

Augmenting Deep Classifiers with Polynomial Neural Networks

Grigorios G. Chrysos*¹[0000-0002-0650-1856], Markos Georgopoulos*²[0000-0001-5928-515X], Jiankang Deng³, Jean Kossaifi⁴, Yannis Panagakis⁵, and Anima Anandkumar⁴

¹ EPFL, Switzerland

² Imperial College London, UK

³ Huawei, UK

⁴ NVIDIA, USA

⁵ University of Athens, Greece

* First two authors have contributed equally.

grigorios.chrysos@epfl.ch, m.georgopoulos@imperial.ac.uk

Abstract. Deep neural networks have been the driving force behind the success in classification tasks, e.g., object and audio recognition. Impressive results and generalization have been achieved by a variety of recently proposed architectures, the majority of which are seemingly disconnected. In this work, we cast the study of deep classifiers under a unifying framework. In particular, we express state-of-the-art architectures (e.g., residual and non-local networks) in the form of different degree polynomials of the input. Our framework provides insights on the inductive biases of each model and enables natural extensions building upon their polynomial nature. The efficacy of the proposed models is evaluated on standard image and audio classification benchmarks. The expressivity of the proposed models is highlighted both in terms of increased model performance as well as model compression. Lastly, the extensions allowed by this taxonomy showcase benefits in the presence of limited data and long-tailed data distributions. We expect this taxonomy to provide links between existing domain-specific architectures.

Keywords: Polynomial neural networks, tensor decompositions, polynomial expansions, classification

1 Introduction

The unprecedented performance of AlexNet [35] in ImageNet classification [50] led to the resurgence of research in the field of neural networks. Since then, an extensive corpus of papers has been devoted to improving classification performance by modifying the architecture of the neural network. However, only a handful (of seemingly disconnected) architectures, such as ResNet [22] or non-local neural networks [60], have demonstrated impressive generalization across different tasks (e.g., [70]), domains (e.g., [63]) and modalities (e.g., [31]). This phenomenon can be attributed to the challenging nature of devising a network

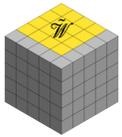
$\mathbf{Wz} + \mathcal{W} \times_2 \mathbf{z} \times_3 \mathbf{z} = \tilde{\mathcal{W}} \times_2 \tilde{\mathbf{z}} \times_3 \tilde{\mathbf{z}}$			
1 st degree	2 nd degree	1 st + 2 nd (shared)	1 st + 2 nd
			
Deep Classifiers			
ResNet, VGG16/19 AlexNet	SENet, SKNet	Π -net	PDC
 zero 1 st degree matrix 2 nd degree tensor shared 1 st +2 nd degree tensor			

Fig. 1: Parameter interactions for different degrees of polynomials. For visualization purpose we assume a second-degree expansion, where we have folded the learnable parameters into a tensor \mathcal{W} and the input is vectorized. The equation on the top is the second-degree polynomial in a tensor-format, while the $\tilde{\mathbf{z}}$ is a padded version of \mathbf{z} . Notice that first-degree polynomials [22, 35, 53] have zeros in a large part of the tensor \mathcal{W} , while similarly second-degree polynomials [25, 36] have the matrices connected to first-order interactions zero. On the contrary, Π -nets along with the proposed PDC capture both interactions. Importantly, we illustrate how PDC learns a more expressive model without the enforced sharing of the Π -net. The notation of mode- m product along with the derivation is conducted in sec. B (supplementary).

and the lack of understanding regarding the assumptions that come with its design, i.e., its inductive bias.

Demystifying the success of deep neural architectures is of paramount importance. In particular, significant effort has been devoted to the study of neural architectures, e.g., depth versus width of the neural network [18, 47] and the effect of residual connections on the training of the network [19, 22, 26]. In this work, we offer a principled approach to study state-of-the-art classifiers as polynomial expansions. We show that polynomials have been a recurring theme in numerous classifiers and interpret their design choices under a unifying framework.

The proposed framework provides a taxonomy for a collection of deep classifiers, e.g., a non-local neural network is a third-degree polynomial and ResNet is a first-degree polynomial. Thus, we provide an intuitive way to study and extend existing networks as visualized in Fig. 1, as well as interpret their gap in performance. Lastly, we design extensions on existing methods and show that we can improve their classification accuracy or achieve parameter reduction. Concretely, our contributions are the following:

- We express a collection of state-of-the-art neural architectures as polynomials. Our unifying framework sheds light on the inductive bias of each

architecture. We experimentally verify the performance of different methods of the taxonomy on four standard benchmarks.

- Our framework allows us to propose intuitive modifications on existing architectures. The proposed new architectures consistently improve upon their corresponding baselines, both in terms of accuracy as well as model compression.
- We evaluate the performance under various changes in the training distribution, i.e., limiting the number of samples per class or creating a long-tailed distribution. The proposed models improve upon the baselines in both cases.
- We release the code as open source to enable the reproduction of our results.

2 Fundamentals on polynomial expansions

In this section, we provide the notation and an intuitive explanation on how a polynomial expansion emerges on various types of variables.

Notation: Below, we symbolize matrices (vectors) with bold capital (lower) letters, e.g. $\mathbf{X}(\mathbf{x})$. A variable that can be either a matrix or a vector is denoted by $\hat{\cdot}$. Tensors are considered as the multidimensional equivalent of matrices. Tensors are symbolized by calligraphic boldface letters e.g., \mathcal{X} . The symbol $\vec{\mathbf{1}}$ denotes a vector of ones and \mathcal{I} is a third-order super-diagonal unit tensor. Due to constrained space, the detailed notation is on sec. A, while in Table 1 the core symbols are summarized.

Table 1: Symbols

Symbol	Dimension(s)	Definition
N	\mathbb{N}	Degree of polynomial expansion.
k	\mathbb{N}	Rank of the decompositions.
\mathbf{z}	\mathbb{R}^d	Vector-form input to the expansion.
$\odot, *$	-	Khatri-Rao product, Hadamard product.

