

Approximating Continuous Convolutions for Deep Network Compression

Theo W. Costain
costain@robots.ox.ac.uk

Victor A. Prisacariu
victor@robots.ox.ac.uk

Active Vision Laboratory
University of Oxford
<https://code.active.vision>

Abstract

We present ApproxConv, a novel method for compressing the layers of a convolutional neural network. Reframing conventional discrete convolution as continuous convolution of parametrised functions over space, we use functional approximations to capture the essential structures of CNN filters with fewer parameters than conventional operations. Our method is able to reduce the size of trained CNN layers requiring only a small amount of fine-tuning. We show that our method is able to compress existing deep network models by half whilst losing only 1.86% accuracy. Further, we demonstrate that our method is compatible with other compression methods like quantisation allowing for further reductions in model size.

1 Introduction

Deep vision models have demonstrated outstanding performance across a range of problems from semantic segmentation of microscopy images[32], to object detection for nature conservation efforts[30]. In many cases, there is strong desire to apply these methods in compute constrained environments, *i.e.* mobile devices, and as a result the desire for space and compute efficient CNN models is as strong as ever.

Despite this, recent works[3, 6] have seen explosive growth in parameter count for highly performant models, and there is no reason to expect this trend to abate soon[4].

Approaches to reduce the size and compute requirements of deep models have been the focus of significant attention. A number of works have proposed models designed explicitly for efficiency[18, 20, 43], and other efforts to improve the efficiency of deep models have included methods such as low-rank factorisation[22], knowledge distillation[7], quantisation[44], and network pruning[10, 24].

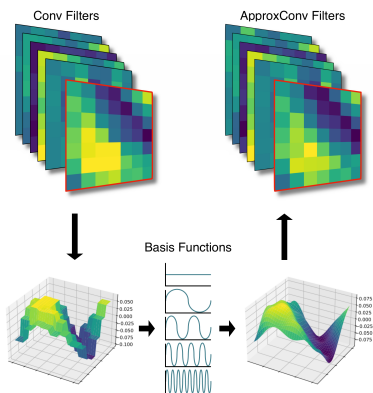


Figure 1: Our method compresses network filter weights, by approximating them using cosine and Chebyshev series.

In this work, we propose to use functional approximations to reduce the number of parameters required to fully describe a CNN layer’s filters. We re-cast conventional discrete convolution on grids as discrete convolution on parametrised continuous functions of space. In this way, we are able to replace the kernel weight tensor with an approximating function that is able to exploit the structures commonly found in CNN filters. Specifically, we make use of cosine and Chebyshev series that are able to preserve the low frequency information that allows the networks to function, whilst minimising the effects of high frequency information that are often exploited by adversarial attacks[45]. Our proposed method is particularly effective for larger kernel sizes, which despite falling out of favour in earlier deep CNN work[15, 56], have recently seen a resurgence[25].

In our experiments, we show that our method is able to significantly reduce the number of parameters in a variety of models, and unlike comparable methods[54, 40] is able to achieve this with only a limited amount of fine-tuning. Further, we demonstrate that our approach is complementary to other network size reduction methods such as quantisation. Complimenting our method with other approaches, we are able to increase the overall reduction in total model size with minimal further impact on accuracy.

In summary, the contributions of this work are as follows:

- A novel method to approximate kernel weights using functional approximations based on cosine and Chebyshev series.
- A framework that allows this method to be used without completely retraining models from scratch.

In the rest of the paper, we discuss works related and relevant to our method(Section 2), our methodology (Section 3), our experiments and results (Section 4), before a final discussion and the conclusion of the work (Section 5).

2 Related Work

A number of works[19, 20, 43] have attempted to reduce the size of models by designing networks with efficiency in mind, such as MobileNets[18] which introduced depthwise convolutions.

Quantisation methods[7, 12, 29, 57, 44] have seen great success in reducing the memory footprints of deep models. Quantising 32-bit floating point numbers down to 8, 4 or even as low as 1 bit, these methods are able to drastically reduce the memory footprint of deep models. Other works[8, 11, 39] demonstrated that non-uniform quantisation schemes, where different layers of the network can be quantised to different numbers of bits, can further reduce the size of the model without compromising accuracy. These non-uniform schemes influence some of our experiments, where we employ a similar approach.

