)}$.

In that case, the expression above simplifies to:

$$\mathcal{F}_{t,y,x} = \sum_{k=1}^C \sum_{j=1}^H \sum_{i=1}^W \sum_{r_0=0}^{R_0-1} \sum_{r_1=0}^{R_1-1} \mathcal{H}_{r_0, r_1, j, i} \mathbf{U}_{t, r_0}^{(T)} \mathbf{U}_{k, r_1}^{(C)} \mathcal{X}_{k, j + y, i + x} \quad (3)$$

In other words, this is equivalence to first transforming the number of channels, then applying a (small) convolution before returning from the rank to the target number of channels. This can be seen by rearranging the terms from Equation 3:

$$\mathcal{F}_{t,y,x} = \sum_{r_0=0}^{R_0-1} \mathbf{U}_{t, r_0}^{(T)} \left[\sum_{j=1}^H \sum_{i=1}^W \sum_{r_1=0}^{R_1-1} \mathcal{H}_{r_0, r_1, j, i} \underbrace{\left[\sum_{k=1}^C \mathbf{U}_{k, r_1}^{(C)} \mathcal{X}(k, j + y, i + x) \right]}_{1 \times 1 \text{ conv}} \right]$$

H × W conv
1 × 1 conv

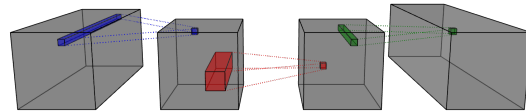


Figure 3: Illustration of a Tucker convolution expressed as a series of small efficient convolutions. Note that this is the approach taken by ResNet for the Bottleneck blocks.

In short, this simplifies to the following expression, also illustrated in Figure 3:

$$\mathcal{F} = \left(\left(\mathcal{X} \times_0 \mathbf{U}^{(C)} \right) \star \mathcal{G} \right) \times_0 \mathbf{U}^{(T)} \quad (4)$$

3.4. Efficient architectures in a tensor framework

While tensor decompositions have been explored in the field of mathematics for decades and in the context of deep learning for years, they are regularly rediscovered and re-introduced in different forms. Here, we revisit popular deep neural network architectures under the lens of tensor factorization. Specifically, we show how these blocks can be obtained from a regular convolution by applying tensor decomposition to its kernel. In practice, batch-normalisation layers and non-linearities are inserted in between the intermediary convolution to facilitate learning from scratch.

ResNet Bottleneck block [9] introduced a block, coined *Bottleneck block* in their seminal work on deep residual networks. It consists in a series of a 1×1 convolution, to reduce the number of channels, a smaller regular (3×3) convolution, and another 1×1 convolution to restore the rank to the desired number of output channels. Based on the equivalence derived in Section 3.3, it is straightforward to see this as applying Tucker decomposition to the kernel of a regular convolution.

ResNext and Xception ResNext [48] builds on this bottleneck architecture, which, as we have shown, is equivalent to applying Tucker decomposition to the convolutional kernel. In order to reduce the rank further, the output is expressed as a sum of such bottlenecks, with a lower-rank. This can be reformulated efficiently using grouped-convolution [48]. In parallel, a similar approach was proposed by [2], but without 1×1 convolution following the grouped depthwise convolution.

MobileNet v1 MobileNet v1 [10] uses building blocks made of a depthwise separable convolutions (spatial part of the convolution) followed by a 1×1 convolution to adjust the number of output channels. This can be readily obtained from a CP decomposition (Section 3.2) as follows: first we write the convolutional weight tensor as detailed in Equation 2, with a rank equal to the number of input channels, i.e. $R = C$. The first depthwise-separable convolution can be obtained by combining the two spatial 1D convolutions $\mathbf{U}^{(H)}$ and $\mathbf{U}^{(W)}$. This results into a single spatial factor $\mathbf{U}^{(S)} \in \mathbb{R}^{H \times W \times R}$, such that $\mathbf{U}_{j,i,r}^{(S)} = \mathbf{U}_{j,r}^{(H)} \mathbf{U}_{i,r}^{(W)}$. The 1×1 convolution is then given by the matrix-product of the remaining factors, $\mathbf{U}^{(F)} = \mathbf{U}^{(T)} (\mathbf{U}^{(C)})^T \in \mathbb{R}^{T \times C}$. This is illustrated in Figure 4.

