

DECORE: Deep Compression with Reinforcement Learning

Manoj Alwani
Element Inc.

ma@discoverelement.com

Yang Wang
Element Inc.

yw@discoverelement.com

Vashisht Madhavan
Element Inc.

vm@discoverelement.com

Abstract

Deep learning has become an increasingly popular and powerful methodology for modern pattern recognition systems. However, many deep neural networks have millions or billions of parameters, making them untenable for real-world applications due to constraints on memory size or latency requirements. As a result, efficient network compression techniques are often required for the widespread adoption of deep learning methods. We present DECORE, a reinforcement learning-based approach to automate the network compression process. DECORE assigns an agent to each channel in the network along with a light policy gradient method to learn which neurons or channels to be kept or removed. Each agent in the network has just one parameter (keep or drop) to learn, which leads to a much faster training process compared to existing approaches. DECORE provides state-of-the-art compression results on various network architectures and various datasets. For example, on the ResNet-110 architecture, it achieves a 64.8% compression and 61.8% FLOPs reduction as compared to the baseline model without any accuracy loss on the CIFAR-10 dataset. It can reduce the size of regular architectures like the VGG network by up to 99% with just a small accuracy drop of 2.28%. For a larger dataset like ImageNet with only 30 epochs of training, it can compress the ResNet-50 architecture by 44.7% and reduce FLOPs by 42.3%, with just a 0.69% drop on Top-5 accuracy of the uncompressed model. We also demonstrate that DECORE can be used to search for compressed network architectures based on various constraints, such as memory and FLOPs.

1. Introduction

Deep neural networks (DNNs) have led to significant advances in image recognition tasks, most notably benchmarked by the ImageNet challenge [42]. Recent trends suggest that deeper architectures with more parameters and more complex computational blocks lead to better results [13, 47]. These increases in model size, however, incur a higher inference cost, making them difficult to be

deployed in embedded systems or mobile devices. Additionally, real-world applications have memory, latency, or throughput constraints which make it difficult to use large convolutional neural networks (CNNs) as is. Therefore, lowering the inference time and memory consumption of large models would benefit the wider adoption of deep learning models for these real-world applications.

There has been a significant amount of work on reducing the inference cost via a more efficient network architecture [5, 7, 27, 43]. Although much progress has been made, the process of finding efficient architectures often requires a lot of experience and manual design. The automation of this process is an active research area of great value, as newer, more powerful DNN architectures are rapidly being discovered. Meanwhile, a large body of work has also been done in the direction of network compression, such as low-rank decomposition [5, 29, 30, 59], parameter quantization [22, 41, 61], and pruning [18, 27, 28, 31, 32, 56].

These compression techniques can be broadly divided into two categories: non-structured and structured pruning. Non-structured pruning methods [9–12] directly prune weights of CNN models to get sparse weight matrices. Along with compression, these approaches can provide faster network inference using specialized software [8] or hardware [8, 21, 39]. However, they often cause irregular memory access that adversely affects the network inference on general-purpose hardware or BLAS libraries. Structured pruning overcomes these limitations by removing structured weights like 2D kernels and feature maps (channels) or layers, which focuses on network architecture changes and can be supported by off-the-shelf deep learning libraries like PyTorch [40], and TensorFlow [1] etc. Pruning structured weights is a complex problem as we have to find out which layers or channels are least important in the network and removing them will not affect the accuracy significantly. For example, in VGG16 [47] the first hidden layer has 64 channels, an exhaustive search for which channels to remove without affecting accuracy involves 2^{64} combinations. For the whole network with around 5,000 channels in total, it is computationally infeasible to go through all combinations. To overcome this problem many techniques have been pro-

posed in the following categories.