Knowledge distillation methods[17, 28, 51, 41], typically use larger (teacher) models to guide the training of smaller (student) ones, often allowing the student models to achieve performance parity despite their smaller size.

Network pruning methods[1, 13, 24, 26] attempt to discover redundant parameters in deep networks. Having identified these parameters, they can then be removed either individually, yielding sparse networks, or as a group, sometimes removing whole filters or channels. Seminal works, like [10, 27], demonstrate that it can be possible to prune networks before training, realising the size reduction benefits for training as well as inference.

Low rank factorisation methods[6, 22, 63, 42] reduce both the total number of parameters in CNN layers, as well as the number of FLOPs required to compute the outputs of the layer. Using a variety of methods, these approaches decompose the weight tensors of layers into a series of smaller tensors. These smaller tensors can then be convolved with the inputs sequentially (in a fashion similar to separable filters from classical methods), yielding an approximation of the original output.

Both methods learn parametrised functions over space, however rather than generalised approximations, they make use of gaussian and gaussian derivative functions. Zamora et al. [40] learns only simple classical image filters (*e.g.* Sobel filter) and Saldanha et al. [54] attempt to replicate Jacobsen et al. [21] using fewer parameters. Using fractional calculus, they are able to efficiently parametrise a family of gaussian and gaussian-derivative functions using only 3[54] and 6[40] parameters per filter. While this approach allows for significant parameter savings, in comparison to our proposed method, the shape of kernels that can be learnt is substantially limited. As a result, whereas their methods require training from scratch, our approach can be applied to pre-trained models, .

In 3D settings, dense regular grid structures are too memory intensive and so receive little attention. Instead, much of the effort focuses on irregular data structures like point-clouds or meshes. In these settings it becomes desirable to have a kernel function that can vary continuously over 3D space. Similar to our method, SplineCNN[9] uses b-Splines to generate a continuous function over space. FlexConv[12] learns a kernel weight function based on the distance between sampled point locations and a hyperplane. Simonovsky and Komodakis [55] use shared MLPs with edge information as input to define the kernel weight functions.

3 Method

We begin by covering the general form of continuous convolution and how this formulation can allow for network compression. Next we outline how we come to our choice of approximation functions, before finally discussing how we apply our method to pre-trained networks.

3.1 Continuous Convolution

In general, discrete n -dimensional convolution can be expressed in the form

$$x' = \sum_{x_i \in \mathcal{N}(x)} x_i \times w_i \tag{1}$$

where $\mathcal{N}(x)$ represents the set of grid locations neighbouring x , and w , typically referred to as the kernel, contains an entry for every neighbouring location. If we replace the discrete grid in the above equation with set of locations sampled from a continuous function over space, we can re-write the above equation as

$$x'(\mathbf{p}) = \sum_{p_i \in \mathcal{N}(x)} x(\mathbf{p}_i) \times w(\mathbf{p}_i) \tag{2}$$

where $\mathbf{p} \in \mathbb{R}^n$, $x : \mathbb{R}^n \mapsto \mathbb{R}^d$, and $w : \mathbb{R}^n \mapsto \mathbb{R}^d$. To recover simple 2D convolution with a 3×3 kernel, we can simply let p be the set of 2-vectors $[(-1, -1), (-1, 0), (-1, 1), \dots]$.

This formulation has two key advantages: i) The same convolutional operation can be used on both regular nD grids, and on arbitrary nD structures (e.g. point-clouds or meshes) over which we can define a neighbourhood function \mathcal{N} ; ii) Replacing the tensor w with a carefully chosen function $w(\cdot)$, we can reduce the memory requirements of 2D convolutional layers by exploiting the structures present in the tensors.

