Accelerating Convolutional Neural Networks via Activation Map Compression

Georgios Georgiadis Samsung Semiconductor, Inc.

g.geor@samsung.com

Abstract

The deep learning revolution brought us an extensive array of neural network architectures that achieve state-ofthe-art performance in a wide variety of Computer Vision tasks including among others, classification, detection and segmentation. In parallel, we have also been observing an unprecedented demand in computational and memory requirements, rendering the efficient use of neural networks in low-powered devices virtually unattainable. Towards this end, we propose a three-stage compression and acceleration pipeline that sparsifies, quantizes and entropy encodes activation maps of Convolutional Neural Networks. Sparsification increases the representational power of activation maps leading to both acceleration of inference and higher model accuracy. Inception-V3 and MobileNet-V1 can be accelerated by as much as $1.6 \times$ with an increase in accuracy of 0.38% and 0.54% on the ImageNet and CIFAR-10 datasets respectively. Quantizing and entropy coding the sparser activation maps lead to higher compression over the baseline, reducing the memory cost of the network execution. Inception-V3 and MobileNet-V1 activation maps, quantized to 16 bits, are compressed by as much as $6 \times$ with an increase in accuracy of 0.36% and 0.55% respectively.

1. Introduction

With the resurgence of Deep Neural Networks occurring in 2012 [33], efforts from the research community have led to a multitude of neural network architectures [20, 55, 59, 60] that have repeatedly demonstrated improvements in model accuracy. The price paid for this improved performance comes in terms of an increase in computational and memory cost as well as in higher power consumption. For example, AlexNet [33] requires 720 million multiplyaccumulate (MAC) operations and features 60 million parameters, while VGG-16 [55] requires a staggering 15 billion MAC's and features 138 million parameters. While the size and number of MAC operations of these networks can be handled with modern desktop computers (mainly due to the advent of Graphical Processing Units (GPU)), low-



Figure 1. Percentage of non-zero activations (above) and compression gain (below) per layer of ResNet-34. Left to right order corresponds to first to last layer in the network. Baseline corresponds to the original model, while sparse and sparse_v2 correspond to our sparsified models.

powered devices such as mobile phones and autonomous driving vehicles do not have the resources to process inputs at an acceptable rate (especially when the application demands real-time processing). Therefore, there is a great need to develop systems and algorithms that would allow us to run these networks under a limited resource budget.

There has been a plethora of works that attempt to meet this need via a variety of different methods. The majority of the approaches [18, 26, 29] attempt to approximate the weights of a neural network in order to reduce the number of parameters (model compression) and/or the number of MAC's (model acceleration). For example, pruning [18] compresses and accelerates a neural network by reducing the number of non-zero weights. Sparse tensors can be more effectively compressed leading to model compression, while multiplication with zero weights can be skipped leading to model acceleration.

Accelerating a model can also be achieved not only by zero-skipping of weights but also by zero-skipping input activation maps, a fact widely used in neural network hardware accelerators [2, 17, 31, 47]. Most modern Convolutional Neural Networks (CNN) use the Rectified Linear Unit (ReLU) [11, 16, 42] as an activation function. As a result, a large percentage of activations are zero and can be safely skipped in multiplications without any loss. However, even though the last few activation maps are typically very sparse, the first few often contain a large percentage of non-zero values (see Fig. 1 for an example based on the ResNet-34 architecture [20]). It also so happens that the first few layers process the largest sized inputs/outputs. Hence, a network would greatly benefit in terms of memory usage and model acceleration by effectively sparsifying even further these activation maps.

However, unlike weights that are trainable parameters and can be easily adapted during training, it is not obvious how one can increase the sparsity of the activations. Hence, the first contribution of this paper is to demonstrate how this can be done without any drop in accuracy. In addition, and if desired, it is also possible to use our method to trade-off accuracy for an increase in sparsity to target a variety of different hardware specifications.

Activation map size also plays an important role in power consumption. In low-powered neural network hardware accelerators (e.g. Intel Movidius Neural Compute Stick¹, DeepPhi DPU² and a plethora of others) on-chip memory is extremely limited. The accelerator is bound to access both weights and activations from the off-chip DRAM, which requires approximately $100 \times$ more power than on-chip access [21], unless the input/output activation maps and weights of the layer fit in the on-chip memory. In such cases, activation map compression is in fact disproportionately much more important than weight compression. Consider the second convolutional layer of Inception-V3, which takes as input a $149 \times 149 \times 32$ tensor and returns one of size $147 \times 147 \times 32$, totalling 1,401,920 values. The weight tensor is of size $32 \times 32 \times 3 \times 3$ totalling just 9,216 values. The total of input/output activations is in fact $150 \times$ larger in size than the number of weights. Compressing activations effectively can reduce dramatically the amount of data transferred between on-chip and off-chip memory, decreasing in turn significantly the power consumed. Thus, our second contribution is to demonstrate an effective (lossy) compression pipeline that can lead to great reductions in the size of activation maps, while still

maintaining the claim of no drop in accuracy. Following the footsteps of weight compression in [18], where the authors propose a three-stage pipeline consisting of pruning, quantization and Huffman coding, we adapt this pipeline to activation compression. Our pipeline also consists of three steps: sparsification, quantization and entropy coding. Sparsification aims to reduce the number of non-zero activations, quantization limits the bitwidth of values and entropy coding losslessly compresses the activations.

Finally, it is important to note that in several applications including, but not limited to, autonomous driving, medical imaging and other high risk operations, compromises in terms of model accuracy might not be tolerated. As a result, these applications demand networks to be compressed losslessly. In addition, lossless compression of activation maps can also benefit training since it can allow increasing the batch size or the size of the input. Therefore, our final contribution lies in the regime of lossless compression. We present a novel entropy coding algorithm, called sparseexponential-Golomb, a variant of exponential-Golomb [61], that outperforms all other tested entropy coders in compressing activation maps. Our algorithm leverages on the sparsity and other statistical properties of the maps to effectively losslessly compress them. This algorithm can stand on its own in scenarios where lossy compression is deemed unacceptable or it can be used as the last step of our threestage lossy compression pipeline.

