

Drop an Octave: Reducing Spatial Redundancy in Convolutional Neural Networks with Octave Convolution

Yunpeng Chen^{†‡}, Haoqi Fan[†], Bing Xu[†], Zhicheng Yan[†], Yannis Kalantidis[†],
Marcus Rohrbach[†], Shuicheng Yan^{‡b}, Jiashi Feng[‡]
[†]Facebook AI, [‡]National University of Singapore, ^bYitu Technology

Abstract

In natural images, information is conveyed at different frequencies where higher frequencies are usually encoded with fine details and lower frequencies are usually encoded with global structures. Similarly, the output feature maps of a convolution layer can also be seen as a mixture of information at different frequencies. In this work, we propose to factorize the mixed feature maps by their frequencies, and design a novel Octave Convolution (OctConv) operation¹ to store and process feature maps that vary spatially “slower” at a lower spatial resolution reducing both memory and computation cost. Unlike existing multi-scale methods, OctConv is formulated as a single, generic, plug-and-play convolutional unit that can be used as a direct replacement of (vanilla) convolutions without any adjustments in the network architecture. It is also orthogonal and complementary to methods that suggest better topologies or reduce channel-wise redundancy like group or depth-wise convolutions. We experimentally show that by simply replacing convolutions with OctConv, we can consistently boost accuracy for both image and video recognition tasks, while reducing memory and computational cost. An OctConv-equipped ResNet-152 can achieve 82.9% top-1 classification accuracy on ImageNet with merely 22.2 GFLOPs.

1. Introduction

The efficiency of Convolutional Neural Networks (CNNs) keeps increasing with recent efforts to reduce the inherent redundancy in dense model parameters [15, 31, 42] and in the channel dimension of feature maps [47, 18, 6, 9]. However, substantial redundancy also exists in the spatial dimension of the feature maps produced by CNNs, where each location stores its own feature descriptor independently, while ignoring common information between adjacent locations that could be stored and processed together.

As shown in Figure 1(a), a natural image can be decom-




Figure 1: (a) Motivation. The spatial frequency model for vision [1, 10] shows that natural image can be decomposed into a low and a high spatial frequency part. (b) The output maps of a convolution layer can also be factorized and grouped by their spatial frequency. (c) The proposed multi-frequency feature representation stores the smoothly changing, low-frequency maps in a low-resolution tensor to reduce spatial redundancy. (d) The proposed Octave Convolution operates directly on this representation. It updates the information for each group and further enables information exchange between groups.

posed into a low spatial frequency component that describes the smoothly changing structure and a high spatial frequency component that describes the rapidly changing fine details [1, 10, 37, 39]. Similarly, we argue that the output feature maps of a convolution layer can also be decomposed into features of different spatial frequencies and propose a novel multi-frequency feature representation which stores high- and low-frequency feature maps into different groups as shown in Figure 1(b). Thus, the spatial resolution of the low-frequency group can be safely reduced by sharing information between neighboring locations to reduce spatial redundancy as shown in Figure 1(c). To accommodate the novel feature representation, we generalize the vanilla convolution, and propose *Octave Convolution (OctConv)* which

¹<https://github.com/facebookresearch/OctConv>

takes in feature maps containing tensors of two frequencies one octave apart, and extracts information directly from the low-frequency maps without the need of decoding it back to the high-frequency as shown in Figure 1(d). As a replacement of vanilla convolution, OctConv consumes substantially less memory and computational resources. In addition, OctConv processes low-frequency information with corresponding (low-frequency) convolutions and effectively enlarges the receptive field in the original pixel space and thus can improve recognition performance.

We design the OctConv in a generic way, making it a plug-and-play replacement for the vanilla convolution. Since OctConv mainly focuses on processing feature maps at multiple spatial frequencies and reducing their spatial redundancy, it is orthogonal and complementary to existing methods that focus on building better CNN topology [22, 41, 35, 33, 29], reducing channel-wise redundancy in convolutional feature maps [47, 9, 34, 32, 21] and reducing redundancy in dense model parameters [42, 15, 31]. Moreover, different from methods that exploit multi-scale information [4, 43, 12], OctConv can be easily deployed as a plug-and-play unit to replace convolution, without the need of changing network architectures or requiring hyper-parameters tuning. Compared to the closely related Multi-grid convolution [25], OctConv provides more insights on reducing the spatial redundancy in CNNs based on the frequency model and adopts more efficient inter-frequency information exchange strategy with better performance. We further integrate the OctConv into a wide variety of backbone architectures (including the ones featuring group, depth-wise, and 3D convolutions) and demonstrate universality of OctConv.

Our experiments demonstrate that by simply replacing the vanilla convolution with OctConv, we can consistently improve the performance of popular 2D CNN backbones including ResNet [16, 17], ResNeXt [47], DenseNet [22], MobileNet [18, 34] and SE-Net [19] on 2D image recognition on ImageNet [11], as well as 3D CNN backbones C2D [44] and I3D [44] on video action recognition on Kinetics [24, 3, 2]. The OctConv-equipped Oct-ResNet-152 can match or outperform state-of-the-art manually designed networks [32, 19] at lower memory and computational cost. Our contributions can be summarized as follows:

- We propose to factorize convolutional feature maps into two groups at different spatial frequencies and process them with different convolutions at their corresponding frequency, one octave apart. As the resolution for low frequency maps can be reduced, this saves both storage and computation. This also helps each layer gain a larger receptive field to capture more contextual information.
- We design a plug-and-play operation named OctConv to replace the vanilla convolution for operating on the new feature representation directly and reducing spatial re-

dundancy. Importantly, OctConv is fast in practice and achieves a speedup close to the theoretical limit.

- We extensively study the properties of the proposed OctConv on a variety of backbone CNNs for image and video tasks and achieve significant performance gain even comparable to the best AutoML networks.