3.2 Network Compression

Our work is concerned with the second of the two advantages above: that a judicious choice of $w(\cdot)$ can permit parameter savings. As has been shown by other works[24, 27] deep networks often contain significant amounts of redundant information, and often filters that contain low-frequency information. Accordingly, it is often possible to substantially reduce the size of a model by removing redundant parameters[13, 16] or preserving only *relatively* low-frequency structures[32].

The choice of function $w(\cdot)$ is primarily constrained by the number of parameters it requires. As discussed previously, methods like FracSRF[54] and Fractional Filters[40] choose efficiently parametrised gaussian and gaussian-derivatives as $w(\cdot)$. But these functions are not able to approximate arbitrary functions.

Separate to these approaches, the field of functional approximation has been of interest for over a century[58], motivated by applications ranging from solving PDEs to numerical integration. As a result a significant body of work is concerned with using various orthogonal basis functions for approximations. From this corpus, we use two common families of functions to allow our method to learn arbitrary kernel functions: cosine and Chebyshev series.

3.3 Cosine and Chebyshev Series

Cosine and Chebyshev series are able to approximate any periodic function that satisfies specific boundary conditions and the Dirichlet conditions. These conditions require that the function: i) is absolutely integrable; ii) is of bounded variation; iii) has a finite number of non-infinite discontinuities. As any discrete weight kernel can be fully described by a piecewise linear function, it is trivial to demonstrate¹ that this piecewise function satisfies the above conditions. Further, the requirement that the function be periodic (in the case of the cosine series) can be trivially satisfied by repeating the function at the boundaries, although careful choice of boundary conditions and approximation interval is necessary.

Although the Fourier series is generally more well known than the cosine series, extensions of the Fourier series to higher dimensions introduces a number of "cross terms" that grow in complexity as the number of dimensions increases. An alternative is to use the sine or cosine series, which are equivalent to the fourier series under certain conditions. Specifically, for an *even* function, the coefficients of the sine terms go to zero leaving only cosine terms, and vice versa for *odd* functions. By shifting the center of our approximation interval, it is possible to coerce the function to behave as an even or odd function. The advantages of the sine and cosine forms is that their extensions into higher dimensions are trivial and, for $\mathbf{x} \in \mathbb{R}^n$ with N harmonics, take the following (for cosine) form

$$\hat{w}(\mathbf{x}) = \sum_{i_0=0}^N \cdots \sum_{i_n=0}^N a_{i_0 \dots i_n} \cos(i_0 x) \cdots \cos(i_n y) \quad (3)$$

¹We leave a complete proof as an exercise for the reader

The choice of cosine over sine is motivated by concerns over boundary conditions. Specifically, for cosine series the boundary conditions require/enforce symmetric periodic repetition which minimises the potential for discontinuity at the boundary.

In our experiments, we also make use of the Chebyshev series as it is closely related to the cosine series through the definition that $T_n(\cos(x)) = \cos(nx)$, which conveniently yields a very similar expression to eq. (3)

$$\hat{w}(\mathbf{x}) = \sum_{i_0=0}^N \cdots \sum_{i_n=0}^N a_{i_0 \dots i_n} T_{i_0}(x_0) \cdots T_{i_n}(x_M) \quad (4)$$

where T_n is the n^{th} Chebyshev polynomial of the first kind, and $\mathbf{x} \in \mathbb{R}^M$. The Chebyshev series can have faster "convergence" (*i.e.* better approximation with fewer terms) than the trigonometric series in eq. (3).

To achieve compression for a $C_{\text{out}}, C_{\text{in}}, K \times K$ kernel, we choose \hat{w} to be a 2D cosine or Chebyshev series with N "harmonics" or "orders", where $N \leq K$. Replacing a 3×3 kernel, with a series 2 harmonics leads to an over 50% reduction in parameters. Our experiment show however, that 2 harmonics *can* be insufficient to properly approximate a 3×3 kernel. However, for 7×7 kernels, which have seen a resurgence in recent work[29], the increased number of harmonics permitted whilst still maintaining a reduction in parameters, can allow for much better compression depending on the complexity of the filters in question. In our experiments, we refer to replacing the kernel weights with a cosine series as **CosConv**, and **ChebConv** when using a Chebyshev series.