Polynomial expansions: Below, polynomials express a relationship between an input variable (e.g., a scalar z) and (learnable) coefficients; this relationship only involves the two core operations of addition and multiplication. When the input variable is in vector form, e.g., $\mathbf{z} \in \mathbb{R}^\delta$ with $\delta \in \mathbb{N}$, then the polynomial captures the relationships between the different elements of the input vector. The input variable can also be a higher-dimensional structure, e.g., a matrix. This is frequently the case in computer vision, where one dimension can express spatial dimensions, while the other can express the features (channels). The polynomial can either capture the interactions across every element of the matrix with every other element, or it can have higher-order interactions between specific elements, e.g., the interactions of a row with each column.

Relationship between polynomial expansions and tensors: Multivariate polynomial expansions are intertwined with tensors. Specifically, the polynomial

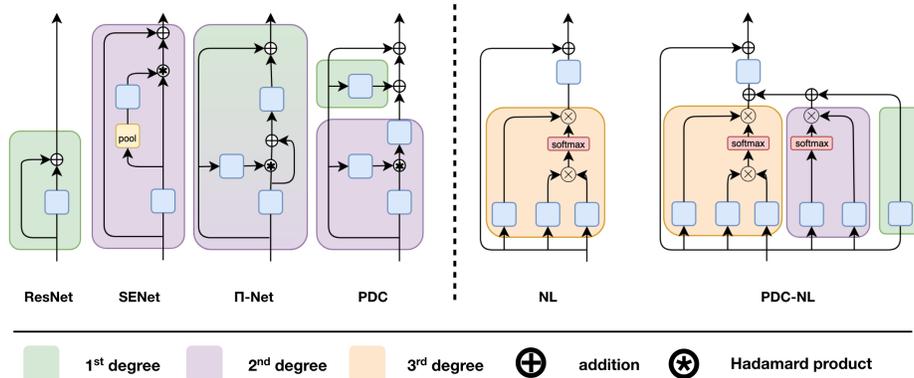


Fig. 2: Blocks (up to third-degree) from different architectures. The layers (i.e., blue boxes) denote any linear operation, e.g., a convolution or a fully-connected layer, depending on the architecture. From left to right, the degree of the polynomial is increasing. Our framework enables also to complete the missing terms of the polynomial (i.e., PDC-NL versus NL).

expansions of interest involve vector, matrix or tensor form as the input. Let us express the output $y \in \mathbb{R}$ as a N^{th} degree polynomial expansion of a d -dimensional input $z \in \mathbb{R}^d$:

$$y = \beta + \sum_{i=1}^d w_i^{[1]} z_i + \sum_{i=1}^d \sum_{j=1}^d w_{i,j}^{[2]} z_i z_j + \dots + \underbrace{\sum_{i=1}^d \dots \sum_{j=1}^d w_{i,\dots,k}^{[N]} z_i \dots z_k}_{N \text{ sums}}, \quad (1)$$

where $\beta \in \mathbb{R}$ is the constant term, and the w terms are the scaling parameters for every degree. Notice that $w^{[n]}$ for $n > 2$ depends on n indices, thus it can be expressed as a tensor. Collecting those tensors, $\{\mathcal{W}^{[n]} \in \mathbb{R}^{\underbrace{d \times \dots \times d}_{n \text{ times}}}\}_{n=1}^N$ represent learnable parameters in (1). Due to the constrained space this relationship is further quantified in sec. B.

3 Polynomials and deep classifiers

The proposed framework unifies recent deep classifiers through the lens of polynomials. We formalize the framework of polynomials below. Formally, let functions $\Phi_i^{[d]}(\hat{Z})$ define a linear (or multi-linear) function over \hat{Z} . The input variable \hat{Z} can either be a vector or a matrix, while d declares the degree and i the index of the function. Then, a polynomial of degree- N is expressed as:

$$\hat{Y} = \hat{\beta} + \Phi_1^{[1]}(\hat{Z}) + \Phi_1^{[2]}(\hat{Z})\Phi_2^{[2]}(\hat{Z}) + \dots + \underbrace{\Phi_1^{[N]}(\hat{Z}) \dots \Phi_{N-1}^{[N]} \Phi_N^{[N]}(\hat{Z})}_{N \text{ terms}} \quad (2)$$

where $\hat{\beta}$ is the constant term. Evidently, each additional degree introduces new parameters that can grow up exponentially with respect to the degree. However, we posit that this introduces a needed inductive bias on the method. In addition, if reduced parameters are required, this can be achieved by using low-rank assumptions or by sharing parameters across different degrees.

Using the formulation in (2), we exhibit how well-established methods can be formulated as polynomials. We present a taxonomy of classifiers based on the degree of the polynomial. In particular, we present first, second, third and higher-degree polynomial expansions. For modeling purposes, we focus on the core block of each architecture, while ignoring any activation functions.

First-degree polynomials include the majority of the feed-forward networks, such as AlexNet [35], VGG [53]. Specifically, the networks that include stacked linear operations (fully-connected or convolutional layers) but do not include any matrix or elementwise multiplication of the representations fit in this category. Such networks can be expressed in the form $\hat{Y} = C\hat{Z} + \hat{\beta}$, where the weight matrix C is a learnable parameter. A special case is ResNet [22]. The idea is to introduce shortcut connections that enable residual learning. Notation-wise this is a re-parametrization of the weight matrix C as $\hat{Y}_r = (I + C)\hat{Z} + \hat{\beta}$ where I is an identity matrix. Thus, $\Phi_1^{[1]}(\hat{Z}) = (I + C)\hat{Z}$.

Second-degree polynomials model self-interactions, i.e., they can selectively maximize the related inputs through second-order interactions. Often the interactions emerge in a particular dimension, e.g., correlations of the channels only, based on the particular application.