2. Related Work

There are few works that deal with activation map compression. Gudovskiy et al. [13] compress activation maps by projecting them down to binary vectors and then applying a nonlinear dimensionality reduction (NDR) technique. However, the method modifies the network structure (which in certain use cases might not be acceptable) and it has only been shown to perform slightly better over simply quantizing activation maps. Dong et al. [10] attempt to predict which output activations are zero to avoid computing them, which can reduce the number of MAC's performed. However, their method also modifies the network structure and in addition, it does not *increase* the sparsity of the activation maps. Alwani et al. [3] reduce the memory requirement of the network by recomputing activation maps instead of storing them, which obviously comes at a computational cost. In [9], the authors perform stochastic activation pruning for adversarial defense. However, due to their sampling strategy, their method achieves best results when $\sim 100\%$ samples are picked, yielding no change in sparsity.

In terms of lossless activation map compression, Rhu *et al.* [52] examine three approaches: run-length encoding [53], zero-value compression (ZVC) and zlib compression [1]. The first two are hardware-friendly, but only achieve competitive compression when sparsity is high, while zlib

https://ai.intel.com/nervana-nnp/

²http://www.deephi.com/technology/

cannot be used in practice due to its high computational complexity. Lossless weight compression has appeared in the literature in the form of Huffman coding (HC) [18] and arithmetic coding [51].

Recently, many lightweight architectures have appeared in the literature that attempt to strike a balance between computational complexity and model accuracy [22, 24, 54, 68]. Typical design choices include the introduction of 1×1 point-wise convolutions and depthwise-separable convolutions. These networks are trained from scratch and offer an alternative to state-of-the-art solutions. When such lightweight architectures do not achieve high enough accuracy, it is possible to alternatively compress and accelerate state-of-the-art networks. Pruning of weights [14, 18, 19, 34, 36, 40, 48, 63, 69] and quantization of weights and activations [6, 7, 12, 15, 18, 23, 50, 64, 67] are the standard compression techniques that are currently used. Other popular approaches include the modification of pre-trained networks by replacing the convolutional kernels with low-rank factorizations [25, 29, 56] or grouped convolutions [26].

One could falsely consider viewing our algorithm as a method to perform activation pruning, an analog to weight pruning [18] or structured weight pruning [38, 39, 43, 66]. The latter prunes weights at a coarser granularity and in doing so also affects the activation map sparsity. However, our method affects sparsity *dynamically*, and not statically as in all other methods, since it does not *permanently* remove any activations. Instead, it encourages a smaller percentage of activations to fire for any given input, while still allowing the network to fully utilize its capacity, if needed.

Finally, activation map regularization has appeared in the literature in various forms such as dropout [57], batch normalization [27], layer normalization [4] and L_2 regularization [41]. In addition, increasing the sparsity of activations has been explored in sparse autoencoders [44] using Kullback-Leibler (KL) divergence and in CNN's using ReLU [11]. In the seminal work of Glorot *et al.* [11], the authors use ReLU as an activation function to induce sparsity on the activation maps and briefly discuss the use of L_1 regularization to enhance it. However, the utility of the regularizer is not fully explored. In this work, we expand on this idea and apply it on CNN's for model acceleration and activation map compression.

3. Learning Sparser Activation Maps

A typical cost function, $E_0(w)$, employed in CNN models is given by:

$$E_0(w) = \frac{1}{N} \sum_{n=1}^{N} c_n(w) + \lambda_w r(w),$$
 (1)

where n denotes the index of the training example, N is the mini-batch size, $\lambda_w \ge 0, w \in \mathbb{R}^d$ denotes the network



Figure 2. Computational graph based on Eq. 3. In red we illustrate the two gradient contributions for x_l .

weights, $c_n(w)$ is the data term (usually the cross-entropy) and r(w) is the regularization term (usually the L_2 norm).

Post-activation maps of training example n and layer $l \in \{1, \ldots, L\}$ are denoted by $x_{l,n} \in \mathbb{R}^{H_l \times W_l \times C_l}$, where H_l, W_l, C_l denote the number of rows, columns and channels of $x_{l,n}$. When the context allows, we write x_l rather than $x_{l,n}$ to reduce clutter. x_0 corresponds to the input of the neural network. Pre-activation maps are denoted by $y_{l,n}$. Note that since ReLU can be computed in-place, in practical applications, $y_{l,n}$ is often only an intermediate result. Therefore, we target to compress x_l rather than y_l . See Fig. 2 for an explanatory illustration of these quantities.

The overwhelming majority of modern CNN architectures achieve sparsity in activation maps through the usage of ReLU as the activation function, which imposes a hard constraint on the intrinsic structure of the maps. We propose to aid training of neural networks by explicitly encoding in the cost function to be minimized, our desire to achieve sparser activation maps. We do so by placing a sparsityinducing prior on x_l for all layers, by modifying the cost function as follows:

$$E(w) = E_0(w) + \frac{1}{N} \sum_{n=1}^{N} \sum_{l=0}^{L} \alpha_l ||x_{l,n}||_1$$

= $\frac{1}{N} \sum_{n=1}^{N} c'_n(w) + \lambda_w r(w),$ (2)

where $\alpha_l \ge 0$ for $l = 1, \dots, L - 1$, $\alpha_0 = \alpha_L = 0$ and c'_n is given by:

$$c'_{n}(w) = c_{n}(w) + \sum_{l=0}^{L} \alpha_{l} \|x_{l,n}\|_{1}.$$
(3)

In Eq. 2 we use the L_1 norm to induce sparsity on x_l that acts as a proxy to the optimal, but difficult to optimize, L_0 norm via a convex relaxation. The technique has been widely used in a variety of different applications including among others sparse coding [45] and LASSO [62].