Whilst it is not necessary to use the same number of harmonics in the x and y directions, doing so significantly increases the hyperparameter space to be evaluated for our method. As we are not aware of any popular or common methods that make use of anisotropic kernels, we do not investigate differing spatial harmonics in this work. However, with an appropriate search method, anisotropic harmonics might permit further increases in memory savings.

Also not investigated is another possibility for reducing parameter count using our method, through variable harmonics for each filter. Although, as above, this is essentially a problem of hyperparameter/architecture search and is outside the scope of this work.

3.4 Finding Weights

Depending on the number of harmonics used in the approximating function, there may exist a closed form solution to find the approximation. In the case of the cosine series, this takes the form of the discrete cosine transform. This closed form solution, however, is only possible when choosing the same number of parameters for the approximation as for the original weight kernel. Accordingly, when approximating the kernel weights for compression, we use a simple iterative gradient descent optimisation to minimise the mean squared error loss

$$\mathcal{L} = \frac{1}{K^2} \sum_{i,j=0}^K (w_{ij} - \hat{w}(\mathbf{p}_{ij}))^2 \quad (5)$$

In conventional numerical approximation, the choice of "sample points" \mathbf{p}_{ij} is extremely important[68]. Normally, performing approximation using Chebyshev polynomials requires careful choice of sample points² to avoid Runge's phenomenon, where approximation error at the ends of the interval increases with the number of harmonics[68]. However, in our case,

²For Chebyshev series, the Chebyshev-Gauss-Lobatto points. For cosine series, equispaced points.

as we only care about the value of the approximation at the individual sample points, we can safely ignore this phenomenon.

In our experiments, we found that the Chebyshev series was significantly more sensitive to the initialisation of the weights before the initial approximation. For ChebConv, we found the best results were achieved by setting the weights of the ‘‘DC’’ harmonics to the mean of the kernel weights for each filter, and the weights of all higher harmonics are set to 0. Conversely, we found CosConv is less sensitive to the initial choice of weights, and that sampling weights from $\mathcal{N}(0, \frac{1}{C_m K^2})$ was sufficient. In both cases, we suspect that there is likely a better motivated initialisation scheme, but we leave this for future work. However for CosConv, we do not expect a better initialisation will improve performance.

4 Experiments & Results

Following this initial approximation, we fine-tune the network for a small number of epochs. This allows the network to correct for any small approximation errors that the initial approximation of weights might have caused. In all our experiments, we fine-tune the networks for 5 epochs. Specific details of training schedules, and the datasets used, are presented in the supplementary material.

4.1 CIFAR

		ResNet-20			ResNet-32				
Layer		Pre Top 1	Post Top 1	Δ	Pre Top 1	Post Top 1	Δ	# Params	Δ %
Conv2D		-	91.73	-	-	92.78	-	0.46M	0.0
Fractional[14]		91.25	91.29	0.04	-	-	-	-	-
FracSRF[15]		-	-	-	92.28	91.60	-1.18	0.16M	34.00
FracSRF†		10.00	38.08	-53.65	9.82	37.91	-54.87	0.15M	33.43
CosConv	3,3,3,2	87.22	90.75	-0.98	87.74	91.51	-1.27	0.27M	57.74
	3,3,2,2	51.43	89.55	-2.18	75.62	90.99	-1.79	0.22M	47.35
	3,2,2,2	21.16	86.99	-4.74	16.52	88.4	-4.38	0.21M	44.57
	2,2,2,3	39.63	86.87	-4.86	20.92	88.64	-4.14	0.40M	86.65
	2,2,2,2	17.55	85.98	-5.75	15.39	87.87	-4.91	0.21M	44.48
ChebConv	3,3,3,2	87.22	90.76	-0.97	87.74	91.48	-1.30	0.27M	57.74
	3,3,2,2	51.43	89.83	-1.90	75.62	91.16	-1.62	0.22M	47.35
	3,2,2,2	21.16	87.79	-3.94	16.52	88.48	-4.30	0.21M	44.57
	2,2,2,3	23.63	87.95	-3.78	20.92	89.26	-3.52	0.40M	86.65
	2,2,2,2	17.55	86.66	-5.07	15.39	88.16	-4.62	0.21M	44.48

Table 1: Results showing the application of our method to the ResNet-20 and ResNet-32 models on CIFAR10. The sequence of numbers for the CosConv and ChebConv rows denote the number of ‘‘harmonics’’ for each of the blocks in the network. † indicates using our regression and retraining approach.