2. Related Work

Improving the efficiency of CNNs. Ever since the pioneering work on AlexNet [26] and VGG [35], researchers have made substantial efforts to improve the efficiency of CNNs. ResNet [16, 17] and DenseNet [22] improve the network topology by adding shortcut connections to early layers. ResNeXt [47] and ShuffleNet [49] use sparsely connected group convolutions to reduce redundancy in inter-channel connectivity. Xception [9] and MobileNet [18, 34] adopt depth-wise convolutions that further reduce the connection density. Meanwhile, NAS [51], PNAS [29] and AmoebaNet [33] propose to atomically find the best network topology for a given task. Pruning methods, such as DSD [15] and ThiNet [31], focus on reducing the redundancy in the model parameters by eliminating the least significant weight or connections in CNNs. Besides, HetConv [36] propose to replace the vanilla convolution filters with heterogeneous convolution filters that are in different sizes. However, all of these methods ignore the redundancy on the spatial dimension of feature maps, which is addressed by the proposed OctConv, making OctConv orthogonal and complementary to these previous methods. Noticeably, OctConv does not change the connectivity between feature maps, making it also different from inception-like multi-path designs [41, 40, 47].

Multi-scale Representation Learning. Prior to the success of deep learning, multi-scale representation has long been applied for local feature extraction, such as the SIFT features [30]. In the deep learning era, multi-scale representation also plays a important role due to its strong robustness and generalization ability. FPN [27] and PSP [50] merge convolutional features from different depths at the end of the networks for object detection and segmentation tasks. MSDNet [20] and HR-Nets [38], proposed carefully designed network architectures that contain multiple branches where each branch has its own spatial resolution. The bL-Net [4] and ELASTIC-Net [43] adopt similar idea, but are designed as a replacement of residual block for ResNet [16, 17] and thus are more flexible and easier to use. But extra expertise and hyper-parameter tuning are still required when adopt them to architectures beyond ResNet, such as MobileNetV1 [18], DenseNet [22]. Multi-grid CNNs [25] propose a multi-grid pyramid feature representation and define the MG-Conv operator as a replacement of convolution operator, which is conceptually similar to our method but is motivated for exploiting multi-

scale features. Compared with MG-Conv, OctConv adopts more efficient design to exchange inter-frequency information with higher performance as can be found in Sec. 3.3 and Sec. 4.3. For video models, the recently proposed Slow-Fast Networks [12] introduce multi-scale pathways on the temporal dimension. As we show in Section 4.4, this is complementary to OctConv which operates on the spatial dimensions.

In a nutshell, OctConv focuses on reducing the spatial redundancy in CNNs and is designed to replace vanilla convolution operations without needing to adjust backbone CNN architecture. We extensively compare OctConv to closely related methods in the sections of method and experiment and show that OctConv CNNs give top results on a number of challenging benchmarks.

3. Method

In this section, we first introduce the octave feature representation and then describe Octave Convolution, which operates directly on it. We also discuss implementation details and show how to integrate OctConv into group and depth-wise convolution architectures.

3.1. Octave Feature Representation

For the vanilla convolution, all input and output feature maps have the same spatial resolution, which may not be necessary since some of the feature maps may represent low-frequency information which is spatially redundant and can be further compressed as illustrated in Figure 1.

To reduce the spatial redundancy, we introduce the *octave feature representation* that explicitly factorizes the feature map tensors into groups corresponding to low and high frequencies. The scale-space theory [28] provides us with a principled way of creating scale-spaces of spatial resolutions, and defines an *octave* as a division of the spatial dimensions by a power of 2 (we only explore 2^1 in this work). We follow this fashion and reduce the spatial resolution of the low-frequency feature maps by an octave.

Formally, let $X \in \mathbb{R}^{c \times h \times w}$ denote the input feature tensor of a convolutional layer, where h and w denote the spatial dimensions and c the number of feature maps or channels. We explicitly factorize X along the channel dimension into $X = \{X^H, X^L\}$, where the high-frequency feature maps $X^H \in \mathbb{R}^{(1-\alpha)c \times h \times w}$ capture fine details and the low-frequency maps $X^L \in \mathbb{R}^{\alpha c \times \frac{h}{2} \times \frac{w}{2}}$ vary slower in the spatial dimensions (w.r.t. the image locations). Here $\alpha \in [0, 1]$ denotes the *ratio* of channels allocated to the low-frequency part and the low-frequency feature maps are defined *an octave lower* than the high frequency ones, *i.e.* at half of the spatial resolution as shown in Figure 1(c).

In the next subsection, we introduce a convolution operator that operates directly on this multi-frequency feature representation and name it *Octave Convolution (OctConv)*.

3.2. Octave Convolution

The octave feature representation presented in Section 3.1 reduces the spatial redundancy and is more compact than the original representation. However, the vanilla convolution cannot directly operate on such a representation, due to differences in spatial resolution in the input features. A naive way of circumventing this is to up-sample the low-frequency part X^L to the original spatial resolution, concatenate it with X^H and then convolve, which would lead to extra costs in computation and memory and diminish all the savings from the compression. In order to fully exploit our compact multi-frequency feature representation, we introduce Octave Convolution, which can directly operate on factorized tensors $X = \{X^H, X^L\}$ without requiring any extra computational or memory overhead.