While it is possible to train a neural network from scratch using the above cost function, our aim is to sparsify activation maps of existing state-of-the-art networks. Therefore, we modify the cost function from Eq. 1 to Eq. 2 during only the fine-tuning process of pre-trained networks. During fine-tuning, we backpropagate the gradients by computing the following quantities:



Figure 3. Inception-V3 histograms of activation maps, before (above) and after (below) sparsification, in log-space from early (Conv2d_1a_3x3, Conv2d_2b_3x3), middle (6b_branch7x7dbl_2) and late (7a_branch7x7x3_3, 7b_branch3x3dbl_3a) layers, extracted from 1000 input images from the ImageNet (ILSVRC2012) [8] training set. Activation maps are quantized to 8 bits (256 bins). The histograms show two important facts that sparse-exponential-Golomb takes advantage of: (1) Sparsity: Zero values are approximately one to two orders of magnitude more likely than any other value and (2) Long tail: Large values have a non-trivial probability. Sparsified activation maps not only have a higher percentage of zero values (enabling acceleration), but also have lower entropy (enabling compression).

$$\frac{\partial c'_n}{\partial w_l} = \frac{\partial c'_n}{\partial x_{l,n}} \cdot \frac{\partial x_{l,n}}{\partial w_l}
= \left[\frac{\partial c'_n}{\partial x_{l+1,n}} \cdot \frac{\partial x_{l+1,n}}{\partial x_{l,n}} + \frac{\partial c'_n}{\partial x_{l,n}} \right] \cdot \frac{\partial x_{l,n}}{\partial w_l},$$
(4)

where w_l corresponds to the weights of layer l, $\partial c'_n / \partial x_{l+1,n}$ is the gradient backpropagated from layer l+1, $\partial x_{l+1,n}/\partial x_{l,n}$ is the output gradient of layer l+1with respect to the input and the *j*-th element of $\partial c'_n / \partial x_{l,n}$ is given by:

$$\frac{\partial c'_{n}}{\partial x^{j}_{l,n}} = \alpha_{l} \frac{\partial \|x_{l,n}\|_{1}}{\partial x^{j}_{l,n}} = \begin{cases} +\alpha_{l}, & \text{if } x^{j}_{l,n} > 0\\ -\alpha_{l}, & \text{if } x^{j}_{l,n} < 0 \\ 0, & \text{if } x^{j}_{l,n} = 0 \end{cases}$$
(5)

where $x_{l,n}^{j}$ corresponds to the *j*-th element of (vectorized) $x_{l,n}$. Fig. 2 illustrates the computational graph and the flow of gradients for an example layer l. Note that x_l can affect c' both through x_{l+1} and *directly*, so we need to sum up both contributions during backpropagation.

4. A Sparse Coding Interpretation

Recently, connections between CNN's and Convolutional Sparse Coding (CSC) have been drawn [46, 58]. In a similar manner, it is also possible to interpret our proposed solution in Eq. 2 through the lens of sparse coding. Let us assume that for any given layer l, there exists an optimal underlying mapping \hat{y}_l that we attempt to learn. Collectively $\{\hat{y}_l\}_{l=1,\dots,L}$ define an optimal mapping from the input of the network to its output, while the network computes $\{y_l = f_l(x_{l-1}; w_l)\}_{l=1,\dots,L}$. We can then think of training a neural network as an attempt to learn the optimal mappings \hat{y}_l by minimizing the difference of the layer outputs:

minimize
$$\frac{1}{N} \sum_{n=1}^{N} \sum_{l=1}^{L} \|\hat{y}_{l,n} - y_{l,n}\|_2^2$$
 (6)
subject to $r(w_l) \le C, l = 1, \dots, L.$

In Eq. 6 we have also explicitly added the regularization term $r(w_l)$. $C \ge 0$ is a pre-defined constant that controls the amount of regularization. In the following, we restrict $r(w_l)$ to the L_2 norm, $r(w_l) = ||w_l||_2^2$, as it is by far the most common regularizer used and add the L_1 prior on x_l :

$$\begin{array}{ll} \underset{w_{l}}{\text{minimize}} & \frac{1}{N} \sum_{n=1}^{N} \left[\sum_{l=1}^{L} \| \hat{y}_{l,n} - y_{l,n} \|_{2}^{2} + \sum_{l=0}^{L} \alpha_{l} \| x_{l,n} \|_{1} \right] \\ \text{subject to} & \| w_{l} \|_{2}^{2} \leq C, l = 1, \dots, L. \end{array}$$

$$(7)$$

We can re-arrange the terms in Eq. 7 (and make use of the fact that $\alpha_L = 0$) as follows:

$$\begin{array}{ll} \underset{w_{l}}{\text{minimize}} & \frac{1}{N} \sum_{n=1}^{N} \sum_{l=1}^{L} P_{l,n} \\ \text{subject to} & \|w_{l}\|_{2}^{2} \leq C, l = 1, \dots, L \end{array}$$

$$\tag{8}$$

where:

S

$$P_{l,n} = \|\hat{y}_{l,n} - y_{l,n}\|_2^2 + \alpha_{l-1} \|x_{l-1,n}\|_1.$$
(9)

For mappings $f_l(x_{l-1}; w_l)$ that are linear, such as those computed by the convolutional and fully-connected layers, $f_l(x_{l-1}; w_l)$ can also be written as a matrix-vector multiplication, by some matrix W_l , $y_l = W_l x_{l-1}$. P_l can then be re-written as $P_l = \|\hat{y}_l - W_l x_{l-1}\|_2^2 + \alpha_{l-1} \|x_{l-1}\|_1$, which can be interpreted as a sparse coding of the optimal mapping \hat{y} . Therefore, Eq. 8 amounts to computing the sparse coding representations of the pre-activation feature maps.

5. Quantization

We quantize floating point activation maps, x_l , to q bits using linear (uniform) quantization:

$$x_l^{\text{quant}} = \frac{x_l - x_l^{\min}}{x_l^{\max} - x_l^{\min}} \times (2^q - 1), \tag{10}$$

where $x_l^{\min} = 0$ and x_l^{\max} corresponds to the maximum value of x_l in layer l in the training set. Values above x_l^{\max} in the testing set are clipped. While we do not retrain our models after quantization, it has been shown by the literature to improve model accuracy [5, 37]. We also believe that a joint optimization scheme where we simultaneously perform quantization and sparsification can further improve our results and leave it for future work.