One special case of the second-degree polynomial is the *Squeeze-and-excitation networks (SENet)* [25]. The motivation lies in improving the channel-wise interactions, since there is already a strong inductive bias for the spatial interactions through convolutional layers. Notation-wise, the squeeze-and-excitation block is expressed as:

$$Y_s = (ZC_1) * r(p(ZC_1)C_2), \quad (3)$$

where $*$ denotes the Hadamard product, p is the global pooling function, r is a function that replicates the channels in the spatial dimensions and C_1, C_2 are weight matrices with learnable parameters. For simplicity, we assume the input is a matrix instead of a tensor, i.e., $Z \in \mathbb{R}^{hw \times c}$ with h the height, w the width and c the channels of the image. Then, Y_s expresses a second-degree term with $\Phi_1^{[2]}(Z) = ZC_1$ and $\Phi_2^{[2]}(Z) = \frac{1}{hw} \mathcal{I} \times_3 (C_2^T C_1^T Z^T \vec{\mathbf{1}})$ where \mathcal{I} is a third-order super-diagonal unit tensor and $\vec{\mathbf{1}}$ is a vector of ones.

The Squeeze-and-excitation block has been extended in the literature. The selective kernels networks [36] introduce a variant, where Z is replaced with the transformed $ZU_1 + ZU_2$. The learnable parameters U_1, U_2 include different receptive fields. However, the degree of the polynomial remains the same, i.e., second-degree.

The *Factorized Bilinear Model* [37] aim to extend the linear transformations of a convolutional layer by modeling the pairwise feature interactions. In particular, every output $y_{fb} \in \mathbb{R}$ is expressed as the following expansion of the input

$\mathbf{z} \in \mathbb{R}^d$:

$$y_{fb} = c_1 + \mathbf{c}_2^T \mathbf{z} + \mathbf{z}^T \mathbf{C}_3^T \mathbf{C}_3 \mathbf{z}, \quad (4)$$

where $c_1 \in \mathbb{R}$, $\mathbf{c}_2 \in \mathbb{R}^d$, $\mathbf{C}_3 \in \mathbb{R}^{k \times d}$ are learnable parameters with $k \in \mathbb{N}$.

Arguably, a more general form of second-degree polynomial expansion is proposed in *SORT* [61]. The idea is to combine different branches with multiplicative interactions. SORT is expressed as:

$$\mathbf{Y}_t = \mathbf{C}_1 \hat{\mathbf{Z}} + \mathbf{C}_2 \hat{\mathbf{Z}} + g((\mathbf{C}_1 \hat{\mathbf{Z}}) * (\mathbf{C}_2 \hat{\mathbf{Z}})), \quad (5)$$

where the $\mathbf{C}_1, \mathbf{C}_2$ are learnable parameters and g an elementwise, differentiable function. When \mathbf{C}_1 is identity, they g is the elementwise square root function.

Third-degree polynomials can encode long-range dependencies that are not captured by local operators, such as convolutions. A popular framework in this category is the *non-local neural networks (NL)* [60]. The non-local block can be expressed as:

$$\mathbf{Y}_n = (\mathbf{Z} \mathbf{C}_1 \mathbf{C}_2 \mathbf{Z}^T) \mathbf{Z} \mathbf{C}_3, \quad (6)$$

where the matrices \mathbf{C}_i for $i = 1, 2, 3$ are learnable parameters and $\mathbf{Z} \in \mathbb{R}^{hw \times c}$ is the input. The third-degree term is then $\Phi_1^{[3]}(\mathbf{Z}) = \mathbf{Z} \mathbf{C}_1$, $\Phi_2^{[3]}(\mathbf{Z}) = \mathbf{C}_2 \mathbf{Z}^T$ and $\Phi_3^{[3]}(\mathbf{Z}) = \mathbf{Z} \mathbf{C}_3$. Recently, *disentangled non-local networks (DNL)* [66] extends the formulation by including a second-degree term and the generated output is:

$$\mathbf{Y}_{dn} = ((\mathbf{Z} \mathbf{C}_1 - \boldsymbol{\mu}_q)(\mathbf{C}_2 \mathbf{Z} - \boldsymbol{\mu}_k)^T + \mathbf{Z} \mathbf{c}_4 \vec{\mathbf{1}}^T) \mathbf{Z} \mathbf{C}_3, \quad (7)$$

where \mathbf{c}_4 is a weight vector, $\boldsymbol{\mu}_q, \boldsymbol{\mu}_k$ are the mean vectors of the keys and queries representations and $\vec{\mathbf{1}} \in \mathbb{R}^{hw \times 1}$ a vector of ones. This translates to a new second-degree term with $\Phi_1^{[2]}(\mathbf{Z}) = \mathbf{Z} \mathbf{c}_4 \vec{\mathbf{1}}^T$ and $\Phi_2^{[2]}(\mathbf{Z}) = \mathbf{Z} \mathbf{C}_3$.

Higher-degree polynomials can (virtually) approximate any smooth functions. The Weierstrass theorem [54] and its extension [43] (pg 19) guarantee that any smooth function can be approximated by a higher-degree polynomial.

A recently proposed framework that leverages high-degree polynomials to approximate functions is *II-net* [8]. Each *II-net* block is a polynomial expansion with a pre-determined degree. The learnable coefficients of the polynomial are represented with higher-order tensors. One of the drawbacks of this method is that the order of the tensor increases linearly with the degree, hence the parameters explode exponentially. To mitigate this issue, a coupled tensor decomposition that allows for sharing among the coefficients is utilized. Thus, the number of parameters is reduced significantly. Although the method is formulated as a complete polynomial, the proposed sharing of the coefficients suppresses its expressive power in favour of model compression. The model used for the classification can be expressed with the following recursive relationship:

$$\mathbf{x}_n = \left(\mathbf{A}_{[n]}^T \mathbf{z} \right) * \left(\mathbf{S}_{[n]}^T \mathbf{x}_{n-1} + \mathbf{B}_{[n]}^T \mathbf{b}_{[n]} \right) + \mathbf{x}_{n-1} \quad (8)$$

for an N^{th} -degree expansion order with $n = 2, \dots, N$. The weight matrices $\mathbf{A}_{[n]}, \mathbf{S}_{[n]}, \mathbf{B}_{[n]}$ and the weight vector \mathbf{b} are learnable parameters, while $\mathbf{x}_1 =$

$(\mathbf{A}_{[1]}^T \mathbf{z}) * (\mathbf{B}_{[1]}^T \mathbf{b}_{[1]})$. The recursive equation can be used to express an arbitrary degree of expansion, while the weight matrices are shared across different degree terms. For instance, $\mathbf{A}_{[2]}$ is shared by both first and second-degree terms when $N = 2$.