Vanilla Convolution. Let $W \in \mathbb{R}^{c \times k \times k}$ denote a $k \times k$ convolution kernel and $X, Y \in \mathbb{R}^{c \times h \times w}$ denote the input and output tensors, respectively. Each feature map in $Y_{p,q} \in \mathbb{R}^c$ can be computed by

$$Y_{p,q} = \sum_{i,j \in \mathcal{N}_k} W_{i+\frac{k-1}{2}, j+\frac{k-1}{2}}^\top X_{p+i, q+j}, \quad (1)$$


where (p, q) denotes the location coordinate and $\mathcal{N}_k = \{(i, j) : i = \{-\frac{k-1}{2}, \dots, \frac{k-1}{2}\}, j = \{-\frac{k-1}{2}, \dots, \frac{k-1}{2}\}\}$ defines a local neighborhood. For simplicity, in all equations we omit the padding, we assume k is an odd number and that the input and output data have the same dimensionality, *i.e.* $c_{in} = c_{out} = c$.

Octave Convolution. The goal of our design is to effectively process the low and high frequency in their corresponding frequency tensor but also enable efficient inter-frequency communication. Let X, Y be the factorized input and output tensors. Then the high- and low-frequency feature maps of the output $Y = \{Y^H, Y^L\}$ will be given by $Y^H = Y^{H \rightarrow H} + Y^{L \rightarrow H}$ and $Y^L = Y^{L \rightarrow L} + Y^{H \rightarrow L}$, respectively, where $Y^{A \rightarrow B}$ denotes the convolutional update from feature map group A to group B . Specifically, $Y^{H \rightarrow H}, Y^{L \rightarrow L}$ denote intra-frequency update, while $Y^{H \rightarrow L}, Y^{L \rightarrow H}$ denote inter-frequency communication.

To compute these terms, we split the convolutional kernel W into two components $W = [W^H, W^L]$ responsible for convolving with X^H and X^L respectively. Each component can be further divided into intra- and inter-frequency part: $W^H = [W^{H \rightarrow H}, W^{L \rightarrow H}]$ and $W^L = [W^{L \rightarrow L}, W^{H \rightarrow L}]$ with the parameter tensor shape shown in Figure 2(b). Specifically for high-frequency feature map, we compute it at location (p, q) by using a regular convolution for the intra-frequency update, and for the inter-frequency communication we can fold the up-sampling over the feature tensor X^L into the convolution, removing the need of explicitly computing and storing the up-sampled



(a) Detailed design of the Octave Convolution. Green arrows correspond to information updates while red arrows facilitate information exchange between the two frequencies.



(b) The Octave Convolution kernel. The $k \times k$ Octave Convolution kernel $W \in \mathbb{R}^{c_{in} \times c_{out} \times k \times k}$ is equivalent to the vanilla convolution kernel in the sense that the two have the exact same number of parameters.

Figure 2: Octave Convolution. We set $\alpha_{in} = \alpha_{out} = \alpha$ throughout the network, apart from the first and last OctConv of the network where $\alpha_{in} = 0, \alpha_{out} = \alpha$ and $\alpha_{in} = \alpha, \alpha_{out} = 0$, respectively.

feature maps as follows:

$$\begin{aligned}
 Y_{p,q}^H &= Y_{p,q}^{H \rightarrow H} + Y_{p,q}^{L \rightarrow H} \\
 &= \sum_{i,j \in \mathcal{N}_k} W_{i+\frac{k-1}{2}, j+\frac{k-1}{2}}^{H \rightarrow H} \top X_{p+i, q+j}^H \\
 &\quad + \sum_{i,j \in \mathcal{N}_k} W_{i+\frac{k-1}{2}, j+\frac{k-1}{2}}^{L \rightarrow H} \top X_{(\lfloor \frac{p}{2} \rfloor + i), (\lfloor \frac{q}{2} \rfloor + j)}^L,
 \end{aligned} \tag{2}$$

where $\lfloor \cdot \rfloor$ denotes the floor operation. Similarly, for the low-frequency feature map, we compute the intra-frequency update using a regular convolution. Note that, as the map is in one octave lower, the convolution is also low-frequency w.r.t. the high-frequency coordinate space. For the inter-frequency communication we can again fold the down-sampling of the feature tensor X^H into the convolution as follows:

$$\begin{aligned}
 Y_{p,q}^L &= Y_{p,q}^{L \rightarrow L} + Y_{p,q}^{H \rightarrow L} \\
 &= \sum_{i,j \in \mathcal{N}_k} W_{i+\frac{k-1}{2}, j+\frac{k-1}{2}}^{L \rightarrow L} \top X_{p+i, q+j}^L \\
 &\quad + \sum_{i,j \in \mathcal{N}_k} W_{i+\frac{k-1}{2}, j+\frac{k-1}{2}}^{H \rightarrow L} \top X_{(2*p+0.5+i), (2*q+0.5+j)}^H,
 \end{aligned} \tag{3}$$

where multiplying a factor 2 to the locations (p, q) performs down-sampling, and further shifting the location by half step is to ensure the down-sampled maps well aligned with the input. However, since the index of X^H can only be an integer, we could either round the index to $(2*p+i, 2*q+j)$ or approximate the value at $(2*p+0.5+i, 2*q+0.5+j)$ by averaging all 4 adjacent locations. The first one is also known as strided convolution and the second one as average pooling. As we discuss in Section 3.3 and Fig. ??, strided convolution leads to misalignment; we therefore use average pooling to approximate this value for the rest of the paper.

An interesting and useful property of the Octave Convolution is the larger receptive field for the low-frequency feature maps. Convoluting the low-frequency part X^L with

$k \times k$ convolution kernels, results in an effective enlargement of the receptive field by a factor of 2 compared to vanilla convolutions. This further helps each OctConv layer capture more contextual information from distant locations and can potentially improve recognition performance.