MobileNet v2 MobileNet v2 [39] employs a similar approach by grouping the spatial factors into one spatial factor $\mathbf{U}^{(S)} \in \mathbb{R}^{H \times W \times R}$, as explained previously for the case of MobileNet. However, the other factors are left untouched. The rank of the decomposition, in this case, corresponds, for each layer, to the expansion factor \times the number of input channels. This results in two 1×1 convolutions and a 3×3 depthwise separable convolution. Finally, the kernel weight

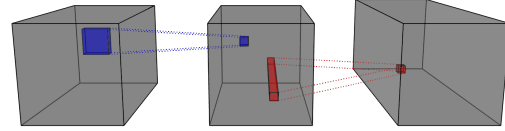


Figure 4: **MobileNet blocks are a special case of CP convolutions**, without the first convolution, and with spatial factors are combined into one.

tensor (displayed graphically in Figure 5) is expressed as:

$$\mathcal{W}_{t,s,j,i} = \sum_{r=0}^{R-1} \mathbf{U}_{t,r}^{(T)} \mathbf{U}_{s,r}^{(C)} \mathbf{U}_{j,i,r}^{(S)} \quad (5)$$

In practice, MobileNet-v2 also includes batch-normalisation layers and non-linearities as well as a skip connection to facilitate learning.

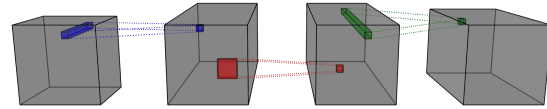


Figure 5: **MobileNet-v2 blocks are a special case of CP convolutions**, with the spatial factors merged into a depthwise separable convolution.

4. Factorized higher-order convolutions

We propose to generalize the framework introduced above to convolutions of any arbitrary order. Specifically, we express, in the general case, separable ND-convolutions as a series of tensor contractions and 1D convolutions. We show how this is derived from a CP convolution on the N -dimensional kernel. We then detail how to expand our proposed factorized higher-order convolutions, trained in N -dimensions to $(N + 1)$ dimensions.

Efficient N-D convolutions via higher-order factorization In particular, here, we consider an $N + 1$ th-order input activation tensor $\mathcal{X} \in \mathbb{R}^{C \times D_0 \times \dots \times D_{N-1}}$ corresponding to N dimensions with C channels. We define a general, high order separable convolution Φ defined by a kernel $\mathcal{W} \in \mathbb{R}^{T \times C \times K_0 \times \dots \times K_{N-1}}$, and expressed as a Kruskal tensor, i.e. $\mathcal{W} = \llbracket \boldsymbol{\lambda}; \mathbf{U}^{(T)}, \mathbf{U}^{(C)}, \mathbf{U}^{(K_0)}, \dots, \mathbf{U}^{(K_{N-1})} \rrbracket$. We can then write:

$$\begin{aligned} \Phi(\mathcal{X})_{t,i_0,\dots,i_{N-1}} &= \sum_{r=0}^R \sum_{s=0}^C \sum_{i_0=0}^{K_0} \dots \\ &\dots \sum_{i_{N-1}=0}^{K_{N-1}} \lambda_r [\mathbf{U}_{t,r}^{(T)} \mathbf{U}_{s,r}^{(C)} \mathbf{U}_{i_0,r}^{(K_0)} \dots \mathbf{U}_{i_{N-1},r}^{(K_{N-1})} \mathcal{X}_{s,i_0,\dots,i_{N-1}}] \end{aligned}$$