6. Entropy Coding

A number of different schemes have been devised to store sparse matrices effectively. Compressed sparse row (CSR) and compressed sparse column (CCS) are two representations that strike a balance between compression and efficiency of arithmetic operation execution. However, such methods assume that the entire matrix is available prior to storage, which is not necessarily true in all of our use cases. For example, in neural network hardware accelerators data are often streamed as they are computed and there is a great need to compress in an on-line fashion. Hence, we shift our focus to algorithms that can encode one element at a time. We present a new entropy coding algorithm, sparseexponential-Golomb (SEG) (Alg. 1) that is based on the effective exponential-Golomb (EG) [61, 65]. SEG leverages on two facts: (1) Most activation maps are sparse and, (2) The first-order probability distribution of the activation maps has a long tail.

Exponential-Golomb is most commonly used in H.264 [28] and HEVC [30] video compression standards with k = 0 as a standard parameter. k = 0 is a particularly effective parameter value when data is sparse since it assigns a code word of length 1 for the value x = 0. However, it also requires that the probability distribution of values falls-off rapidly. Activation maps have a non-trivial number of large values (especially for q = 8, 12, 16), rendering k = 0 ineffective. Fig. 3 demonstrates this fact by showing histograms of activation maps for a sample of layers from Inception-V3. While this issue could be solved by using larger values of k, the consequence is that the value x = 0 is no longer encoded with a 1 bit code word (see Table 1). SEG solves this problem by dedicating the code word '1' for x = 0 and by pre-appending the code word generated

Algorithm 1 Sparse-exponential-Golomb

```
Input: Non-negative integer x, Order k
Output: Bitstream y
function encode_sparse_exp_Golomb (x, k)
{
  If k == 0:
     y = \text{encode}_\exp_\text{Golomb}(x, k)
  Else:
     If x == 0:
           Let y = '1'
     Else:
          Let y = 0^{\circ} + \text{encode exp } \text{Golomb}(x - 1, k)
  Return y
Input: Bitstream x, Order k
Output: Non-negative integer y
function decode_sparse_exp_Golomb (x, k)
{
  If k == 0:
     y = \text{decode}_\exp_\text{Golomb}(x, k)
  Else:
     If x[0] == 1:
          Let y = 0
     Else:
           Let y = 1 + \text{decode\_exp\_Golomb}(x[1:], k)
  Return y
}
```

Algorithm	k = 0	k = 4	k = 8	k = 12
EG	1	5	9	13
SEG	1	1	1	1

Table 1. Code word length comparison of x = 0 between EG and SEG for different values of k.

by EG with a '0' for x > 0. Sparse-exponential golomb can be found in Alg. 1 while exponential-Golomb [61] is provided in the supplementary material³.

7. Experiments

In the experimental section, we investigate two important applications: (1) acceleration of computation and (2) compression of activation maps. We carry our experiments on three different datasets, MNIST [35], CIFAR-10 [32] and ImageNet ILSVRC2012 [8] and five different networks, a LeNet-5 [35] variant⁴, MobileNet-V1 [22], Inception-V3 [60], ResNet-18 [20] and ResNet-34 [20]. These networks cover a wide variety of network sizes, complexity in design and efficiency in computation. For example, unlike AlexNet [33] and VGG-16 [55] that are over-parameterized and easy

³https://georgios0.github.io/cvpr2019/ ⁴https://github.com/pytorch/examples/

Dataset	Model	Variant	Top-1 Acc.	Top-5 Acc.	Acts. (%)	Speed-up
MNICT	LaNat 5	Baseline	98.45%	-	53.73%	1.0 imes
MINIST	Lenet-5	Sparse	98.48% (+0.03%)	-	23.16%	2.32 ×
CIEAD 10	MahilaNat V1	Baseline	89.17%	-	47.44%	1.0 imes
CIFAR-10	MobileNet-VI	Sparse	89.71% (+0.54%)	-	29.54%	1.61 ×
		Baseline	75.76%	92.74%	53.78%	$1.0 \times$
	Inception-V3	Sparse	76.14% (+0.38%)	92.83% (+0.09%)	33.66%	$1.60 \times$
		Sparse_v2	68.94% (-6.82%)	88.52% (-4.22%)	25.34%	2.12 ×
		Baseline	69.64%	88.99%	60.64%	$1.0 \times$
ImageNet	ResNet-18	Sparse	69.85% (+0.21%)	89.27% (+0.28%)	49.51%	$1.22 \times$
		Sparse_v2	68.62% (-1.02%)	88.41% (-0.58%)	34.29%	1.77 ×
	ResNet-34	Baseline	73.26%	91.43%	57.44%	$1.0 \times$
		Sparse	73.95%(+0.69%)	91.61% (+0.18%)	46.85%	$1.23 \times$
		Sparse_v2	67.73% (-5.53%)	87.93% (-3.50%)	29.62%	1.94 ×

Table 2. Accelerating neural networks via sparsification. Numbers in brackets indicate change in accuracy. Acts. (%) shows the percentage of non-zero activations.

Network	Algorithm	Top-1 Acc. Change	Top-5 Acc. Change	Speed-up
	Ours (Sparse)	+0.21%	+0.28%	18.4%
	Ours (Sparse_v2)	-1.02%	-0.58%	43.5%
ResNet-18	LCCL [10]	-3.65%	-2.30%	34.6%
	BWN [50]	-8.50%	-6.20%	50.0%
	XNOR [50]	-18.10%	-16.00%	98.3%
	Ours (Sparse)	+0.69%	+0.18%	18.4%
D N-+ 24	Ours (Sparse_v2)	-5.53%	-3.50%	48.4%
Resinet-34	LCCL [10]	-0.43%	-0.17%	24.8%
	PFEC [36]	-1.06%	-	24.2%
	Ours (Sparse)	+0.03%	-	56.9%
LeNet-5	[18] (p = 70%)	-0.12%	-	7.3%
	[18] (p = 80%)	-0.57%	-	14.7%

Table 3. Comparison between various state-of-the-art acceleration methods on ResNet-18/34 and LeNet-5. For ResNet-18/34, we emphasized in bold Sparse_v2 as a good compromise between acceleration and accuracy. p = 70% indicates that the pruned network has 30% non-zero weights. Speed-up calculation follows the convention from [10], where it is reported as 1 - (non-zero activations of sparse model) / (non-zero activations of baseline).

to compress, MobileNet-V1 is an architecture that is already designed to reduce computation and memory consumption and as a result, it presents a challenge for achieving further savings. Furthermore, compressing and accelerating Inception-V3 and the ResNet architectures has tremendous usefulness in practical applications since they are among the state-of-the-art image classification networks.