We first conduct experiments on CIFAR-10[[23](#)] to validate our method. We use the smaller ResNet-20 and ResNet-32 networks from He et al. [[15](#)]. These models consist of 4 ‘‘blocks’’ of filters, with different numbers of channels for each block. The results of our experiments are presented in table 1. The sequence of numbers in the layer column next to the CosConv and ChebConv labels represent the number of harmonics used for the layers in

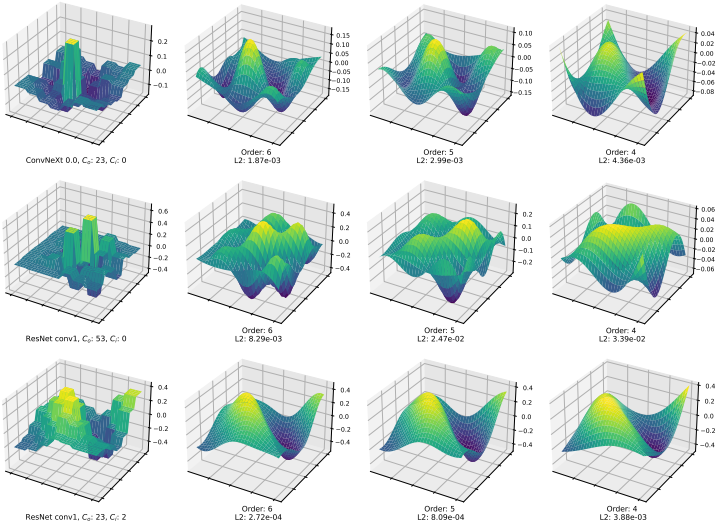


Figure 2: Filters from the ConvNeXt and ResNet networks compressed using our method. Left: Original Kernel, Right (in columns): Compressed version of this kernel using CosConv with 6, 5, and 4 orders/harmonics respectively.

each block. The columns show the accuracy after the initial approximation (Pre Top 1), and the accuracy after 5 epochs of fine tuning. The Δ columns show: i) the difference in Top 1 accuracy between the method/configuration and the baseline 2D convolution; ii) the percentage change in number of parameters. Our method is able to reduce the size of ResNet-20 by 42% losing only 1% accuracy. For both the 20 and 32 variants, the final block contains the majority of the parameters because of the larger channel dimension. Accordingly, reducing the number of parameters of this final block alone has more of an effect on the overall number of parameters than all of the other blocks combined (see $\Delta\%$ for 2,2,2,2 v.s. 2,2,2,3). Whilst not designed to approximate previously learned filters, we investigate applying our “regression and re-training” approach to both Fractional Filters[[11](#)] as well as FracSRF[[12](#)]. FracSRF was not able to approximate the kernels initially, but was able to recover some of the accuracy during fine-tuning. We found that Fractional Filters were not able to approximate the kernels to any degree even after fine tuning, and the results were no better than random guessing (10% Top 1).

4.2 ImageNet

We conduct further experiments on the ImageNet[[9](#)] dataset. These experiments use both the ResNet-18[[13](#)] and ConvNeXt-T[[14](#)] models. The results of our experiments on ResNet-18 are presented in table 2. We see that our method is able to reduce the size of the model, with a modest cost to accuracy. Again, like ResNet-20 and ResNet-32, the structure of ResNet-18 has increasing parameter counts through the blocks, and that compressing the final layers of the network provides outside reduction in parameter count. Although the top configuration reduces the number of parameters in the first layer by 30% with no loss in Top 1 accuracy, the distribution of parameters in the network mean this has a barely noticeable effect on the overall parameter count (0.02%). However, decreasing the order further (config 5,3,3,3,3),