By rearranging the terms, this expression can be rewrit-

ten as:

$$\mathcal{F} = \left(\rho \left(\mathcal{X} \times_0 \mathbf{U}^{(T)} \right) \right) \times_0 \left(\text{diag}(\boldsymbol{\lambda}) \mathbf{U}^{(C)} \right) \quad (6)$$

where ρ applies the 1D spatial convolutions:

$$\rho(\mathcal{X}) = \left(\mathcal{X} \star_1 \mathbf{U}^{(K_0)} \star_2 \mathbf{U}^{(K_1)} \star \dots \star_{N+1} \mathbf{U}^{(K_{N-1})} \right)$$

Tensor decompositions (and, in particular, decomposed convolutions) are notoriously hard to train end-to-end [12, 27, 46]. As a result, most existing approaches rely on first training an uncompressed network, then decomposing the convolutional kernels before replacing the convolution with their efficient rewriting and fine-tuning to recover lost performance. However, this approach is not suitable for higher-order convolutions where it might not be practical to train the full N-D convolution. It is possible to facilitate training from scratch by absorbing the magnitude of the factors into the weight vector $\boldsymbol{\lambda}$. We can also add non-linearities Ψ (e.g. batch normalisation combined with RELU), leading to the following expression, resulting in an efficient higher-order CP convolution:

$$\mathcal{F} = \rho \left(\Psi \left(\mathcal{X} \times_0 \mathbf{U}^{(T)} \right) \right) \times_0 \left(\text{diag}(\boldsymbol{\lambda}) \mathbf{U}^{(C)} \right) \quad (7)$$

Skip connection can also be added by introducing an additional factor $\mathbf{U}^{(S)} \in \mathbb{R}^{C \times T}$ and using $\mathcal{F}' = \mathcal{X} + (\mathcal{F} \times_0 \mathbf{U}^{(S)})$.

This formulation is significantly more efficient than that of a regular convolution. Let's consider an N-dimensional convolution, with C input channels and T output channels, i.e. a weight of size $\mathcal{W} \in \mathbb{R}^{C \times T \times I_0 \times \dots \times I_{N-1}}$. Then a regular 3D convolution has $C \times T \times \left(\prod_{k=0}^{N-1} I_k \right)$ parameters. By contrast, our HO-CP convolution with rank R has only $R \left(C + T + \sum_{k=0}^{N-1} I_k \right) + 1$ parameters. The +1 term accounts for the weights $\boldsymbol{\lambda}$. For instance, for a 3D convolution with a cubic kernel (of size $K \times K \times K$, a regular 3D convolution would have CTK^3 parameters, versus only $R(C + T + 3K)$ for our proposed factorized version.

This reduction in the number of parameters translates into much more efficient operation in terms of floating point operations (FLOPs). We show, in Figure 6, a visualisation of the number of Giga FLOPs (GFLOPs, with 1GFLOP = $1e^9$ FLOPs), for both a regular 3D convolution and our proposed approach, for an input of size $32 \times 32 \times 16$, varying the number of input and output channels, with a kernel size of $3 \times 3 \times 3$.

Higher-Order Transduction Here, we introduce *transduction*, which allows to first train an N -dimensional convolution and expand it to $(N + K)$ dimensions, $K > 0$.

Thanks to the efficient formulation introduced in Equation 7, we now have an effective way to go from N dimensions to $N + 1$. We place ourselves in the same setting as for Equation 7, where we have a regular N-D convolution with spatial dimensions and extend the model to