Acceleration. In Table 2, we summarize our speed-up results. Baselines were obtained as follows: LeNet-5 and MobileNet-V1⁵ were trained from scratch, while Inception-V3 and ResNet-18/34 were obtained from the PyTorch [49] repository⁶. The results reported were computed on the validation sets of the corresponding datasets. The speed-up factor is calculated by dividing the number of non-zero activations of the baseline by the number of non-zero activa-



(a) MobileNet-V1 (b) Inception-V3 Figure 4. Evolution of Top-1 accuracy and activation map sparsity of MobileNet-V1 and Inception-V3 during training on the validation sets of CIFAR-10 and ILSVRC2012 respectively. The red arrow indicates the checkpoint selected for the sparse model.

tions of the sparse models. For Inception-V3 and ResNet-18/34, we present two sparse variants, one targeting minimum reduction in accuracy (sparse) and the other targeting high sparsity (sparse_v2).

Many of our sparse models not only have an increased sparsity in their activation maps, but also demonstrate increased accuracy. This alludes to the well-known fact that sparse activation maps have strong representational power [11] and the addition of the sparsity-inducing prior does not necessarily trade-off accuracy for sparsity. When the regularization parameter is carefully chosen, the data term and the prior can work together to improve both the Top-1 accuracy and the sparsity of the activation maps. Inception-V3 can be accelerated by as much as $1.6 \times$ with an increase in accuracy of 0.38%, while ResNet-18 achieves a speed-up of $1.8 \times$ with a modest 1% accuracy drop. LeNet-5 can be accelerated by $2.3 \times$, MobileNet-V1 by $1.6 \times$ and finally ResNet-34 by $1.2 \times$, with all networks exhibiting ac*curacy increase.* The regularization parameters, α_l , can be adapted for each layer independently or set to a common constant. We experimented with various configurations and

⁵https://github.com/kuangliu/pytorch-cifar/

⁶https://pytorch.org/docs/stable/torchvision/



Figure 5. Percentage of non-zero activations (above) and compression gain (below) per layer for various network architectures before and after sparsification.

Dataset	Model	Algorithm	t=1,000	t=10,000	t=30,000	t=60,000
MNIST	LeNet-5	SEG	$1.700 \times$	$1.700 \times$	$1.701 \times$	$1.701 \times$
		ZVC[52]	$1.665 \times$	$1.666 \times$	$1.666 \times$	$1.667 \times$
Dataset	Model	Algorithm	t = 500	t=1,000	t=2,000	t = 5,000
ImageNet	Inception-V3	SEG	$1.763 \times$	$1.769 \times$	$1.774 \times$	$1.779 \times$
magerver		ZVC[52]	1.652×	1.655×	1.661×	1.667×

Table 4. Effect of evaluation set size on compression performance. Varying the size yields minor compression gain changes, indicating that a smaller dataset can serve as a good benchmark for evaluation purposes.

the selected parameters are shared in the supplementary material. To determine the selected α_l , we used grid-based hyper-parameter optimization. The number of epochs required for fine-tuning varied for each network as Fig. 4 illustrates. We chose to typically train up to 90-100 epochs and then selected the most appropriate result. In Fig. 3, we show the histograms of a selected number of Inception-V3 layers before and after sparsification. Histograms of the sparse model have a greater proportion of zero-values, leading to model acceleration as well as lower entropy, leading to higher activation map compression.

In Fig. 4, we show how the accuracy and activation map sparsity of MobiletNet-V1 and Inception-V3 changes during training evaluated on the validation sets of CIFAR-10 and ILSVRC2012 respectively. On the same figure, we show the checkpoint selected to report results in Table 2. Selecting a checkpoint trades-off accuracy with sparsity and which point is chosen depends on the application at hand. In the top row of Fig. 5 we show the percentage of non-zero activations per layer for the five networks. While some layers (e.g. last few layers in MobileNet-V1) exhibit tremendous decrease in non-zero activations, most of the acceleration comes by effectively sparsifying the first few layers of a network. Finally, in Table 3 we compare our approach to other state-of-the-art methods as well as to a weight pruning algorithm [18]. In particular, note that [18] is not effective in increasing the sparsity of activations. Overall, our approach can achieve both high acceleration (with a slight accuracy drop) and high accuracy (with lower acceleration).

Model	Variant	Measurement	float32	uint16	uint12	uint8
LeNet-5	Baseline	Top-1 Acc. Compression	98.45%	98.44% (-0.01%) 3.40× (1.70×)	98.44% (-0.01%) 4.40× (1.64×)	98.39% (-0.06%) 6.32× (1.58×)
(MNIST)	Sparse	Top-1 Acc. Compression	98.48% (+0.03%)	98.48% (+0.03%) 6.76× (3.38×)	98.49% (+0.04%) 8.43× (3.16×)	98.46% (+0.01%) 11.16× (2.79×)
MobiletNet-V1 (CIFAR-10)	Baseline	Top-1 Acc. Compression	89.17%	89.18% (+0.01%) 5.52× (2.76×)	89.15% (-0.02%) 7.09× (2.66×)	89.16% (-0.01%) 9.76× (2.44×)
	Sparse	Top-1 Acc. Compression	89.71% (+0.54%)	89.72% (+0.55%) 5.84× (2.92×)	87.72 (+0.55%) 7.33× (2.79×)	89.62% (+0.45%) 10.24× (2.56×)

Table 5. Effect of quantization on compression on SEG. LeNet-5 is compressed by $11 \times$ and MobileNet-V1 by $10 \times$. In brackets, we report change in accuracy and compression gain over the float32 baseline.