Config	CosConv			ChebConv			# Params.	Δ %
	Pre Top 1	Post Top 1	Δ %	Pre Top 1	Post Top 1	Δ %		
Conv2D	-	69.76	0.0	-	69.76	0.0	11.68M	0.0
6,3,3,3,3	69.46	70.19	0.43	69.32	70.11	0.35	11.68M	99.98
6,3,3,3,2	64.33	68.97	-0.79	64.11	68.62	-1.14	7.09M	60.70
6,3,3,2,2	56.59	67.90	-1.86	56.35	67.23	-2.53	5.94M	50.88
6,3,2,2,2	14.56	66.89	-2.87	14.68	65.63	-4.13	5.66M	48.43
5,3,3,3,3	66.71	69.80	0.04	65.79	69.74	-0.02	10.99M	94.10
5,3,3,3,2	60.81	68.53	-1.23	59.71	68.15	-1.61	7.09M	60.68
5,3,3,2,2	51.92	67.48	-2.28	13.58	65.44	-4.32	5.94M	50.86
4,3,3,3,2	56.21	68.05	-1.71	54.37	67.34	-2.42	7.09M	60.67
4,3,3,2,2	45.99	66.84	-2.92	11.35	64.66	-5.1	5.94M	50.85

Table 2: Results showing the application of our method to ResNet-18 on the ImageNet dataset. The sequence of numbers for the CosConv and ChebConv rows denote the number of “harmonics” for each of the blocks in the network.

	Pre Top 1	Post Top 1	Δ %	# Comp. Params.	Δ %
Conv2D	-	82.10		0.33M	0.0
$7_{\times 3}, 7_{\times 3}, 7_{\times 9}, 6_{\times 3},$	80.27	81.19	-0.91	0.30M	90.90
$7_{\times 3}, 7_{\times 3}, 7_{\times 4}, 6_{\times 5}, 6_{\times 3}$	79.10	80.87	-1.23	0.27M	83.32
$7_{\times 3}, 7_{\times 3}, 6_{\times 9}, 6_{\times 3}$	70.10	79.94	-2.16	0.25M	77.25
$7_{\times 3}, 6_{\times 3}, 6_{\times 9}, 6_{\times 3}$	49.60	79.44	-2.66	0.25M	74.98
$7_{\times 3}, 7_{\times 3}, 5_{\times 9}, 5_{\times 3}$	62.40	79.28	-2.82	0.19M	58.01
$7_{\times 3}, 7_{\times 3}, 5_{\times 9}, 4_{\times 3}$	57.67	78.39	-3.71	0.17M	51.71
$7_{\times 3}, 7_{\times 3}, 4_{\times 9}, 4_{\times 3}$	20.30	75.94	-6.16	0.14M	42.26
$6_{\times 3}, 6_{\times 3}, 6_{\times 9}, 6_{\times 3}$	24.75	76.26	-5.84	0.24M	73.84
$5_{\times 3}, 5_{\times 3}, 5_{\times 9}, 5_{\times 3}$	14.94	77.29	-4.81	0.17M	51.71

Table 3: Results showing our method (CosConv) used on the 7×7 depth-wise filters from ConvNeXt, showing the accuracy after the initial approximation and then after fine tuning for 5 epochs. Comp. Params. (Compressible Parameters) shows the number of parameters making up the 7×7 convolutions in the network.