3.3. Implementation Details

As discussed in the previous subsection, the index $\{(2*p+0.5+i), (2*q+0.5+j)\}$ has to be an integer for Eq. 3. Instead of rounding it to $\{(2*p+i), (2*q+j)\}$, i.e. conduct convolution with stride 2 for down-sampling, we adopt average pooling to get more accurate approximation. This helps alleviate misalignments that appear when aggregating information from different scales, as shown in Appendix A. and Appendix C.. We can now rewrite the output $Y = \{Y^H, Y^L\}$ of the Octave Convolution using average pooling for down-sampling as:

$$\begin{aligned}
 Y^H &= f(X^H; W^{H \rightarrow H}) + \text{upsample}(f(X^L; W^{L \rightarrow H}), 2) \\
 Y^L &= f(X^L; W^{L \rightarrow L}) + f(\text{pool}(X^H, 2); W^{H \rightarrow L}),
 \end{aligned} \tag{4}$$

where $f(X; W)$ denotes a convolution with parameters W , $\text{pool}(X, k)$ is an average pooling operation with kernel size $k \times k$ and stride k . $\text{upsample}(X, k)$ is an up-sampling operation by a factor of k via nearest interpolation.

The details of the OctConv operator implementation are shown in Figure 2. It consists of four computation paths that correspond to the four terms in Eq. (4): two green paths correspond to information updating for the high- and low-frequency feature maps, and two red paths facilitate information exchange between the two octaves.

Group and Depth-wise convolutions. The Octave Convolution can also be adopted to other popular variants of the vanilla convolution such as group [47] or depth-wise [18] convolutions. For the group convolution case, we simply set all four convolution operations that appear inside the design of the OctConv to group convolutions. Similarly, for the depth-wise convolution case, the convolution opera-

ratio (α)	.0	.125	.25	.50	.75	.875	1.0
#FLOPs Cost	100%	82%	67%	44%	30%	26%	25%
Memory Cost	100%	91%	81%	63%	44%	35%	25%

Table 1: Relative theoretical gains for the proposed multi-frequency feature representation over vanilla feature maps for varying choices of the ratio α of channels used by the low-frequency feature. When $\alpha = 0$, no low-frequency feature is used which is the case of vanilla convolution.

tions are depth-wise and therefore the information exchange paths are eliminated, leaving only two depth-wise convolution operations. We note that both group OctConv and depth-wise OctConv reduce to their respective vanilla versions if we do not compress the low-frequency part.

Efficiency analysis. Table 1 shows the theoretical computational cost and memory consumption of OctConv over the vanilla convolution and vanilla feature map representation. More information on deriving the theoretical gains presented in Table 1 can be found in the supplementary material. We note the theoretical gains are calculated per convolutional layer. In Section 4 we present the corresponding practical gains on real scenarios and show that our OctConv implementation can sufficiently approximate the theoretical numbers.

Integrating OctConv into backbone networks. OctConv is backwards compatible with vanilla convolution and can be inserted to regular convolutional networks without special adjustment. To convert a vanilla feature representation to a multi-frequency feature representation, *i.e.* at the first OctConv layer, we set $\alpha_{in} = 0$ and $\alpha_{out} = \alpha$. In this case, OctConv paths related to the low-frequency input is disabled, resulting in a simplified version which only has two paths. To convert the multi-frequency feature representation back to vanilla feature representation, *i.e.* at the last OctConv layer, we set $\alpha_{out} = 0$. In this case, OctConv paths related to the low-frequency output is disabled, resulting in a single full resolution output.

Comparison to Multi-grid Convolution [25]. The multi-grid conv (MG-Conv) [25] is a bi-directional and cross-scale convolution operator. Though being conceptually similar, our OctConv is different from MG-Conv in both the core motivation and design. MG-Conv aims to exploit multi-scale information in existing CNNs, while OctConv is focusing on reducing spatial redundancy among neighborhood pixels. In terms of design, MG-Conv adopts max-pooling for down-sampling. This requires extra memory for storing the index of the maximum value during training and further yields lower accuracy (see Appendix C.). MG-Conv also first up-samples and then convolves with the enlarged feature maps. Differently, OctConv aims for reducing spatial redundancy and is a naive extension of convolution op-

erator. It uses average pooling to distill low-frequency features without extra memory cost and its upsampling operation follows the convolution, and is thus more efficient than MG-Conv. The meticulous design of the lateral paths are essential for OctConv to be much more memory and computationally efficient than MG-Conv and improve accuracy without increasing the network complexity. We compare OctConv to MG-Conv in Table 4.

4. Experimental Evaluation

In this section, we validate the effectiveness and efficiency of the proposed Octave Convolution for both 2D and 3D networks. We first present ablation studies for image classification on ImageNet [11] and then compare it with the state-of-the-art. Then, we show the proposed OctConv also works in 3D CNNs using Kinetics-400 [24, 3] and Kinetics-600 [2] datasets. The best results per category/block are highlighted in bold font throughout the paper.

4.1. Experimental Setups

Image classification. We examine OctConv on a set of most popular CNNs [18, 34, 16, 17, 22, 47, 19] by replacing the regular convolutions with OctConv (except the first convolutional layer before the max pooling). The resulting networks only have one global hyper-parameter α , which denotes the ratio of low frequency part. We do apple-to-apple comparison and reproduce all baseline methods by ourselves under the same training/testing setting for internal ablation studies. All networks are trained with naive softmax cross entropy loss except that the MobileNetV2 also adopts the label smoothing [40], and the best ResNet-152 adopts both label smoothing and mixup [48] to prevent overfitting. Same as [4], all networks are trained from scratch and optimized by SGD with cosine learning rate [13]. Standard accuracy of single central crop [16, 17, 47, 4, 43] on validation set is reported.

Video action recognition. We use both Kinetics-400 [24, 3] and Kinetics-600 [2] for human action recognition. We choose standard baseline backbones from Inflated 3D ConvNet [44] and compare them with the OctConv counterparts. We follow the setting from [45] using frame length of 8 as standard input size, training 300k iterations in total, and averaging the predictions over 30 crops during inference time. To make fair comparison, we report the performance of the baseline and OctConv under precisely the same settings.