From blocks to architecture: The core blocks of different architectures, and their polynomial counterparts, are analyzed above. The final network (in each case) is obtained by concatenating the respective blocks in a cascade. That is, the output of the first block is used as the input for the next block and so on. Each block expresses a polynomial expansion, thus, the final architecture expresses a product of polynomials.

4 Novel architectures based on the taxonomy

The taxonomy offers a new perspective on how to modify existing architectures in a principled way. We showcase how new architectures arise in a natural way by modifying the popular Non-local neural network and the recent Π -nets.

4.1 Higher-degree ResNet blocks

As shown in the previous section, the ResNet block is a first-degree polynomial. We can extend the polynomial expansion degree in order to enable higher-order correlations. A general N^{th} -degree polynomial is expressed as:

$$\mathbf{y} = \sum_{n=1}^N \left(\mathcal{W}^{[n]} \prod_{j=2}^{n+1} \times_j \mathbf{z} \right) + \beta \quad (9)$$

where $\mathbf{z} \in \mathbb{R}^\delta$, \times_m denotes the mode- m vector product, $\{\mathcal{W}^{[n]} \in \mathbb{R}^{\delta \times \prod_{m=1}^n \times_m \delta}\}_{n=1}^N$ are the tensor parameters. To reduce the learnable parameters, we assume a low-rank CP decomposition [32] on each tensor. By applying Lemma 1, we obtain:

$$\mathbf{y} = \beta + \mathbf{C}_{1,[1]}^T \mathbf{z} + \underbrace{\left(\mathbf{C}_{1,[2]}^T \mathbf{z} \right) * \left(\mathbf{C}_{2,[2]}^T \mathbf{z} \right) + \dots + \left(\mathbf{C}_{1,[N]}^T \mathbf{z} \right) * \dots * \left(\mathbf{C}_{N,[N]}^T \mathbf{z} \right)}_{N \text{ Hadamard products}} \quad (10)$$

where all \mathbf{C}_i are learnable parameters. Our proposed model in (10) is modular and can be designed for arbitrary polynomial degree. A schematic of (10) is depicted in Fig 3. The proposed model differentiates itself from (8) by not assuming any parameter sharing across the different degree terms.

4.2 Polynomial non-local blocks of different degrees

In this section, we demonstrate how the proposed taxonomy can be used to design a new architecture based on non-local blocks (NL). NL includes a third-degree

term, while disentangled non-local network (DNL) includes both a third-degree and a second-degree term.

We begin by including an additional first-degree term with learnable weights. The formed PDC- $NL^{[3]}$ block is:

$$\mathbf{Y}_{ours}^{[3]} = (\mathbf{Z}\mathbf{C}_1\mathbf{C}_2\mathbf{Z}^T)\mathbf{Z}\mathbf{C}_3 + \mathbf{Z}\mathbf{C}_4\mathbf{Z}\mathbf{C}_5 + \mathbf{Z}\mathbf{C}_6 \quad (11)$$

where \mathbf{C}_i are learnable parameters. In practice, a softmax activation function is added for the second and third degree factors, similar to the baseline. Besides the first-degree term, the proposed model differentiates itself from DNL by removing the sharing between the factors of third and second degree (i.e., \mathbf{C}_3 and \mathbf{C}_5) as well as utilizing a full factor matrix instead of a vector for the latter (i.e., \mathbf{C}_4). In our implementation the matrices \mathbf{C}_i matrices (for $i \neq 4$) compress the channels of the input by a factor of 4 for all models.

Building on (11), we propose to expand the PDC- $NL^{[3]}$ to a fourth degree polynomial expansion. We multiply the right hand side of (11) with a linear function with respect to the input \mathbf{Z} . That is, the PDC- $NL^{[4]}$ block is:

$$\mathbf{Y}_{ours}^{[4]} = \mathbf{Y}_{ours}^{[3]} + \mathbf{Y}_{ours}^{[3]} * r(p(\mathbf{Z}\mathbf{C}_7)\mathbf{C}_8), \quad (12)$$

where the term $r(p(\mathbf{Z}\mathbf{C}_7)\mathbf{C}_8)$ is similar to squeeze-and-excitation right hand side term. The new term captures the channel-wise correlations.

5 Experimental evaluation

In this section, we study how the degree of the polynomial expansion affects the expressive power of the model. Our goal is to illustrate how the taxonomy enables improvements in strong-performing models with minimal and intuitive modifications. This approach should also shed light on the inductive bias of each model, as well as the reasons behind the gap in performance and model compression. Unless mentioned otherwise, the degree of the polynomial for each block is considered the second-degree expansion. The proposed variants are referred to as PDC. Experiments with higher-degree polynomials, which are denoted as PDC $^{[k]}$ for k^{th} degree, are also conducted.

Training details: For fair comparison all the aforementioned experiments have the same optimization-related hyper-parameters, e.g., the layer initialization. The training is run for 120 epochs with batch size 128 and the SGD optimizer is used. The initial learning rate is 0.1, while the learning rate is multiplied with a factor of 0.1 in epochs 40, 60, 80, 100. Classification accuracy is utilized as the evaluation metric.