Compression. We carry our compression experiments on a subset of each network's activation maps, since caching activation maps of the entire training and testing sets is prohibitive due to storage constraints. In Table 4 we show the effect of testing size on the compression performance of two algorithms, SEG and zero-value compression (ZVC) [52] on MNIST and ImageNet. Evaluating compression yields minor differences as a function of size. Similar observations can be made on all investigated networks and datasets. In subsequent experiments, we use the entire testing set of MNIST to evaluate LeNet-5, the entire testing set of CIFAR-10 to evaluate MobileNet-V1 and 5,000 randomly selected images from the validation set of ILSVRC2012 to evaluate Inception-V3 and ResNet-18/34. The Top-1/Top-5 accuracy, however, is measured on the entire validation sets.

In Table 6, we summarize the compression results. Compression gain is defined as the ratio of activation size before and after compression. We evaluate our method (SEG) against exponential-Golomb (EG) [61], Huffman Coding (HC) used in [18], zero-value compression (ZVC) [52] and ZLIB [1]. We report the Top-1/Top-5 accuracy of the quantized models at 16 bits. Parameter choice for SEG and EG is provided in Table 7. SEG outperforms all methods in both the baseline and sparse models validating the claim that it is a very effective encoder for distributions with high sparsity and long tail. SEG achieves almost $7 \times$ compression gain on LeNet-5, almost $6 \times$ on MobileNet-V1 and Inception-V3 and more than $4 \times$ compression gain in the ResNet architectures while also featuring an increase in accuracy.

It is also clear that our sparsification step leads to greater

Dataset	Model	Variant	Bits	Top-1 Acc.	Top-5 Acc.	SEG	EG [61]	HC [18]	ZVC [52]	ZLIB [1]
MNIST		Baseline	float32	98.45%	-	-	-	-	-	-
	LeNet-5	Baseline	wint16	98.44% (-0.01%)	-	$3.40 \times (1.70 \times)$	$2.30 \times (1.15 \times)$	$2.10 \times (1.05 \times)$	$3.34 \times (1.67 \times)$	$2.42 \times (1.21 \times)$
		Sparse	unitio	98.48% (+0.03%)	-	6.76× (3.38×)	$4.54 \times (2.27 \times)$	$3.76 \times (1.88 \times)$	$6.74 \times (3.37 \times)$	$3.54 \times (1.77 \times)$
		Baseline	float32	89.17%	-	-	-	-	-	-
CIFAR-10	MobileNet-V1	Baseline	wint16	89.18% (+0.01%)	-	5.52 imes (2.76 imes)	$3.70 \times (1.85 \times)$	$2.90 \times (1.45 \times)$	$5.32 \times (2.66 \times)$	$3.76 \times (1.88 \times)$
		Sparse	unitio	89.72% (+0.55%)	-	5.84× (2.92×)	$3.90 \times (1.95 \times)$	$3.00 \times (1.50 \times)$	$5.58 \times (2.79 \times)$	$3.90 \times (1.95 \times)$
		Baseline	float32	75.76%	92.74%	-	-	-	-	-
	Incontion V2	Baseline		75.75% (-0.01%)	92.74% (+0.00%)	3.56× (1.78×)	$2.42 \times (1.21 \times)$	$2.66 \times (1.33 \times)$	$3.34 \times (1.67 \times)$	$2.66 \times (1.33 \times)$
	inception-v 5	Sparse	uint16	76.12% (+0.36%)	92.83% (+0.09%)	$5.80 \times (2.90 \times)$	$4.10 \times (2.05 \times)$	$4.22 \times (2.11 \times)$	$5.02 \times (2.51 \times)$	$3.98 \times (1.99 \times)$
		Sparse_v2		68.96% (-6.80%)	88.54% (-4.20%)	6.86 imes (3.43 imes)	5.12 imes (2.56 imes)	$5.12 \times (2.56 \times)$	$6.36 \times (3.18 \times)$	$4.90 \times (2.45 \times)$
1	ResNet-18	Baseline	float32	69.64%	88.99%	-	-	-	-	-
ImageNet		Baseline		69.64% (+0.00%)	88.99% (+0.00%)	$3.22 \times (1.61 \times)$	$2.32 \times (1.16 \times)$	$2.54 \times (1.27 \times)$	$3.00 \times (1.50 \times)$	$2.32 \times (1.16 \times)$
inagertet		Sparse	uint16	69.85% (+0.21%)	89.27% (+0.28%)	$4.00 \times (2.00 \times)$	$2.70 \times (1.35 \times)$	$3.04 \times (1.52 \times)$	$3.60 \times (1.80 \times)$	$2.68 \times (1.34 \times)$
		Sparse_v2		68.62% (-1.02%)	88.41% (-0.58%)	5.54× (2.77×)	$3.80 \times (1.90 \times)$	$4.02 \times (2.01 \times)$	$4.94 \times (2.47 \times)$	$3.54 \times (1.77 \times)$
		Baseline	float32	73.26%	91.43%	-	-	-	-	-
	DasNat 24	Baseline		73.27% (+0.01%)	91.43% (+0.00%)	$3.38 \times (1.69 \times)$	$2.38 \times (1.19 \times)$	$2.56 \times (1.28 \times)$	$3.14 \times (1.57 \times)$	$2.46 \times (1.23 \times)$
	Resinet-54	Sparse	uint16	73.96% (+0.70%)	91.61% (+0.18%)	$4.18 \times (2.09 \times)$	$2.84 \times (1.42 \times)$	$3.04 \times (1.52 \times)$	$3.78 \times (1.89 \times)$	$2.84 \times (1.42 \times)$
		Sparse_v2		67.74% (-5.52%)	87.90% (-3.53%)	$6.26 \times (3.13 \times)$	$4.38 \times (2.19 \times)$	$4.32 \times (2.16 \times)$	$5.58 \times (2.79 \times)$	$4.02 \times (2.01 \times)$

Table 6. Compressing activation maps. We report the Top-1/Top-5 accuracy, with the numbers in brackets indicating the change in accuracy. The total compression gain is reported for various state-of-the-art algorithms (in brackets we also report the compression gain without including gains from quantization). SEG outperforms other state-of-the-art algorithms in all models and datasets.