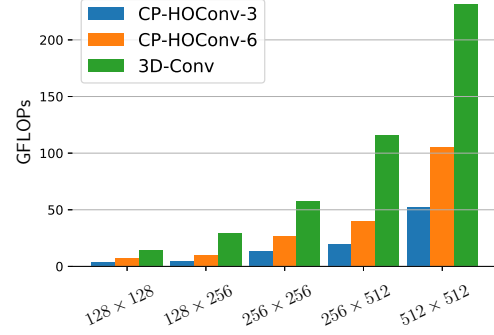


Figure 6: **Comparison of the number of Giga-FLOPs between regular 3D convolutions and our proposed method.** We consider inputs of size $32 \times 32 \times 16$, and vary the number of the input and output channels (the x-axis shows input \times output channels). Our proposed CP-HO convolution, here for a rank equal to 6 and 3 times the input channels (*CP-HOConv-6* and *CP-HOConv-3*), has significantly less FLOPs than regular convolution (*3D-Conv*).

$N + 1$ dimensions. To do so, we introduce a new factor $\mathbf{U}^{(K_{N+1})} \in \mathbb{R}^{(K_{N+1} \times R)}$ corresponding to the new $N + 1$ th dimension. The final formulation is then :

$$\mathcal{F} = \hat{\rho} \left(\Psi \left(\mathcal{X} \times_0 \mathbf{U}^{(T)} \right) \right) \times_0 \left(\text{diag}(\boldsymbol{\lambda}) \mathbf{U}^{(C)} \right), \quad (8)$$

with $\hat{\rho}(\mathcal{X}) = \rho(\mathcal{X}) \star_{N+1} \mathbf{U}^{(K_{N+1})}$.

Note that only the new factor needs to be trained, e.g. transduction can be done by simply training $K_{N+1} \times R$ additional parameters.

Automatic Rank Selection Our proposed factorized higher-order convolutions introduce a new additional parameter, corresponding to the rank of the factorization. This can be efficiently incorporated in the formulation by introducing a vector of weights, represented by $\boldsymbol{\lambda}$ in Equation 8. This allows us to automatically tune the rank of each of the layers by introducing an additional *Lasso* term in the loss function, e.g., an ℓ_1 regularization on $\boldsymbol{\lambda}$. Let $\boldsymbol{\lambda}_l$ be the vector of weights for each layer $l \in [0 \dots L - 1]$ of our neural network, associated with a loss \mathcal{L} . The overall loss with regularization will become $\mathcal{L}_{\text{reg}} = \mathcal{L} + \gamma \sum_{l=0}^{L-1} |\boldsymbol{\lambda}_l|$, where γ controls the amount of sparsity in the weights.

5. Experimental setting

Datasets for In-The-Wild Affect Estimation We validate the performance of our approach on established large scale datasets for continuous affect estimation in-the-wild.

AffectNet [30] is a large static dataset of human faces and labelled in terms of facial landmarks, emotion categories as well as valence and arousal values. It contains

more than a million images, including 450,000 manually labelled by twelve expert annotators.

AFEW-VA [23] is composed of video clips taken from feature films and accurately annotated *per frame* in terms of continuous levels of valence and arousal. Besides 68 accurate facial landmarks are also provided for each frame.

SEWA [24] is the largest video dataset for affect estimation in-the-wild. It contains over 2000 minutes of audio and video data annotated in terms of facial landmarks, valence and arousal values. It contains 398 subjects from six different cultures, is gender balanced and uniformly spans the age range 18 to 65.

Implementation details We implemented all models using PyTorch [33] and TensorLy [22]. In all cases, we divided the dataset in subject independent training, validation and testing sets. For our factorized higher-order convolutional network, we further removed the flattening and fully-connected layers and replace them with a single tensor regression layer [21] in order to fully preserve the structure in the activations. For training, we employed an Adam optimizer [16] and validated the learning rate in the range $[10^{-5}; 0.01]$, the beta parameters in the range $[0.0; 0.999]$ and the weight decay in the range $[0.0; 0.01]$ using a randomized grid search. We also decreased the learning rate by a factor of 10 every 15 epochs. The regularization parameter γ was validated in the range $[10^{-4}; 1.0]$ on AffectNet and consequently set to 0.01 for all other experiments.