5.1 Image classification with residual blocks

The standard benchmarks of CIFAR10 [34], CIFAR100 [33] and ImageNet [50] are utilized to evaluate our framework on image classification. To highlight the

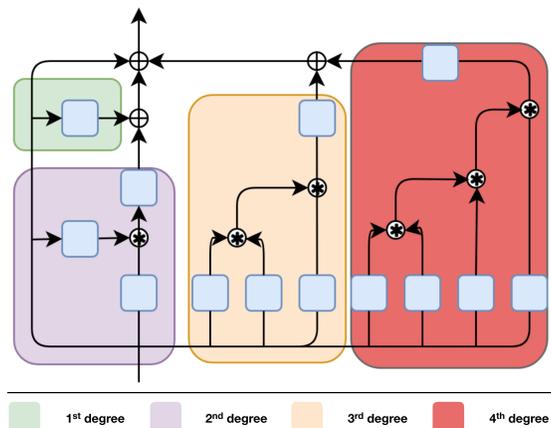


Fig. 3: Illustration of the modular nature of the propose polynomial expansion. Notice how we can trivially add new blocks in PDC to reach the pre-determined degree of the complete polynomial. The N^{th} -degree term includes $N + 1$ new layers (i.e., blue boxes), which might make it impractical to expand over the fourth-degree. In practice, higher-degree polynomials make the model more expressive and hence allow for reduced number of channels.

impact of the degree of the polynomial, we experiment with first (ResNet), second (SENet and \mathcal{H} -net^{||}) as well as higher-degree polynomials. PDC relies on eq. 10; that is, we do not assume shared factor matrices across the different degree terms.

The first experiment is conducted on CIFAR100. The ResNet18 and the respective SENet are the baselines, while the \mathcal{H} -net-ResNet is the main comparison method. To exhibit the efficacy of PDC, we implement various versions: (a) PDC-channels is the variant that retains the same channels as the original ResNet18, (b) PDC-param has approximate the same parameters as the corresponding baselines, (c) PDC-comp has the least parameters that can achieve a performance similar to the ResNet18 (which is achieved by reducing the channels), (d) the PDC which is the variant that includes both reduced parameters and increased performance with respect to ResNet18, (e) the higher-degree variants of PDC^[3] and PDC^[4] that modify the PDC such that the degree of each respective block is three and four respectively (see Fig. 3).

The accuracy of each method is reported in Table 2. SENet improves upon the baseline of ResNet18, which verifies the benefit of second-degree correlations in the channels. This is further demonstrated with \mathcal{H} -net-ResNet that captures second-degree information on all dimensions and can achieve the same performance with reduced parameters. Our method can further reduce the parameters by 60% over ResNet18 and achieve the same accuracy. This can be attributed

^{||}The default \mathcal{H} -net block is designed as second-degree polynomial. Higher-degree blocks are denoted with the symbol $\cdot^{[k]}$.

to the no-sharing scheme that enables more flexibility in the learned decompositions. The PDC-param further improves the accuracy by increasing the parameters, while the PDC-channels achieves the best accuracy by maintaining the same channels per residual block as ResNet18. The experiment is repeated with ResNet34 as the baseline. The accuracies in Table 9 demonstrate similar pattern, i.e., the parameters can be substantially reduced without sacrificing the accuracy. The results exhibit that the expressive power of PDC can allow it to be used both for compression and for improving the performance by maintaining the same number of parameters.

Table 2: Image classification on CIFAR100 with variants of ResNet18. The symbol ‘p’ denotes parameters.

Model	#p $\times 10^6$	Accuracy
ResNet18	11.2	0.756
SENet	11.6	0.760
<i>H</i> -net-ResNet	6.1	0.760
PDC-comp	4.3	0.760
PDC-channels	19.2	0.773
PDC-param	11.4	0.770
PDC	8.0	0.765
PDC ^[3]	16.8	0.766
PDC ^[4]	28.0	0.771

Table 3: Image classification on CIFAR10. The symbol ‘p’ denotes parameters.

Model	#p $\times 10^6$	Accuracy
ResNet18	11.2	0.945
SENet	11.5	0.946
<i>H</i> -net-ResNet	6.0	0.945
PDC	8.0	0.946
PDC-comp	4.3	0.945
ResNet34	21.3	0.948
<i>H</i> -net-ResNet	13.0	0.949
PDC	10.5	0.948

A similar experiment is conducted on CIFAR10. The same networks as above are used as baselines, i.e., ResNet18 and SENet. The methods *H*-net-ResNet and the PDC variants are also implemented as above. The results in Table 3 verify that both the *H*-net-ResNet and PDC-comp can compress substantially the number of parameters required while retaining the same accuracy. Importantly, our method achieves a parameter reduction of 28% over *H*-net-ResNet.

The last experiment is conducted on ImageNet [12], which contains 1.28M training images and 50K validation images from 1000 classes. We follow the standard settings to train deep networks on the training set and report the single-crop top-1 and the top-5 errors on the validation set. Our pre-processing and augmentation strategy follows the settings of the baseline (i.e., ResNet18). All models are trained for 100 epochs on 8 GPUs with 32 images per GPU (effective batch size of 256) with synchronous SGD of momentum 0.9. The learning rate is initialized to 0.1, and decays by a factor of 10 at the 30th, 60th, and 90th epochs. The results in Table 4 reflect the patterns that emerged above. The proposed PDC can reduce the number of parameters required while retaining the same accuracy. The variant of PDC that has similar number of parameters to the baseline ResNet18 can achieve an increase in the accuracy. More specifically, we set the stem channel number to 52 for PDC-ResNet18-cmp and 60 for

PDC18, respectively. In addition, we also decrease the channel to 1/4 before the Hadamard product in the proposed PDC block to save the computation cost.

Table 4: Image classification on ImageNet with variants of ResNet18.

Model	Top-1 Accuracy	Top-5 Accuracy	Flops	$\#\mathbf{p} \times 10^6$
ResNet18	0.698	0.891	1.82G	11.69M
SE-ResNet18	0.706	0.896	1.82G	11.78 M
<i>II</i> -net-ResNet	0.707	0.895	1.8207G	11.96M
PDC-cmp	0.698	0.893	1.30G	7.51M
PDC	0.710	0.899	1.67G	10.69M

5.2 Image classification with non-local blocks

In this section, an experimental comparison and validation is conducted with non-local blocks. The benchmark of CIFAR100 is selected, while ResNet18 is referred as the baseline. The original non-local network (NL) and the disentangled non-local (DNL) are the main compared methods from the literature. As a reminder we perform two extensions in DNL: i) we add a first-degree term (PDC- $NL^{[3]}$), ii) we add a fourth degree term (PDC- $NL^{[4]}$).