4.2. Ablation Study on ImageNet

We conduct a series of ablation studies aiming to answer the following questions: 1) *Does OctConv have better FLOPs-Accuracy trade-off than vanilla convolution?* 2) *In which situation does the OctConv work the best?*




Figure 3: Ablation study results on ImageNet. OctConv-equipped models are more efficient and accurate than baseline models. Markers in black in each line denote the corresponding baseline models without OctConv. The colored numbers are the ratio α . Numbers in X axis denote FLOPs in logarithmic scale.

Results on ResNet-50. We begin with using the popular ResNet-50 [17] as the baseline CNN and replacing the regular convolution with our proposed OctConv to examine the flops-accuracy trade-off. In particular, we vary the global ratio $\alpha \in \{0.125, 0.25, 0.5, 0.75\}$ to compare the image classification accuracy versus computational cost (i.e. FLOPs) [16, 17, 47, 7] with the baseline. The results are shown in Figure 3 in pink.

We make following observations. 1) The flops-accuracy trade-off curve is a concave curve, where the accuracy first rises up and then slowly goes down. 2) We can see two sweet spots: The first at $\alpha = 0.5$, where the network gets similar or better results even when the FLOPs are reduced by about half; the second at $\alpha = 0.125$, where the network reaches its best accuracy, 1.2% higher than baseline (black circle). We attribute the increase in accuracy to OctConv’s effective design of multi-frequency processing and the corresponding enlarged receptive field which provides more contextual information to the network. While reaching the accuracy peak at 0.125, the accuracy does not suddenly drop but decreases slowly for higher ratios α , indicating reducing the resolution of the low frequency part does not lead to significant information loss. Interestingly, 75% of the feature maps can be compressed to half the resolution with only 0.3% accuracy drop, which demonstrates effectiveness of grouping and compressing the smoothly changed feature maps for reducing the spatial redundancy in CNNs. In Table 2 we demonstrate the theoretical FLOPs saving of OctConv is also reflected in the actual CPU inference time in practice. For ResNet-50, we are close to obtaining theoretical FLOPs speed up. These results indicate OctConv is able to deliver important practical benefits, rather than only saving FLOPs in theory.

ratio (α)	Top-1 (%)	#FLOPs (G)	Inference Time (ms)	Backend
N/A	77.0	4.1	119	MKLDNN
N/A	77.0	4.1	115	TVM
.125	78.2	3.6	116	TVM
.25	78.0	3.1	99	TVM
.5	77.4	2.4	74	TVM
.75	76.7	1.9	61	TVM

Table 2: Results of ResNet-50. Inference time is measured on Intel Skylake CPU at 2.0 GHz (single thread). We report Intel(R) Math Kernel Library for Deep Neural Networks v0.18.1 (MKLDNN) [23] inference time for vanilla ResNet-50. Because vanilla ResNet-50 is well optimized by Intel, we also show MKLDNN results as additional performance baseline. OctConv networks are compiled by TVM [5] v0.5.




Figure 4: Frequency analysis for activation maps in different groups. ‘Baseline’ refers to vanilla ResNet. 10k activation maps are sampled from ResNet-101(Res3).

Results on more CNNs. To further examine if the proposed OctConv works for other networks with different depth/wide/topology, we select the currently most popular networks as baselines and repeat the same ablation study. These networks are ResNet-(26;50;101;200) [17], ResNeXt-(50,32x4d;101,32x4d) [47], DenseNet-121 [22] and SE-ResNet-50 [19]. The ResNeXt is chosen for assessing the OctConv on group convolution, while the SE-Net [19] is used to check if the gain of SE block found on vanilla convolution based networks can also be seen on OctConv. As shown in Figure 3, OctConv equipped networks for different architecture behave similarly to the Oct-ResNet-50, where the FLOPs-Accuracy trade-off is in a concave curve and the performance peak also appears at ratio $\alpha = 0.125$ or $\alpha = 0.25$. The consistent performance gain on a variety of backbone CNNs confirms that OctConv is a good replacement of vanilla convolution.

Frequency Analysis. Figure 4 shows the frequency analysis results. We conducted the Fourier transform for each group of feature maps and visualized the averaged results. From the energy map, the low frequency group does not contain high frequency signal, while the high frequency group contains both low and high frequency signals. This confirms that low-frequency group indeed captures low-frequency information as expected. Note that OctConv gives the high frequency group the flexibly to store both low and high frequency signals for better learning capacity.

Summary. 1) OctConv can help CNNs improve the accuracy while decreasing the FLOPs, deviating from other methods that reduce the FLOPs with a cost of lower accuracy. 2) At test time, the gain of OctConv over baseline

Method	ratio (α)	#Params (M)	#FLOPs (M)	CPU (ms)	Top-1 (%)
CondenseNet ($G = C = 8$) [21]	-	2.9	274	-	71.0
1.5 ShuffleNet (v1) [49]	-	3.4	292	-	71.5
1.5 ShuffleNet (v2) [32]	-	3.5	299	-	72.6
0.75 MobileNet (v1) [18]	-	2.6	325	13.4	70.3*
0.75 Oct-MobileNet (v1) (ours)	.375	2.6	213	11.9	70.5
1.0 Oct-MobileNet (v1) (ours)	.5	4.2	321	18.4	72.5
1.0 MobileNet (v2) [34]	-	3.5	300	24.5	72.0
1.0 Oct-MobileNet (v2) (ours)	.375	3.5	256	17.1	72.0
1.125 Oct-MobileNet (v2) (ours)	.5	4.2	295	26.3	73.0

Table 3: ImageNet classification results for *Small* models. * indicates it is better than original reproduced by MXNet GluonCV v0.4 [14]. The inference speed is tested using TVM on Intel Skylake processor (2.0GHz, single thread)².