bits	Config	CosConv			ChebConv			# Params	Model size (MB)
		Top1	Quant. Top1	Delta	Top1	Quant. Top1	Delta		
32	Conv2D	69.76	69.76	0.0	69.76	69.76	0.0	11.68M	46.7
8	Conv2D	69.76	69.74	-0.02	69.76	69.74	-0.02	11.68M	23.4
8	6,3,3,3,2	68.97	68.12	-0.85	67.23	67.39	0.16	7.09M	14.2
8	6,3,3,2,2	67.90	65.23	-2.67	65.63	65.20	-0.43	5.94M	11.9
8	5,3,3,3,2	68.53	67.56	-0.97	68.15	66.80	-1.35	7.09M	14.2
8	5,3,3,2,2	67.48	64.69	-2.79	65.44	64.47	-0.97	5.94M	11.9
4	Conv2D	69.76	69.28	-0.48	69.76	69.28	-0.48	11.68M	17.5
4	6,3,3,3,2	68.97	66.70	-2.27	67.23	66.50	-0.73	7.09M	10.6
4	6,3,3,2,2	67.90	64.10	-3.80	65.63	63.88	-1.75	5.94M	8.9
4	5,3,3,3,2	68.53	66.38	-2.15	68.15	65.30	-2.85	7.09M	10.6
4	5,3,3,2,2	67.48	62.76	-4.72	65.44	63.10	-2.34	5.94M	8.9

Table 4: Results showing how our method applied to ResNet-18 on the ImageNet dataset performs under quantisation.

reduces the number of parameters in the first layer by $\approx 50\%$ with no loss to Top 1 accuracy, but still only resulting in a 6% reduction in model size. The deficiencies in the initialization of the ChebConv layers is more apparent here, where the same configuration loses more accuracy than an equivalent CosConv.

Figure 2 shows two filters from the first layer of ResNet-18 and one from the ConvNeXt-T model. We show CosConv approximations of various orders, as well as the L2 error between the approximation and the original kernel. Whilst our method is less able to recreate the strong discontinuities found in the depth-wise ConvNeXt kernels than the less discontinuous ResNet filters, our method is able to preserve much of the filter’s structure.

To better investigate the capability of our method to reduce the parameter counts of layers with a larger kernel size, we evaluate our method on the 7×7 depthwise kernels that are present throughout the network, and present these results in table 3. Because the vast majority of the parameters in ConvNeXt pertain to operations are either pointwise/linear operations, or 2×2 downsampling kernels for which our method cannot be meaningfully applied, we consider only the “compressible parameters” which are the parameters of the 7×7 depth wise convolutions. The results show that our method is able to compress the applicable layers of the network by 42% with a 2.8% loss in Top 1 accuracy. We suspect that the impact of our method on ConvNeXt is not as strong as on ResNet-18, because ConvNeXt contains significantly more discontinuous kernels.

4.3 Quantization

To show that our method can complement other existing methods for reducing the size of deep models, we perform experiments to show the application of our method in tandem with quantisation. We use a basic quantisation scheme and apply it on top of our method, quantizing the parameters of our kernel functions, as well as the activations of the network. Further details of the setup for these experiments are presented in the supplementary material. Table 4 shows our experiments quantising ResNet-18 as well as applying our method. The ‘Top 1’ column shows the performance of the unquantised configuration, and ‘Quant. Top1’ shows the results after fine-tuning, applying both our method and quantisation. The results show that a further 3-4 \times reduction in model size, compared to unquantised models (40% reduction compared to quantised conventional convolution), is achieved with a 0.37% reduction in accuracy.

5 Conclusion

In summary, we have presented a novel method to reduce the number of parameters required to represent a conventional 2D convolution. Using functional approximations in the form of cosine and Chebyshev series, we are able to represent the weights of a kernel using fewer parameters. We outline an approach to learn an initial set of weights based on pre-trained models that can be further refined through a small amount of fine-tuning. Through experiments, we demonstrated that our method is able to reduce the size of common deep vision models by as much as $\sim 50\%$ with a minor reduction in accuracy, as well as the compatibility between our method and quantisation.

In future work, we hope to examine how anisotropic and variable numbers of harmonics per filter might improve the compression.

Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council [EP/R513295/1]. We would also like to thank Henry Howard-Jenkins for constructive discussions regarding this work.