Table 5 contains the accuracy of each method. Notice that DNL improves upon the baseline NL, while PDC- $NL^{[3]}$ improves upon both DNL and NL. Interestingly, the fourth-degree variant, i.e., PDC- $NL^{[4]}$ outperforms all the compared methods by a considerable margin without increasing the number of parameters significantly. The results verify our intuition that the different polynomial terms enable additional flexibility to the model.

Besides the experiments on CIFAR100, we also test the proposed Non-local blocks in ImageNet. The training setup remains the same as in sec. 5.1. As shown in Table 6, the proposed Non-local blocks with different polynomial degrees significantly improve the top-1 accuracy on ImageNet while the total parameter numbers are lower than the baseline ResNet18. More specifically, we set the bottleneck ratio of 4 for the proposed PDC- $NL^{[3]}$ and PDC- $NL^{[4]}$, and we apply non-local blocks in multiple layers (c3+c4+c5) to better capture long-range dependency with only a slight increase in the computation cost. PDC- $NL^{[4]}$ significantly outperforms NL by 1.03%, showing the advantages of the polynomial information fusion.

5.3 Audio classification

Besides image classification, we conduct a series of experiments on audio classification to test the generalization of the polynomial expansion in different types of signals. The popular dataset of Speech Commands [62] is selected as our benchmark. The dataset consists of 60,000 audio files containing a single word each.

Table 5: Classification on CIFAR100 with non-local blocks. Table 6: Classification on ImageNet with non-local blocks.

Model	#p $\times 10^6$	Accuracy
ResNet18	11.2	0.756
NL	11.57	0.769
DNL	11.57	0.771
PDC-NL ^[3]	11.87	0.773
PDC-NL ^[4]	12.00	0.779

Model	#p $\times 10^6$	Accuracy
ResNet18	11.69	0.698
NL	12.02	0.702
PDC	10.69	0.710
PDC-NL ^[3]	11.35	0.712
PDC-NL ^[4]	11.51	0.716

The total number of words is 35, while there are at least 1,500 different files for each word. Every audio file is converted into a mel-spectrogram.

ResNet7 includes one residual block per group, while ResNet18 includes two residual blocks per group. The accuracy for each model (in both ResNet7 and ResNet18 comparisons) is reported in Table 7. All the compared methods have accuracy over 0.97, while the polynomial expansions of *II*-net-ResNet and PDC are able to reduce the number of parameters required to achieve the same accuracy, due to their expressiveness.

Table 7: Speech classification with ResNet variants. Four residual blocks are used in ResNet7 instead of eight of ResNet18. Nevertheless, the respective PDC7 can reduce even further the parameters to achieve the same performance.

Model	# param ($\times 10^6$)	Accuracy
ResNet7	4.9	0.974
SENet	5.1	0.974
PDC	3.9	0.975
ResNet18	11.2	0.977
SENet	11.5	0.977
<i>II</i> -net-ResNet	6.0	0.977
PDC	8.0	0.978

5.4 Image classification on long-tailed distributions

We scrutinize the performance of the proposed models on long-tailed image recognition. We convert CIFAR10, which has 5,000 number of samples per class, to a long-tailed version, called CIFAR10-LT. The imbalance factor (IF) is defined as the ratio of the largest class to the smallest class. The imbalance factor varies from 10 – 200, following similar benchmarks tailored to long-tailed distributions [11]. We note that the models are as defined above; the only change is the data distribution. The accuracy of each method is reported in Table 8. The results exhibit the benefits of the proposed models, which outperform the baselines.

Table 8: Accuracy on image classification on CIFAR10-LT. Each column corresponds to a different imbalance factor (IF).

Model \ IF	200	100	50	20	10
ResNet18	0.645	0.696	0.784	0.844	0.877
SENet	0.636	0.713	0.784	0.844	0.878
Π -net-ResNet	0.653	0.718	0.783	0.845	0.879
PDC-comp	0.653	0.727	0.786	0.848	0.882
PDC-param	0.665	0.726	0.792	0.851	0.886

6 Related work

Classification benchmarks [50] have a profound impact in the progress observed in machine learning. Such benchmarks been a significant testbed for new architectures [?, 22, 25, 35] and for introducing novel tasks, such as adversarial perturbations [57]. The architectures originally designed for classification have been applied to diverse tasks, such as image generation [3, 70].

ResNet [22] has been among the most influential works of the last few years which can be attributed both in its simplicity and its stellar performance. ResNet has been studied in both its theoretical capacity [2, 19, 51, 68] and its empirical performance [56, 64, 69]. A number of works focus on modifying the shortcut connection and the concatenation operation [7, 26, 59].

The Squeeze-and-Excitation network (SENet) [25] has been extended by [48] to capture second-degree correlations in both the spatial and the channel dimensions. [24] extend SENet by replacing the pooling operation with alternative operators that aggregate contextual information. [65] inspired by the SENet propose a gated convolution module. [49] improve SENet by introducing long-range dependencies. SENet has also been used as a drop-in module to improve the performance of residual blocks [14].

Non-local neural networks have been used extensively for capturing long-range interactions in both image-related tasks [28, 38, 71] and video-related tasks [6, 17]. Non-local networks are also related with self-attention [58] that is widely used in both vision [44] and natural language processing [13, 42]. [5] study when the long-range interactions emerge in non-local neural networks. They also frame a simplified non-local block, which in our terminology is a second-degree polynomial. Naturally, they observe that this resembles the SENet and merge their simplified non-local block and SENet block.

A promising line of research is that of polynomial activation functions. For instance, the element-wise quadratic activation function f^2 applied to a linear operation $C\hat{Z}$ is $(C\hat{Z})^2$. Both the theoretical [30, 40] and the empirical results [41, 45] support that polynomial activation functions can be beneficial. Our work is orthogonal to the polynomial activation functions, as they express a polynomial expansion of the representation, while we model a polynomial expansion of the input.

A well-established line of research is that of considering second or higher-order moments for various tasks. For instance, second-order statistics have been used for normalization methods [27], learning latent variable models [1]. However, our work focuses on classification methods that can be expressed as polynomials and not on the method of moments.