Pruning using trained network statistics: These approaches mostly consist of two steps, first pruning the weights using network statistics and then fine-tuning the compressed network to recover from accuracy loss. Among these techniques, magnitude-based pruning has been proposed to find filters [27] or channels (feature maps) [18] with the lowest l_1 -norm as they are least informative and can be removed. Yang et al. [15] calculates the geometric median of each layer and prunes the filters close to it. Greedy pruning [16, 35] is also a common strategy that prunes filters by considering statistics computed from the next layer in a greedy layer-wise manner. More recently, Lin et al. [28] proposed to find the channel’s importance by calculating the rank of channels. All these methods calculate channel or filter importance at one or neighboring layers without considering all other layers in a deep neural network. Yu et al. [56] found that channels that were assumed less important and pruned in early layers could have a significant impact on the pruning of later layers. To address this issue, they calculate the importance score of the final layer and propagate it to each filter in the network to get its importance score.

Finding least important channels using trained network statistics is fairly straightforward, but finding how many channels or filters to drop from each layer (or the network) to achieve a certain compression rate requires manual effort. For example, in [28, 56] a fixed compression rate is used before fine-tuning which limits exploration for different compression rates, as each combination requires fine-tuning to recover from accuracy loss which could be time-consuming.

Pruning by learning channel importance: Another group of techniques is learning which channels are important by applying compression constraints while retraining the network, such as sparse scaling parameters in batch normalization layers [34, 52, 60]. Extension to these approaches has also taken into account of resource constraints such as latency and computation [6, 7, 38]. Huang et al. [20] and Lin et al. [32] introduced a new binary parameter to each channel (mask) to find out if the channel is important or not, but learning this binary mask involves an NP-hard optimization. This joint optimization for both compression and accuracy leads to higher accuracy performance but since the (combined) loss function is different than the original network, it requires specialized optimizers [20, 32] and needs iterative fine-tuning which increases time for both model training and manual adjustments.

Pruning by architecture search: Given that designing deep network architectures requires significant human effort, there have been some explorations on automatically learning network architectures. These approaches build networks from the ground up using a set of custom building blocks, relying on variations of a trial-and-error search to

find promising architectures [23, 33, 51, 52, 62], or compressed ones. In [3, 14], Reinforcement Learning (RL) is used to find compressed architectures but these techniques often take a long time to train. For example, in [3] a knowledge distillation [4, 17] approach is used to find a compressed student network from a teacher network using RL. The method utilizes the architecture search to reduce the depth of a network and size of layers which can take up to 2,500 epochs. This search cost is however prohibitively expensive for regular deep learning practitioners and further exacerbated by the growth in the dataset and network complexities. To reduce the architecture search time, generative adversarial networks have been used in [32] to find compressed structures. It shortens the search time significantly but still needs many iterations to train both generative and adversarial networks. Although these methods have led to a high compression rate, they often require a large amount of GPU time to obtain high-performing architectures.

Although the approaches mentioned above provide state-of-the-art results, they are computationally expensive and time-consuming as they require multiple iterations of pruning and fine-tuning, and often need a lot of manual effort. In this paper, we present DECORE, a multi-agent reinforcement learning (RL) framework for network compression which overcomes these limitations. In RL, an agent learns to perform a sequence of actions in an environment that maximizes some notion of a cumulative reward [49]. In our approach, we assign dedicated agents to all channels in the network which take actions to drop or keep a channel in the network. The agent gets a positive reward when it drops the channel for higher compression, but it gets a negative reward (penalty) if the accuracy is decreased due to pruning. Using the REINFORCE algorithm [54], we optimize agents’ policies to find out which channels to drop at each layer without affecting accuracy significantly and in turn maximizing the reward. Training of agents using the REINFORCE algorithm is independent of the network training loss which helps speed up the compression process.

Our main contributions are summarized as follows:

- We propose DECORE, a flexible and powerful approach to automating structure search and model compression. DECORE learns which channels are important by jointly training the network to provide high compression and FLOPs reduction rates.
- DECORE assigns an agent to each channel in the network (multi-agent learning) while each agent learns only one parameter, as opposed to other RL-based methods where the policy involves training another neural network. Learning a single parameter for each agent makes training much faster and more efficient.
- We demonstrate that our approach is able to find the

most important channels in the network which can be used to search for compressed network architectures.