Method	ratio (α)	Depth	#Params (M)	#FLOPs (G)	Top-1 (%)
R-MG-34 [25]	-	34	32.9	5.8	75.5
Oct-ResNet-26 (ours)	.25	26	16.0	1.9	76.1
Oct-ResNet-50 (ours)	.5	50	25.6	2.4	77.4
ResNet-50 + GloRe [8] (+3 blocks Res4)	-	50	30.5	5.2	78.4
Oct-ResNet-50 (ours) + GloRe [8] (+3 blocks Res4)	.5	50	30.5	3.1	78.8
ResNeXt-50 + Elastic [43]	-	50	25.2	4.2	78.4
Oct-ResNeXt-50 (32 \times 4d) (ours)	.25	50	25.0	3.2	78.8
ResNeXt-101 + Elastic [43]	-	101	44.3	7.9	79.2
Oct-ResNeXt-101 (32 \times 4d) (ours)	.25	101	44.2	5.7	79.6
bL-ResNet-50 [‡] ($\alpha = 4, \beta = 4$) [4]	-	50 (+3)	26.2	2.5	76.9
Oct-ResNet-50 [‡] (ours)	.5	50 (+3)	25.6	2.5	77.8
Oct-ResNet-50 (ours)	.5	50	25.6	2.4	77.4
bL-ResNeXt-50 [‡] (32 \times 4d) [4]	-	50 (+3)	26.2	3.0	78.4
Oct-ResNeXt-50 [‡] (32 \times 4d) (ours)	.5	50 (+3)	25.1	2.7	78.6
Oct-ResNeXt-50 (32 \times 4d) (ours)	.5	50	25.0	2.4	78.4
bL-ResNeXt-101 [‡] (32 \times 4d) [4]	-	101 (+1)	43.4	4.1	78.9
Oct-ResNeXt-101 [‡] (32 \times 4d) (ours)	.5	101 (+1)	40.1	4.2	79.4
Oct-ResNeXt-101 [‡] (32 \times 4d) (ours)	.5	101 (+1)	44.2	4.2	79.1
Oct-ResNeXt-101 (32 \times 4d) (ours)	.5	101	44.2	4.0	78.9

Table 4: ImageNet Classification results for *Middle* sized models. ‡ refers to method that replaces “Max Pooling” by extra convolution layer(s) [4]. § refers to method that uses balanced residual block distribution [4].

models increases as the test image resolution grows because OctConv can detect large objects better due to its larger receptive field, see Appendix C. 3) Both the information exchanging paths are important, since removing any of them can lead to accuracy drop, see Appendix C. 4) Shallow networks, *e.g.* ResNet-26, have a rather limited receptive field, and can especially benefit from OctConv, which greatly enlarges their receptive field.

4.3. Comparing with SOTAs on ImageNet

Small models. We adopt the most popular light weight networks as baselines and examine if OctConv works well on these compact networks with depth-wise convolution. In particular, we use the “0.75 MobileNet (v1)” [18] and “1.0 MobileNet (v2)” [34] as baseline and replace the regular convolution with our proposed OctConv. The results are shown in Table 3. We find that OctConv can reduce the FLOPs of MobileNetV1 by 34%, and provide better accuracy and faster speed in practice; it is able to reduce the FLOPs of MobileNetV2 by 15%, achieving the same accuracy with faster speed. When the computation budget is

fixed, one can adopt wider models to increase the learning capacity because OctConv can compensate the extra computation cost. In particular, our OctConv equipped networks achieve 2% improvement on MobileNetV1 under the same FLOPs and 1% improvement on MobileNetV2.

Medium models. In the above experiment, we have compared and shown that OctConv is complementary with a set of state-of-the-art CNNs [16, 17, 47, 22, 18, 34, 19]. In this part, we compare OctConv with MG-Conv [25], GloRe [8], Elastic [43] and bL-Net [4] which share a similar idea as our method. Seven groups of results are shown in Table 4. In group 1, our Oct-ResNet-26 shows 0.6% better accuracy than R-MG-34 while costing only one third of FLOPs and half of #Params. Also, our Oct-ResNet-50, which costs less than half of FLOPs, achieves 1.9% higher accuracy than R-MG-34. In group 2, adding our OctConv to GloRe network reduces the FLOPs with better accuracy. In group 3, our Oct-ResNeXt-50 achieves better accuracy than the Elastic [43] based method (78.8% v.s. 78.4%) while reducing the computational cost by 31%. In group 4, the Oct-ResNeXt-101 also achieves higher accuracy than the Elastic based method (79.6% v.s. 79.2%) while costing 38% less computation. When compared to the bL-Net [4], OctConv equipped methods achieve better FLOPs-Accuracy trade-off without bells and tricks. When adopting the tricks used in the baseline bL-Net [4], our Oct-ResNet-50 achieves 0.9% higher accuracy than bL-ResNet-50 under the same computational budget (group 5), and Oct-ResNeXt-50 (group 6) and Oct-ResNeXt-101 (group 7) get better accuracy under comparable or even lower computational budget. This is because MG-Conv [25], Elastic-Net [43] and bL-Net [4] are designed following the principle of introducing multi-scale features without considering reducing the spatial redundancy. In contrast, OctConv is born for solving the high spatial redundancy problem in CNNs, uses more efficient strategies to store and process the information throughout the network, and can thus achieve better efficiency and performance.