Tensors and tensor decompositions are related to our work [52]. Tensor decompositions such as the CP or the Tucker decompositions [32] are frequently used for model compression in computer vision. Tensor decompositions have also been used for modeling the components of deep neural networks. [10] interpret a whole convolutional network as a tensor decomposition, while the recent Einconv [20] focus on a single convolutional layer and model them with tensor decompositions. In our work the focus is not in the tensor decomposition used, but on the polynomial expansion that provides insights on the correlations that are captured by each model.

A line of research that is related to ours is that of multiplicative data fusion [?, 15, 16, 29, 31, 46, 67]. Even though multiplicative interactions can be considered as second-degree polynomials, data fusion of the aforementioned works is not our focus.

7 Conclusion

In this work, we study popular classification networks under the unifying perspective of polynomials. Notably, the popular ResNet, SENet and non-local networks are expressed as first, second and third degree polynomials respectively. The common framework provides insights on the inductive biases of each model and enables natural extensions building upon their polynomial nature. We conduct an extensive evaluation on image and audio classification benchmarks. We show how intuitive extensions to existing networks, e.g., converting the third-degree non-local network into a fourth degree, can improve the performance. Such natural extensions can be used for designing new architectures based on the proposed taxonomy. Importantly, our experimental evaluation highlights the dual utility of the polynomial framework: the networks can be used either for model compression or increased model performance. We expect this to be a significant feature when designing architectures for edge devices. Our experimentation in the presence of limited data and long-tailed data distributions highlights the benefits of the proposed taxonomy and provides a link to real-world applications, where massive data annotation is challenging.

References

1. Anandkumar, A., Ge, R., Hsu, D., Kakade, S.M., Telgarsky, M.: Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research* **15**, 2773–2832 (2014)
2. Balduzzi, D., Frean, M., Leary, L., Lewis, J., Ma, K.W.D., McWilliams, B.: The shattered gradients problem: If resnets are the answer, then what is the question? In: *International Conference on Machine Learning (ICML)*. pp. 342–350 (2017)
3. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: *International Conference on Learning Representations (ICLR)* (2019)
4. Cai, Z., Vasconcelos, N.: Cascade R-CNN: Delving into high quality object detection. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
5. Cao, Y., Xu, J., Lin, S., Wei, F., Hu, H.: Gcnet: Non-local networks meet squeeze-excitation networks and beyond. In: *International Conference on Computer Vision Workshops (ICCV'W)*. pp. 0–0 (2019)
6. Chen, Y., Kalantidis, Y., Li, J., Yan, S., Feng, J.: A²-nets: Double attention networks. In: *Advances in neural information processing systems (NeurIPS)*. pp. 352–361 (2018)
7. Chen, Y., Li, J., Xiao, H., Jin, X., Yan, S., Feng, J.: Dual path networks. In: *Advances in neural information processing systems (NeurIPS)*. pp. 4467–4475 (2017)
8. Chrysos, G., Moschoglou, S., Bouritsas, G., Panagakis, Y., Deng, J., Zafeiriou, S.: π -nets: Deep polynomial neural networks. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
9. Chrysos, G., Moschoglou, S., Panagakis, Y., Zafeiriou, S.: Polygan: High-order polynomial generators. *arXiv preprint arXiv:1908.06571* (2019)
10. Cohen, N., Shashua, A.: Convolutional rectifier networks as generalized tensor decompositions. In: *International Conference on Machine Learning (ICML)*. pp. 955–963. PMLR (2016)
11. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9268–9277 (2019)
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 248–255 (2009)
13. Galassi, A., Lippi, M., Torrioni, P.: Attention, please! a critical review of neural attention models in natural language processing. *arXiv preprint arXiv:1902.02181* (2019)
14. Gao, S., Cheng, M.M., Zhao, K., Zhang, X.Y., Yang, M.H., Torr, P.H.: Res2net: A new multi-scale backbone architecture. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)* (2019)
15. Georgopoulos, M., Chrysos, G., Pantic, M., Panagakis, Y.: Multilinear latent conditioning for generating unseen attribute combinations. *International Conference on Machine Learning (ICML)* (2020)
16. Georgopoulos, M., Oldfield, J., Nicolaou, M.A., Panagakis, Y., Pantic, M.: Mitigating demographic bias in facial datasets with style-based multi-attribute transfer. *International Journal of Computer Vision (IJCV)* (May 2021)
17. Girdhar, R., Carreira, J., Doersch, C., Zisserman, A.: Video action transformer network. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 244–253 (2019)

18. Hanin, B.: Universal function approximation by deep neural nets with bounded width and relu activations. *Mathematics* **7**(10), 992 (2019)
19. Hardt, M., Ma, T.: Identity matters in deep learning. In: *International Conference on Learning Representations (ICLR)* (2017)
20. Hayashi, K., Yamaguchi, T., Sugawara, Y., Maeda, S.i.: Einconv: Exploring unexplored tensor network decompositions for convolutional neural networks. In: *Advances in neural information processing systems (NeurIPS)* (2019)
21. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: *International Conference on Computer Vision (ICCV)* (2017)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
23. Hornik, K., Stinchcombe, M., White, H., et al.: Multilayer feedforward networks are universal approximators. *Neural networks* **2**(5), 359–366 (1989)
24. Hu, J., Shen, L., Albanie, S., Sun, G., Vedaldi, A.: Gather-excite: Exploiting feature context in convolutional neural networks. In: *Advances in neural information processing systems (NeurIPS)*. pp. 9401–9411 (2018)
25. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7132–7141 (2018)
26. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4700–4708 (2017)
27. Huang, L., Yang, D., Lang, B., Deng, J.: Decorrelated batch normalization. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 791–800 (2018)
28. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: *International Conference on Computer Vision (ICCV)*. pp. 603–612 (2019)
29. Jayakumar, S.M., Czarnecki, W.M., Menick, J., Schwarz, J., Rae, J., Osindero, S., Teh, Y.W., Harley, T., Pascanu, R.: Multiplicative interactions and where to find them. In: *International Conference on Learning Representations (ICLR)* (2020)
30. Kileel, J., Trager, M., Bruna, J.: On the expressive power of deep polynomial neural networks. In: *Advances in neural information processing systems (NeurIPS)* (2019)
31. Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. In: *Advances in neural information processing systems (NeurIPS)*. pp. 1564–1574 (2018)
32. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM review* **51**(3), 455–500 (2009)
33. Krizhevsky, A., Nair, V., Hinton, G.: Cifar-100 (canadian institute for advanced research) <http://www.cs.toronto.edu/~kriz/cifar.html>
34. Krizhevsky, A., Nair, V., Hinton, G.: The cifar-10 dataset. online: <http://www.cs.toronto.edu/~kriz/cifar.html> **55** (2014)
35. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems (NeurIPS)*. pp. 1097–1105 (2012)
36. Li, X., Wang, W., Hu, X., Yang, J.: Selective kernel networks. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 510–519 (2019)
37. Li, Y., Wang, N., Liu, J., Hou, X.: Factorized bilinear models for image recognition. In: *International Conference on Computer Vision (ICCV)*. pp. 2079–2087 (2017)
38. Li, Y., Jin, X., Mei, J., Lian, X., Yang, L., Xie, C., Yu, Q., Zhou, Y., Bai, S., Yuille, A.L.: Neural architecture search for lightweight non-local networks. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10297–10306 (2020)

39. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: European Conference on Computer Vision (ECCV) (2014)
40. Livni, R., Shalev-Shwartz, S., Shamir, O.: On the computational efficiency of training neural networks. In: Advances in neural information processing systems (NeurIPS). pp. 855–863 (2014)
41. Lokhande, V.S., Tasneeyapant, S., Venkatesh, A., Ravi, S.N., Singh, V.: Generating accurate pseudo-labels in semi-supervised learning and avoiding overconfident predictions via hermite polynomial activations. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11435–11443 (2020)
42. Ma, X., Zhang, P., Zhang, S., Duan, N., Hou, Y., Zhou, M., Song, D.: A tensorized transformer for language modeling. In: Advances in neural information processing systems (NeurIPS). pp. 2232–2242 (2019)
43. Nikol'skii, S.: Analysis III: Spaces of Differentiable Functions. Encyclopaedia of Mathematical Sciences, Springer Berlin Heidelberg (2013)
44. Parmar, N., Ramachandran, P., Vaswani, A., Bello, I., Levskaya, A., Shlens, J.: Stand-alone self-attention in vision models. In: Advances in neural information processing systems (NeurIPS). pp. 68–80 (2019)
45. Ramachandran, P., Zoph, B., Le, Q.V.: Searching for activation functions. arXiv preprint arXiv:1710.05941 (2017)
46. Reed, S., Sohn, K., Zhang, Y., Lee, H.: Learning to disentangle factors of variation with manifold interaction. In: International Conference on Machine Learning (ICML). pp. 1431–1439 (2014)
47. Rolnick, D., Tegmark, M.: The power of deeper networks for expressing natural functions. In: International Conference on Learning Representations (ICLR) (2018)
48. Roy, A.G., Navab, N., Wachinger, C.: Concurrent spatial and channel ‘squeeze & excitation’ in fully convolutional networks. In: International conference on medical image computing and computer-assisted intervention. pp. 421–429. Springer (2018)
49. Ruan, D., Wen, J., Zheng, N., Zheng, M.: Linear context transform block. In: AAAI. pp. 5553–5560 (2020)
50. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* **115**(3), 211–252 (2015)
51. Shamir, O.: Are resnets provably better than linear predictors? In: Advances in neural information processing systems (NeurIPS). pp. 507–516 (2018)
52. Sidiropoulos, N.D., De Lathauwer, L., Fu, X., Huang, K., Papalexakis, E.E., Faloutsos, C.: Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing* **65**(13), 3551–3582 (2017)
53. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (ICLR) (2015)
54. Stone, M.H.: The generalized weierstrass approximation theorem. *Mathematics Magazine* **21**(5), 237–254 (1948)
55. Styan, G.P.: Hadamard products and multivariate statistical analysis. *Linear algebra and its applications* **6**, 217–240 (1973)
56. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A.: Inception-v4, inception-resnet and the impact of residual connections on learning. In: AAAI Conference on Artificial Intelligence (2017)

57. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. In: International Conference on Learning Representations (ICLR) (2014)
58. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems (NeurIPS). pp. 5998–6008 (2017)
59. Wang, W., Li, X., Yang, J., Lu, T.: Mixed link networks. In: International Joint Conferences on Artificial Intelligence (IJCAI) (2018)
60. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7794–7803 (2018)
61. Wang, Y., Xie, L., Liu, C., Qiao, S., Zhang, Y., Zhang, W., Tian, Q., Yuille, A.: Sort: Second-order response transform for visual recognition. In: International Conference on Computer Vision (ICCV). pp. 1359–1368 (2017)
62. Warden, P.: Speech commands: A dataset for limited-vocabulary speech recognition. arXiv preprint arXiv:1804.03209 (2018)
63. Won, M., Chun, S., Serra, X.: Toward interpretable music tagging with self-attention. arXiv preprint arXiv:1906.04972 (2019)
64. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1492–1500 (2017)
65. Yang, Z., Zhu, L., Wu, Y., Yang, Y.: Gated channel transformation for visual recognition. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11794–11803 (2020)
66. Yin, M., Yao, Z., Cao, Y., Li, X., Zhang, Z., Lin, S., Hu, H.: Disentangled non-local neural networks. In: European Conference on Computer Vision (ECCV) (2020)
67. Yu, Z., Yu, J., Fan, J., Tao, D.: Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In: International Conference on Computer Vision (ICCV). pp. 1821–1830 (2017)
68. Zaeemzadeh, A., Rahnavard, N., Shah, M.: Norm-preservation: Why residual networks can become extremely deep? arXiv preprint arXiv:1805.07477 (2018)
69. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016)
70. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: International Conference on Machine Learning (ICML) (2019)
71. Zhu, Z., Xu, M., Bai, S., Huang, T., Bai, X.: Asymmetric non-local neural networks for semantic segmentation. In: International Conference on Computer Vision (ICCV). pp. 593–602 (2019)