For our baseline we use both a 3D ResNet and a ResNet 2+1D [47], both with a ResNet-18 backbone. For our method, we use the same ResNet-18 architecture but replace each of the convolutional layers with our proposed higher-order factorized convolution. We initialised the rank so that the number of parameters would be the same as the original convolution. When performing transduction, the added temporal factors to the CP convolutions are initialized to a constant value of one. In a first step, these factors are optimized while keeping the remaining parameters fixed. The whole network is then fine-tuned. This avoids the transducted factors to pollute what has already been learnt in the static case. The full process is summarized in Figure 1.

Performance metrics and loss function In all cases, we report performance for the *RMSE*, *SAGR*, *PCC*, and *CCC* which are common metrics employed in affect estimation. Let \mathbf{y} be a ground-truth signal and $\hat{\mathbf{y}}$ the associated prediction by the model.

The **RMSE** is the well known Root Mean Square Error: $RMSE(\mathbf{y}, \hat{\mathbf{y}}) = \sqrt{\mathbb{E}((\mathbf{y} - \hat{\mathbf{y}})^2)}$.

The **SAGR** assesses whether the sign of the two signals agree: $SAGR(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{n} \sum_{i=1}^n \delta(\text{sign}(\mathbf{y}_i), \text{sign}(\hat{\mathbf{y}}_i))$.

The **PCC** is the Pearson product-moment correlation coefficient and measures how correlated the two signals are: $PCC(\mathbf{y}, \hat{\mathbf{y}}) = \frac{\mathbb{E}(\mathbf{y} - \mu_{\mathbf{y}})(\hat{\mathbf{y}} - \mu_{\hat{\mathbf{y}}})}{\sigma_{\mathbf{y}}\sigma_{\hat{\mathbf{y}}}}$.

The **CCC** is Lin’s Concordance Correlation Coefficient and assesses the correlation of the two signals but also how close the two signals are: $CCC(\mathbf{y}, \hat{\mathbf{y}}) = \frac{2\sigma_{\mathbf{y}}\sigma_{\hat{\mathbf{y}}}PCC(\mathbf{y}, \hat{\mathbf{y}})}{\sigma_{\hat{\mathbf{y}}}^2 + \sigma_{\mathbf{y}}^2 + (\mu_{\mathbf{y}} - \mu_{\hat{\mathbf{y}}})^2}$.

The goal for continuous affect estimation is typically to maximize the correlations coefficients PCC and CCC. However minimizing the RMSE also helps maximizing the correlations as it gives a lower error in each individual prediction. Our regression loss function reflects this by incorporating three terms: $\mathcal{L} = \frac{1}{\alpha + \beta + \gamma} (\alpha \mathcal{L}_{RMSE} + \beta \mathcal{L}_{PCC} + \gamma \mathcal{L}_{CCC})$, with $\mathcal{L}_{RMSE} = RMSE_{\text{valence}} + RMSE_{\text{arousal}}$, $\mathcal{L}_{PCC} = 1 - \frac{PCC_{\text{valence}} + PCC_{\text{arousal}}}{2}$ and $\mathcal{L}_{CCC} = 1 - \frac{CCC_{\text{valence}} + CCC_{\text{arousal}}}{2}$. The coefficients α , β and γ are shake-shake regularization coefficients [5] sampled randomly in the range $[0; 1]$ following a uniform distribution. This ensures none of the terms are ignored during optimization. On AffectNet, where discrete classes of emotions are available, we jointly perform a regression of the valence and arousal values and a classification of the emotional class by adding a cross entropy to the loss function.

6. Performance evaluation

In this section, we report the performance of our methods for facial affect estimation in the wild. First, we report results on static images. We then show how the higher-order transduction allows us to extend these models to the temporal domain. In all cases, we compare with the state-of-the-art and show superior results.