References

- [1] Sajid Anwar and Wonyong Sung. Coarse pruning of convolutional neural networks with random masks. 2016. URL <https://openreview.net/forum?id=HkvS3Mqxe>.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9650–9660, October 2021.
- [4] 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, 2009. doi: 10.1109/CVPR.2009.5206848.
- [5] Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. *Advances in neural information processing systems*, 27, 2014.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [7] Mario Drumond, Tao Lin, Martin Jaggi, and Babak Falsafi. Training dnns with hybrid block floating point. *Advances in Neural Information Processing Systems*, 31, 2018.
- [8] Ahmed T Elthakeb, Prannoy Pilligundla, Fatemehsadat Miresghallah, Amir Yazdanbakhsh, and Hadi Esmaeilzadeh. Releq: A reinforcement learning approach for automatic deep quantization of neural networks. *IEEE micro*, 40(5):37–45, 2020.
- [9] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 869–877, 2018.

- [10] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2018.
- [11] Marcelo Gennari do Nascimento, Theo W Costain, and Victor Adrian Prisacariu. Finding non-uniform quantization schemes using multi-task gaussian processes. In *European Conference on Computer Vision*, pages 383–398. Springer, 2020.
- [12] Fabian Groh, Patrick Wieschollek, and Hendrik P. A. Lensch. Flex-convolution (million-scale point-cloud learning beyond grid-worlds). In *Asian Conference on Computer Vision (ACCV)*, Dezember 2018.
- [13] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.
- [14] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)*, 2016.
- [15] 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.
- [16] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1389–1397, 2017.
- [17] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [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.
- [19] 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.
- [20] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [21] Jorn-Henrik Jacobsen, Jan Van Gemert, Zhongyu Lou, and Arnold WM Smeulders. Structured receptive fields in cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2610–2619, 2016.
- [22] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *Proceedings of the British Machine Vision Conference. BMVA Press*, 2014.

- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [24] Yann LeCun, John Denker, and Sara Solla. Optimal brain damage. *Advances in neural information processing systems*, 2, 1989.
- [25] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11976–11986, June 2022.
- [26] 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*, 41(10):2525–2538, 2018.
- [27] Eran Malach, Gilad Yehudai, Shai Shalev-Schwartz, and Ohad Shamir. Proving the lottery ticket hypothesis: Pruning is all you need. In *International Conference on Machine Learning*, pages 6682–6691. PMLR, 2020.
- [28] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI conference on artificial intelligence*, pages 5191–5198, 2020.
- [29] Marcelo Gennari do Nascimento, Roger Fawcett, and Victor Adrian Prisacariu. Dsconv: efficient convolution operator. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5148–5157, 2019.
- [30] Mohammad Sadegh Norouzzadeh, Anh Nguyen, Margaret Kosmala, Alexandra Swanson, Meredith S Palmer, Craig Packer, and Jeff Clune. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences*, 115(25):E5716–E5725, 2018.
- [31] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [33] Tara N Sainath, Brian Kingsbury, Vikas Sindhwani, Ebru Arisoy, and Bhuvana Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6655–6659. IEEE, 2013.
- [34] Nikhil Saldanha, Silvia L Pintea, Jan C van Gemert, and Nergis Tomen. Frequency learning for structured cnn filters with gaussian fractional derivatives. In *BMVC*, 2021.
- [35] Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3693–3702, 2017.

- [36] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- [37] Zhouhui Song, Zhenyu Liu, and Dongsheng Wang. Computation error analysis of block floating point arithmetic oriented convolution neural network accelerator design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [38] Lloyd N Trefethen. *Approximation theory and approximation practice [electronic resource]*. Other titles in applied mathematics. SIAM, Philadelphia, extended edition, edition, 2019. ISBN 9781611975949 (ebook).
- [39] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8612–8620, 2019.
- [40] Julio Zamora, Jesus A. Cruz Vargas, Anthony Rhodes, Lama Nachman, and Narayan Sundararajan. Convolutional filter approximation using fractional calculus. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 383–392, October 2021.
- [41] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3713–3722, 2019.
- [42] Xiangyu Zhang, Jianhua Zou, Kaiming He, and Jian Sun. Accelerating very deep convolutional networks for classification and detection. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):1943–1955, 2015.
- [43] 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.
- [44] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.
- [45] Yue Zhou, Xiaofang Hu, Jiaqi Han, Lidan Wang, and Shukai Duan. High frequency patterns play a key role in the generation of adversarial examples. *Neurocomputing*, 459:131–141, 2021.