Large models. Table 5 shows the results of OctConv in large models. Here, we choose the ResNet-152 as the backbone CNN, replacing the first 7×7 convolution by three 3×3 convolution layers and removing the max pooling by a lightweight residual block [4]. We report results for Oct-ResNet-152 with and without the SE-block [19]. As can be seen, our Oct-ResNet-152 achieves accuracy comparable to the best manually designed networks with less FLOPs (10.9G v.s. 12.7G). Since our model does not use group or depth-wise convolutions, it also requires significantly less GPU memory, and runs faster in practice compared to the SE-ShuffleNet v2-164 and AmoebaNet-A (N=6, F=190) which have low FLOPs in theory but run slow in practice

²For small models, we should notice according to arithmetic intensity [46], real execution time is not only bounded by FLOPs.

Method	#Params (M)	Training			Testing (224 × 224)			Testing (320 × 320 / 331 × 331)		
		Input Size	Memory Cost (MB)	Speed (im/s)	#FLOPs (G)	Top-1 (%)	Top-5 (%)	#FLOPs (G)	Top-1 (%)	Top-5 (%)
NASNet-A (N=6, F=168) [51] [◇]	88.9		> 32,480	43 [‡]	-	-	-	23.8	82.7	96.2
AmoebaNet-A (N=6, F=190) [33] [◇]	86.7		> 32,480	47 [‡]	-	-	-	23.1	82.8	96.1
PNASNet-5 (N=4, F=216) [29] [◇]	86.1	331 × 331 / 320 × 320	> 32,480	38 [‡]	-	-	-	25.0	82.9	96.2
Squeeze-Excite-Net [19]	115.1		> 32,480	43 [‡]	-	-	-	42.3	83.1	96.4
AmoebaNet-A (N=6, F=448) [33] [◇]	469		> 32,480	15 [§]	-	-	-	104	83.9	96.6
Dual-Path-Net-131 [7]	79.5		31,844	83	16.0	80.1	94.9	32.0	81.5	95.8
SE-ShuffleNet v2-164 [32]	69.9		> 32,480	70 [†]	12.7	81.4	-	-	-	-
Squeeze-Excite-Net [19]	115.1	224 × 224	28,696	78	21	81.3	95.5	42.3	82.7	96.2
Oct-ResNet-152 , $\alpha = 0.125$ (ours)	60.2		15,566	162	10.9	81.4	95.4	22.2	82.3	96.0
Oct-ResNet-152 + SE ³ , $\alpha = 0.125$ (ours)	66.8		21,885	95	10.9	81.6	95.7	22.2	82.9	96.3

Table 5: ImageNet Classification results for *Large* models. The names of OctConv-equipped models are in bold font and performance numbers for related works are copied from the corresponding papers. Networks are evaluated using CuDNN v10.0⁴ in flop16 on a *single* Nvidia Titan V100 (32GB) for their training memory cost and speed. Works that employ neural architecture search are denoted by [◇]. We set batch size to 128 in most cases, but had to adjust it to 64 (noted by [†]), 32 (noted by [‡]) or 8 (noted by [§]) for networks that are too large to fit into GPU memory.

due to the use of group and depth-wise convolutions. Our proposed method is also complementary to Squeeze-and-excitation [19], where the accuracy can be further boosted when the SE-Block is added (last row).

4.4. Experiments of Video Recognition on Kinetics

In this subsection, we evaluate the effectiveness of OctConv for action recognition in videos and demonstrate that our spatial OctConv is sufficiently generic to be integrated into 3D convolution to decrease #FLOPs and increase accuracy at the same time. As shown in Table 6, OctConv consistently decreases FLOPs and meanwhile improves the accuracy when added to C2D and I3D [44, 45], and is also complementary to the Non-local [44]. This is observed for models pre-trained on ImageNet [11] as well as models trained from scratch on Kinetics. The higher accuracy, lower FLOPs and the ability of being complimentary to existing methods, *e.g.* Non-local method, confirm the effectiveness of the proposed OctConv method. Performance further increases when combining OctConv with the SlowFast Networks [12]. Specifically, we apply OctConv on the spatial dimensions and SlowFast on the temporal dimension.

5. Conclusion

In this work, we address the problem of reducing spatial redundancy that widely exists in vanilla CNN models, and propose a novel Octave Convolution operation to store and process low- and high-frequency features separately to improve the model efficiency. Octave Convolution is suffi-

³The auto-tune is set to *off* when evaluating the memory cost for more accurate result, and is set to *on* when measuring speed for fastest speed.

⁴An extra BatchNorm is added at the beginning of each residual function, otherwise the gradient will easily diverged due to the newly added SE module. This costs more memory and slows down the speed but can provide higher accuracy.

⁵Note that [12] reports 36.1 GFLOPs at a spatial size of 256², while we report (training) GFLOPs at 224² for all methods.

Method	ImageNet Pretrain	#FLOPs (G)	Top-1 (%)
(a) Kinetics-400 [3]			
I3D		28.1	72.6
Oct-I3D, $\alpha=0.1$, (ours)		25.6	73.6 (+1.0)
Oct-I3D, $\alpha=0.2$, (ours)		22.1	73.1 (+0.5)
Oct-I3D, $\alpha=0.5$, (ours)		15.3	72.1 (-0.5)
C2D	✓	19.3	71.9
Oct-C2D, $\alpha=0.1$, (ours)	✓	17.4	73.8 (+1.9)
I3D	✓	28.1	73.3
Oct-I3D, $\alpha=0.1$, (ours)	✓	25.6	74.6 (+1.3)
I3D + Non-local	✓	33.3	74.7
Oct-I3D + Non-local, $\alpha=0.1$, (ours)	✓	28.9	75.7 (+1.0)
SlowFast-R50 [12]		27.6 ⁵	75.6
Oct-SlowFast-R50, $\alpha=0.1$, (ours)		24.5	76.2 (+0.6)
Oct-SlowFast-R50, $\alpha=0.2$, (ours)		22.9	75.8 (+0.2)
(b) Kinetics-600 [2]			
I3D	✓	28.1	74.3
Oct-I3D, $\alpha=0.1$, (ours)	✓	25.6	76.0 (+1.7)

Table 6: Action Recognition in videos, ablation study, all models with ResNet50 [16].

ciently generic to replace the regular convolution operation in-place, and can be used in most 2D and 3D CNNs without model architecture adjustment. Beyond saving a substantial amount of computation and memory, Octave Convolution can also improve the recognition performance by effective communication between the low- and high-frequency and by enlarging the receptive field size which contributes to capturing more global information. Our extensive experiments on image classification and video action recognition confirm the superiority of our method for striking a much better trade-off between recognition performance and model efficiency, not only in FLOPs, but also in practice.