Static affect analysis in-the-wild with factorized CNNs First we show the performance of our models trained and tested on individual (static) imagery. We train our method on AffectNet, which is the largest database but consists of static images only. There, our method outperforms all other works by a large margin (Table 1). We observe similar results on SEWA (Table 2) and AFEW-VA (Table 3). In the supplementary document, we also report results on LSEMSW [11] and CIFAR10 [25].

Temporal prediction via higher-order transduction We then apply transduction as described in the method section to convert the static model from SEWA (Table 2, temporal case) and AFEW-VA (Table 3, temporal case) to the temporal domain, where we simply train the added temporal factors. This approach allows to efficiently train temporal models on even small datasets. Our method outperforms other approaches, despite having only 11 million parameters, compared to 33 million for the corresponding 3D ResNet18, and 31 million parameters for a (2+1)D ResNet-18. Interestingly, in all cases, we notice that valence is better predicted than arousal, which, is in line with finding by psychologists that humans are better at estimating valence from visual data [38, 6].

Using the automatic rank selection procedure detailed in

Table 1: Results on the AffectNet dataset

Network	Acc.	Valence				Arousal			
		RMSE	SAGR	PCC	CCC	RMSE	SAGR	PCC	CCC
AffectNet baseline [30]	0.58	0.37	0.74	0.66	0.60	0.41	0.65	0.54	0.34
Face-SSD [13]	-	0.44	0.73	0.58	0.57	0.39	0.71	0.50	0.47
VGG-Face+2M imgs [18]	0.60	0.37	0.78	0.66	0.62	0.39	0.75	0.55	0.54
Baseline ResNet-18	0.55	0.35	0.79	0.68	0.66	0.33	0.8	0.58	0.57
Ours	0.59	0.35	0.79	0.71	0.71	0.32	0.8	0.63	0.63

Table 2: Results on the SEWA database

Case	Network	Valence				Arousal			
		RMSE	SAGR	PCC	CCC	RMSE	SAGR	PCC	CCC
Static	[24]	-	-	0.32	0.31	-	-	0.18	0.20
	VGG16+TRL [29]	0.33	-	0.50	0.47	0.39	-	0.44	0.39
	ResNet-18	0.37	0.62	0.33	0.29	0.52	0.62	0.28	0.19
	Ours	0.33	0.65	0.64	0.6	0.39	0.75	0.48	0.44
Temporal	ResNet-3D	0.37	0.59	0.47	0.41	0.41	0.69	0.29	0.21
	ResNet-(2+1)D	0.35	0.63	0.59	0.49	0.41	0.63	0.39	0.31
	Ours – scratch	0.33	0.63	0.62	0.54	0.40	0.72	0.42	0.32
	Ours – transduction	0.24	0.69	0.84	0.75	0.32	0.80	0.60	0.52

Table 3: Results on the AFEW-VA database

Case	Network	Valence				Arousal			
		RMSE	SAGR	PCC	CCC	RMSE	SAGR	PCC	CCC
Static	RF Hybrid DCT [23]	0.27	-	0.407	-	0.23	-	0.45	-
	ResNet50+TRL [29]	0.40	-	0.33	0.33	0.41	-	0.42	0.4
	ResNet-18	0.43	0.42	0.05	0.03	0.41	0.68	0.06	0.05
	Ours	0.24	0.64	0.55	0.55	0.24	0.77	0.57	0.52
Temporal	Baseline ResNet-18-3D	0.26	0.56	0.19	0.17	0.22	0.77	0.33	0.29
	ResNet-18-(2+1)D	0.31	0.50	0.17	0.16	0.29	0.73	0.33	0.20
	AffWildNet [19]	-	-	0.51	0.52	-	-	0.575	0.556
	Ours – scratch	0.28	0.53	0.12	0.11	0.19	0.75	0.23	0.15
	Ours – transduction	0.20	0.67	0.64	0.57	0.21	0.79	0.62	0.56

section 4, we let the model learn end-to-end the rank of each of the factorized higher-order convolutions. We found that on average, 8 to 15% of the parameters can be set to zero for optimal performance. In practice, about 1 million (of the 11 million) parameters were set to zero by the Lasso regularization. An in-depth study of the effect of the automatic rank selection is provided in the supplementary document.