Dataset	Model	Variant	SEG	EG [<mark>61</mark>]
MNIST	LeNet-5	Baseline Sparse	$\begin{array}{l} k=12\\ k=14 \end{array}$	k = 9 $k = 0$
CIFAR-10	MobileNet-V1	Baseline Sparse	$\begin{array}{l} k=13\\ k=13 \end{array}$	k = 0 $k = 0$
	Inception-V3	Baseline Sparse Sparse_v2	k = 12 $k = 10$ $k = 13$	k = 7 $k = 0$ $k = 0$
ImageNet	ResNet-18	Baseline Sparse Sparse_v2	k = 12 $k = 12$ $k = 11$	k = 10 $k = 0$ $k = 0$
	ResNet-34	Baseline Sparse Sparse_v2	k = 12 $k = 12$ $k = 11$	k = 9 $k = 0$ $k = 0$

Table 7. SEG and EG parameter values. Parameters were selected on a separate training set composed of 1,000 randomly selected images from each dataset. SEG fits the data distribution better by splitting the values into two sets: zero and non-zero. When searching for the optimal parameter value, we fit the parameter to the non-zero value distribution. EG fits the parameter to both sets simultaneously resulting in a sub-optimal solution.

compression gains, as can be seen by comparing the baseline and sparse models of each network. When compared to the baseline, the sparse models of LeNet-5, MobileNet-V1, Inception-V3, ResNet-18 and ResNet-34 can be *compressed* by $2.0 \times , 1.1 \times , 1.6 \times , 1.2 \times , 1.2 \times$ respectively *more* than the *compressed* baseline, while also featuring an *increase* in accuracy and an *accelerated* computation of $2.3 \times , 1.6 \times , 1.2 \times , 1.2 \times$ respectively. Table 6 also demonstrates that our pipeline can achieve even greater compression gains if slight accuracy drops are acceptable. The sparsification step induces higher sparsity and lower entropy in the distribution of values, which both lead to these additional gains (Fig. 3). While all compression algorithms benefit from the sparsification step, SEG is the most effective in exploiting the resulting distribution of values. In the bottom row of Fig. 5 we show the compression gain per layer for various networks. Finally, we study the effect of quantization on compression for SEG in Table 5. We evaluate the baseline and sparse models of LeNet-5 and MobileNet-V1 at q = 16, 12, 8 bits and report compression gain and accuracy. We report the accuracy change and compression gain compared to the floating-point baseline model. LeNet-5 can be compressed by as much as $11 \times$ with a 0.01% accuracy gain, while MobileNet-V1 can be compressed by $10 \times$ with a 0.45% accuracy gain.

8. Conclusion

We have presented a three-stage compression and acceleration pipeline that sparsifies, quantizes and encodes activation maps of CNN's. The sparsification step increases the number of zero values leading to model acceleration on specialized hardware, while the quantization and encoding stages lead to compression by effectively utilizing the lower entropy of the sparser activation maps. The pipeline demonstrates an effective strategy in reducing the computational and memory requirements of modern neural network architectures, taking a step closer to realizing the execution of state-of-the-art neural networks on low-powered devices. At the same time, we have demonstrated that adding a sparsity-inducing prior on the activation maps does not necessarily come at odds with the accuracy of the model, alluding to the well-known fact that sparse activation maps have a strong representational power. Furthermore, we motivated our proposed solution by drawing connections between our approach and sparse coding. Finally, we believe that an optimization scheme in which we jointly sparsify and quantize neural networks can lead to further improvements in acceleration, compression and accuracy of the models.

Acknowledgments. We would like to thank Hui Chen, Weiran Deng and Ilia Ovsiannikov for valuable discussions.

References

- Zlib compressed data format specification version 3.3. https://tools.ietf.org/html/rfc1950. Accessed: 2018-05-17. 2, 7, 8
- [2] J. Albericio, P. Judd, T. Hetherington, T. Aamodt, N. E. Jerger, and A. Moshovos. Cnvlutin: Ineffectual-neuron-free deep neural network computing. In 2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture, 2016. 2
- [3] M. Alwani, H. Chen, M. Ferdman, and P. Milder. Fusedlayer cnn accelerators. In *IEEE/ACM International Sympo*sium on Microarchitecture, 2016. 2
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. arXiv preprint arXiv:1607.06450, 2016. 3
- [5] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. arXiv preprint arXiv:1702.00953, 2017. 5
- [6] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Training deep neural networks with low precision multiplications. arXiv preprint arXiv:1412.7024, 2014. 3
- [7] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In Advances in Neural Information Processing Systems, 2015. 3
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009. 4, 5
- [9] Guneet S Dhillon, Kamyar Azizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. Stochastic activation pruning for robust adversarial defense. arXiv preprint arXiv:1803.01442, 2018.
- [10] Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. More is less: A more complicated network with less inference complexity. In *The IEEE International Conference on Computer Vision*, 2017. 2, 6
- [11] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *International Conference* on Artificial Intelligence and Statistics, 2011. 2, 3, 6
- [12] Yunchao Gong, Liu Liu, Ming Yang, and Lubomir D. Bourdev. Compressing deep convolutional networks using vector quantization. arXiv preprint arXiv:1412.6115, 2014. 3
- [13] D. A. Gudovskiy, A. Hodgkinson, and L. Rigazio. DNN Feature Map Compression using Learned Representation over GF(2). arXiv preprint arXiv:1808.05285, 2018. 2
- [14] Yiwen Guo, Anbang Yao, and Yurong Chen. Dynamic network surgery for efficient dnns. arXiv preprint arXiv:1608.04493, 2016. 3
- [15] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, 2015. 3
- [16] Richard HR Hahnloser, Rahul Sarpeshkar, Misha A Mahowald, Rodney J Douglas, and H Sebastian Seung. Digi-

tal selection and analogue amplification coexist in a cortexinspired silicon circuit. *Nature*, 2000. 2