Acknowledgement. We would like to thank Min Lin and Xin Zhao for helpful discussions of the code development.

References

- [1] Fergus W Campbell and JG Robson. Application of fourier analysis to the visibility of gratings. *The Journal of physiology*, 197(3):551–566, 1968. 1
- [2] Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018. 2, 5, 8
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 2, 5, 8
- [4] Chun-Fu Chen, Quanfu Fan, Neil Mallinar, Tom Sercu, and Rogerio Feris. Big-little net: An efficient multi-scale feature representation for visual and speech recognition. *Proceedings of the Seventh International Conference on Learning Representations*, 2019. 2, 5, 7
- [5] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. {TVM}: An automated end-to-end optimizing compiler for deep learning. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 578–594, 2018. 6
- [6] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. Multi-fiber networks for video recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 352–367, 2018. 1
- [7] Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks. In *Advances in Neural Information Processing Systems*, pages 4467–4475, 2017. 6, 8
- [8] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Shuicheng Yan, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 7
- [9] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 1, 2
- [10] Russell L. De Valois and Karen K. De Valois. Spatial vision. *Oxford psychology series, No. 14.*, 1988. 1
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2, 5, 8
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. *ICCV*, 2019. 2, 3, 8
- [13] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 5
- [14] Jian Guo, He He, Tong He, Leonard Lausen, Mu Li, Haibin Lin, Xingjian Shi, Chenguang Wang, Junyuan Xie, Sheng Zha, Aston Zhang, Hang Zhang, Zhi Zhang, Zhongyue Zhang, and Shuai Zheng. Gluoncv and gluonnlp: Deep learning in computer vision and natural language processing. *arXiv preprint arXiv:1907.04433*, 2019. 7
- [15] Song Han, Jeff Pool, Sharan Narang, Huizi Mao, Enhao Gong, Shijian Tang, Erich Elsen, Peter Vajda, Manohar Paluri, John Tran, et al. Dsd: Dense-sparse-dense training for deep neural networks. *arXiv preprint arXiv:1607.04381*, 2016. 1, 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 2, 5, 6, 7, 8
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 2, 5, 6, 7
- [18] 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. *arXiv preprint arXiv:1704.04861*, 2017. 1, 2, 4, 5, 7
- [19] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 2, 5, 6, 7, 8
- [20] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. *ICLR*, 2018. 2
- [21] Gao Huang, Shichen Liu, Laurens Van der Maaten, and Kilian Q Weinberger. Condensenet: An efficient densenet using learned group convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2752–2761, 2018. 2, 7
- [22] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 2, 5, 6, 7
- [23] Intel. Math kernel library for deep neural networks (mkl-dnn). <https://github.com/intel/mkl-dnn/tree/7de7e5d02bf687f971e7668963649728356e0c20>, 2018. 6
- [24] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 2, 5
- [25] Tsung-Wei Ke, Michael Maire, and Stella X Yu. Multigrid neural architectures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6665–6673, 2017. 2, 5, 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] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid

- networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017. 2
- [28] Tony Lindeberg. *Scale-space theory in computer vision*, volume 256. Springer Science & Business Media, 2013. 3
- [29] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018. 2, 8
- [30] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 2
- [31] Jian-Hao Luo, Hao Zhang, Hong-Yu Zhou, Chen-Wei Xie, Jianxin Wu, and Weiyao Lin. Thinet: pruning cnn filters for a thinner net. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1, 2
- [32] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 116–131, 2018. 2, 7, 8
- [33] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 2019. 2, 8
- [34] 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*, pages 4510–4520, 2018. 2, 5, 7
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [36] Pravendra Singh, Vinay Kumar Verma, Piyush Rai, and Vinay P Nambodiri. Hetconv: Heterogeneous kernel-based convolutions for deep cnns. *arXiv preprint arXiv:1903.04120*, 2019. 2
- [37] Mallat Stephane. A wavelet tour of signal processing. 1
- [38] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 2
- [39] Wim Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM journal on mathematical analysis*, 29(2):511–546, 1998. 1
- [40] 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. 2, 5
- [41] 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*, pages 1–9, 2015. 2
- [42] Frederick Tung and Greg Mori. Clip-q: Deep network compression learning by in-parallel pruning-quantization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7873–7882, 2018. 1, 2
- [43] Huiyu Wang, Aniruddha Kembhavi, Ali Farhadi, Alan Yuille, and Mohammad Rastegari. Elastic: Improving cnns with instance specific scaling policies. *arXiv preprint arXiv:1812.05262*, 2018. 2, 5, 7
- [44] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 5, 8
- [45] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. <https://github.com/facebookresearch/video-nonlocal-net>, 2018. 5, 8
- [46] Samuel Webb Williams. *Auto-tuning performance on multicore computers*. University of California, Berkeley, 2008. 7
- [47] 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*, pages 1492–1500, 2017. 1, 2, 4, 5, 6, 7
- [48] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *Proceedings of the Sixth International Conference on Learning Representations*, 2018. 5
- [49] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6848–6856, 2018. 2, 7
- [50] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. 2
- [51] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018. 2, 8