7. Conclusion

We established the link between tensor factorizations and efficient convolutions, in a unified framework. Based on

this, we proposed a factorized higher-order (N -dimensional) convolutional block. This results in efficient models that outperform traditional networks, while being more computationally and memory efficient. We also introduced a higher-order transduction algorithm for converting the expressive power of trained N -dimensional convolutions to any $N + K$ dimensions. We then applied our approach to continuous facial affect estimation in naturalistic conditions. Using transduction, we transferred the performance of the models trained on static images to the temporal domain and reported state-of-the-art results in both cases.

References

- [1] Marcella Astrid and Seung-Ik Lee. Cp-decomposition with tensor power method for convolutional neural networks compression. *CoRR*, abs/1701.07148, 2017. [3](#)
- [2] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Computer Vision and Pattern Recognition*, 2017. [5](#)
- [3] Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on Learning Theory*, pages 698–728, 2016. [2](#)
- [4] Timur Garipov, Dmitry Podoprikin, Alexander Novikov, and Dmitry Vetrov. Ultimate tensorization: compressing convolutional and fc layers alike. *NIPS workshop: Learning with Tensors: Why Now and How?*, 2016. [2](#)
- [5] Xavier Gastaldi. Shake-shake regularization. *arXiv preprint arXiv:1705.07485*, 2017. [7](#)
- [6] Michael Grimm and Kristian Kroschel. Emotion estimation in speech using a 3d emotion space concept. In Michael Grimm and Kristian Kroschel, editors, *Robust Speech*, chapter 16. IntechOpen, Rijeka, 2007. [7](#)
- [7] Julia Gusak, Maksym Kholiavchenko, Evgeny Ponomarev, Larisa Markeeva, Philip Blagoveschensky, Andrzej Cichocki, and Ivan Oseledets. Automated multi-stage compression of neural networks. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. [2](#)
- [8] Kaifeng He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, 2016. [3](#)
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, 2016. [5](#)
- [10] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. [3](#), [5](#)
- [11] Guosheng Hu, Li Liu, Yang Yuan, Zehao Yu, Yang Hua, Zhihong Zhang, Fumin Shen, Ling Shao, Timothy Hospedales, Neil Robertson, et al. Deep multi-task learning to recognise subtle facial expressions of mental states. In *ECCV*, pages 103–119, 2018. [7](#)
- [12] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. In *British Machine Vision Conference*, 2014. [3](#), [6](#)
- [13] Youngkyoon Jang, Hatice Gunes, and Ioannis Patras. Registration-free face-ssd: Single shot analysis of smiles, facial attributes, and affect in the wild. *Computer Vision and Image Understanding*, 2019. [8](#)
- [14] Majid Janzamin, Rong Ge, Jean Kossaifi, Anima Anandkumar, et al. Spectral learning on matrices and tensors. *Foundations and Trends® in Machine Learning*, 12(5-6):393–536, 2019. [2](#)
- [15] Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. *ICLR*, 2016. [3](#), [4](#)
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [7](#)
- [17] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM REVIEW*, 51(3):455–500, 2009. [2](#)
- [18] Dimitrios Kollias, Shiyang Cheng, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Generating faces for affect analysis. *CoRR*, abs/1811.05027, 2018. [8](#)
- [19] Dimitrios Kollias, Panagiotis Tzirakis, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, Björn Schuller, Irene Kotsia, and Stefanos Zafeiriou. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision*, 127(6-7):907–929, 2019. [3](#), [8](#)
- [20] Jean Kossaifi, Aran Khanna, Zachary Lipton, Tommaso Furlanello, and Anima Anandkumar. Tensor contraction layers for parsimonious deep nets. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017. [2](#)
- [21] Jean Kossaifi, Zachary C. Lipton, Aran Khanna, Tommaso Furlanello, and Anima Anandkumar. Tensor regression networks. *CoRR*, abs/1707.08308, 2018. [2](#), [7](#)
- [22] Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. Tensorly: Tensor learning in python. *Journal of Machine Learning Research*, 20(26):1–6, 2019. [7](#)
- [23] Jean Kossaifi, Georgios Tzimiropoulos, Sinisa Todorovic, and Maja Pantic. Afew-va database for valence and arousal estimation in-the-wild. *Image and Vision Computing*, 65:23–36, 2017. [3](#), [7](#), [8](#)
- [24] J. Kossaifi, R. Walecki, Y. Panagakis, J. Shen, M. Schmitt, F. Ringeval, J. Han, V. Pandit, A. Toisoul, B. W. Schuller, K. Star, E. Hajiyev, and M. Pantic. Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019. [7](#), [8](#)
- [25] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009. [3](#), [7](#)
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. [2](#)
- [27] Vadim Lebedev, Yaroslav Ganin, Maksim Rakhuba, Ivan V. Oseledets, and Victor S. Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. In *ICLR*, 2015. [3](#), [4](#), [6](#)
- [28] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998. [2](#)
- [29] A. Mitenkova, J. Kossaifi, Y. Panagakis, and M. Pantic. Valence and arousal estimation in-the-wild with tensor methods. In *IEEE International Conference on Automatic Face Gesture Recognition (FG 2019)*, 2019. [8](#)
- [30] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017. [3](#), [6](#), [8](#)