- [17] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally. Eie: Efficient inference engine on compressed deep neural network. *SIGARCH Computer Architecture News*, 2016. 2
- [18] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149, 2015. 1, 2, 3, 6, 7, 8
- [19] Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeo. In Advances in Neural Information Processing Systems, 1993. 3
- [20] 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, 2016. 1, 2, 5
- [21] M. Horowitz. Computing's energy problem (and what we can do about it). In *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, 2014. 2
- [22] 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. 3, 5
- [23] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. arXiv preprint arXiv:1609.07061, 2016. 3
- [24] Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Song Han, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size. arXiv preprint arXiv:1602.07360, 2016. 3
- [25] Yani Ioannou, Duncan Robertson, Jamie Shotton, Roberto Cipolla, and Antonio Criminisi. Training cnns with lowrank filters for efficient image classification. arXiv preprint arXiv:1705.08665, 2015. 3
- [26] Yani Ioannou, Duncan P. Robertson, Roberto Cipolla, and Antonio Criminisi. Deep roots: Improving CNN efficiency with hierarchical filter groups. arXiv preprint arXiv:1605.06489, 2016. 1, 3
- [27] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167, 2015. 3
- [28] ITU-T. Itu-t recommendation h. 264: Advanced video coding for generic audiovisual services. 5
- [29] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. arXiv preprint arXiv:1405.3866, 2014. 1, 3
- [30] JCT-VC. Itu-t recommendation h. 265 and iso. 5
- [31] D. Kim, J. Ahn, and S. Yoo. Zena: Zero-aware neural network accelerator. *IEEE Design Test*, 2018. 2
- [32] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.5
- [33] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural net-

works. In Advances in Neural Information Processing Systems, 2012. 1, 5

- [34] Vadim Lebedev and Victor Lempitsky. Fast convnets using group-wise brain damage. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
 3
- [35] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradientbased learning applied to document recognition. *Proceedings of the IEEE*, 1998. 5
- [36] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710, 2016. 3, 6
- [37] Darryl Lin, Sachin Talathi, and Sreekanth Annapureddy. Fixed point quantization of deep convolutional networks. In *International Conference on Machine Learning*, 2016. 5
- [38] Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. arXiv preprint arXiv:1705.08665, 2017. 3
- [39] Christos Louizos, Max Welling, and Diederik P. Kingma. Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312*, 2017. 3
- [40] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. arXiv preprint arXiv:1707.06342, 2017. 3
- [41] Stephen Merity, Bryan McCann, and Richard Socher. Revisiting activation regularization for language rnns. arXiv preprint arXiv:1708.01009, 2017. 3
- [42] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.
 2
- [43] Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Structured bayesian pruning via log-normal multiplicative noise. arXiv preprint arXiv:1705.07283, 2017. 3
- [44] Andrew Ng. Cs294a lecture notes, 2011. https: //web.stanford.edu/class/cs294a/ sparseAutoencoder.pdf. 3
- [45] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? Vision research, 1997. 3
- [46] Vardan Papyan, Yaniv Romano, and Michael Elad. Convolutional neural networks analyzed via convolutional sparse coding. *The Journal of Machine Learning Research*, 18(1):2887–2938, 2017. 4
- [47] A. Parashar, M. Rhu, A. Mukkara, A. Puglielli, R. Venkatesan, B. Khailany, J. Emer, S. W. Keckler, and W. J. Dally. Scnn: An accelerator for compressed-sparse convolutional neural networks. In 2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture, 2017. 2
- [48] Jongsoo Park, Sheng Li, Wei Wen, Ping Tak Peter Tang, Hai Li, Yiran Chen, and Pradeep Dubey. Faster cnns with direct sparse convolutions and guided pruning. *arXiv preprint arXiv:1608.01409*, 2016. 3
- [49] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6

- [50] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference* on Computer Vision, 2016. 3, 6
- [51] Brandon Reagen, Udit Gupta, Robert Adolf, Michael M. Mitzenmacher, Alexander M. Rush, Gu-Yeon Wei, and David Brooks. Weightless: Lossy weight encoding for deep neural network compression. arXiv preprint arXiv:1711.04686. 3
- [52] Minsoo Rhu, Mike O'Connor, Niladrish Chatterjee, Jeff Pool, and Stephen W. Keckler. Compressing DMA engine: Leveraging activation sparsity for training deep neural networks. arXiv preprint arXiv:1705.01626, 2017. 2, 7, 8
- [53] AH Robinson and Colin Cherry. Results of a prototype television bandwidth compression scheme. *Proceedings of the IEEE*, 1967. 2
- [54] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. arXiv preprint arXiv:1801.04381, 2018. 3
- [55] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 1, 5
- [56] A. Sironi, B. Tekin, R. Rigamonti, V. Lepetit, and P. Fua. Learning separable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015. 3
- [57] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 2014. 3
- [58] Jeremias Sulam, Vardan Papyan, Yaniv Romano, and Michael Elad. Multi-layer convolutional sparse modeling: Pursuit and dictionary learning. arXiv preprint arXiv:1708.08705, 2017. 4
- [59] Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *arXiv preprint arXiv:1602.07261*, 2016. 1
- [60] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 1, 5
- [61] Jukka Teuhola. A compression method for clustered bitvectors. *Information processing letters*, 1978. 2, 5, 7, 8
- [62] Robert Tibshirani. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological), 1996. 3
- [63] Karen Ullrich, Edward Meeds, and Max Welling. Soft weight-sharing for neural network compression. arXiv preprint arXiv:1702.04008, 2017. 3
- [64] Vincent Vanhoucke, Andrew Senior, and Mark Z Mao. Improving the speed of neural networks on cpus. In Proceedings Deep Learning and Unsupervised Feature Learning NIPS Workshop, 2011. 3

- [65] Jiangtao Wen and John D Villasenor. Reversible variable length codes for efficient and robust image and video coding. In *Data Compression Conference*, 1998. 5
- [66] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, Hai Li, Christos Louizos, Max Welling, and Diederik P. Kingma. Learning structured sparsity in deep neural networks. arXiv preprint arXiv:1608.03665, 2016. 3
- [67] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 3
- [68] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. arXiv preprint arXiv:1707.01083, 2017. 3
- [69] Hao Zhou, Jose M Alvarez, and Fatih Porikli. Less is more: Towards compact cnns. In *European Conference on Computer Vision*, 2016. 3