- [31] Alexander Novikov, Dmitry Podoprikin, Anton Osokin, and Dmitry Vetrov. Tensorizing neural networks. In *Neural Information Processing Systems*, 2015. 2
- [32] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33(5):2295–2317, Sept. 2011. 2
- [33] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017. 3, 7
- [34] T Poggio and Q Liao. Theory i: Deep networks and the curse of dimensionality. *Bulletin of the Polish Academy of Sciences: Technical Sciences*, pages 761–773, 01 2018. 2
- [35] Jonathan Posner, James A Russell, and Bradley S Peterson. The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and psychopathology*, 17(3):715–734, 2005. 3
- [36] R. Rigamonti, A. Sironi, V. Lepetit, and P. Fua. Learning separable filters. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, June 2013. 3
- [37] Fabien Ringeval, Björn Schuller, Michel Valstar, Nicholas Cummins, Roddy Cowie, Leili Tavabi, Maximilian Schmitt, Sina Alisamir, Shahin Amiriparian, Eva-Maria Messner, et al. Audio/visual emotion challenge 2019: State-of-mind, detecting depression with ai, and cross-cultural affect recognition. 2019. 3
- [38] J.A. Russell, Bachorowski J.A., and J.M. Fernandez-Dols. Facial and vocal expressions of emotions. *Annu. Rev. Psychol.*, 54:329–349, 2003. 3, 7
- [39] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 3, 5
- [40] Evangelos Sariyanidi, Hatice Gunes, and Andrea Cavallaro. Automatic analysis of facial affect: A survey of registration, representation, and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(6):1113–1133, 2015. 3
- [41] Nicholas D Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E Papalexakis, and Christos Faloutsos. Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, 2017. 2
- [42] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [43] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 3
- [44] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015. 3
- [45] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 3
- [46] Cheng Tai, Tong Xiao, Xiaogang Wang, and Weinan E. Convolutional neural networks with low-rank regularization. *CoRR*, abs/1511.06067, 2015. 3, 6
- [47] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. 1, 7
- [48] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 3, 5
- [49] Yongxin Yang and Timothy M. Hospedales. Deep multi-task representation learning: A tensor factorisation approach. In *ICLR*, 2017. 2
- [50] Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotzia. Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017. 3