- Results on both CIFAR-10 and ImageNet datasets are reported with various networks to demonstrate that DECORE is able to achieve better accuracy, as well as higher compression and acceleration (FLOPs reduction) rates, compared to other existing methods.

2. Related Work

2.1. Network Pruning

Apart from weight pruning as discussed in section 1, a few approaches have also been proposed to leverage sparsity in networks by removing entire channels [27, 44, 53]. Other approaches prune channels by analyzing trained network statistics [46, 57], or learning the importance of each channel or filter in the network [7, 20, 28, 34]. Our approach is similar to these methods in terms of operating on entire channels, but using RL to identify important channels instead of relying on network statistics. Moreover, our approach learns channel importance while training the network without changing the loss function like in [20, 32], which speeds up the overall training process.

2.2. Architecture Search

Modern architecture search approaches have led to state-of-the-art results for many tasks, but often require a significant amount of GPU time to obtain a high-performing architecture [3, 14, 33, 51, 52, 62]. Our use case is different, as we take already trained architectures and search for compressed architecture while maintaining accuracy. This drastically reduces the search space and thus makes our approach more efficient.

Reinforcement learning has also been used to find compressed architectures from trained networks [3, 14, 58]. These approaches focus on finding a fixed set of custom rules for architecture design and then fine-tune the best architecture found. Although the result is a compressed network, these approaches might still take a large amount of GPU hours to train. However, our proposed approach does not use any custom rules to look for compressed architectures. Like some of the regularization approaches, we simply learn which channels are more or less important for prediction. We obtain a compressed network by removing the channels that have little effect on accuracy.

3. Approach

3.1. DNN compression in RL framework

A Markov Decision Process (MDP) provides a mathematical framework to model decision making in stochastic environments. At each time step, t , the agent is at state s_t ,

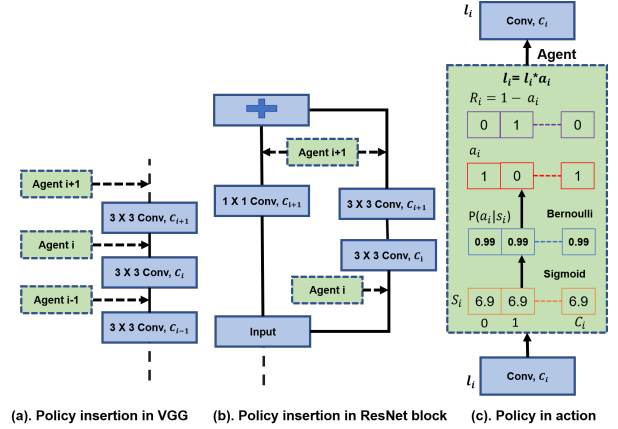


Figure 1. **DECORE Policies** (a) The policy is inserted before each convolution layer of VGG. (b) For ResNet blocks, we insert the policy sequentially (agent i) and also in the parallel path (agent $i+1$). We keep the same policy (agent $i+1$) for the last layer of parallel paths so that it doesn't affect future operations. (c) Our agent samples actions a_i and multiplies them with layer l_i to drop channels and compute the reward R_i .

where it can choose any available action a_t and move to the next state s_{t+1} , while accruing reward $R(s_t, a_t)$. The probability of moving to a new state s_{t+1} , given s_t and a_t , is determined by the state transition probability $P(s_{t+1} | s_t, a_t)$. The objective of RL is to find a policy $\pi(a_t | s_t)$, which specifies the action an agent will take at s_t to maximize some notion of cumulative reward [49].

We can therefore frame the network compression process as an RL problem. At each layer l_i with C_i channels, we assign dedicated agents to all channels in that layer. Each agent decides to drop or keep its channel and subsequently receives reward $R_{i,j}$ (channel j at layer i). The neural network does inference at layer l_i with channels which are not dropped and reaches the next state/layer l_{i+1} . Dropped channels can then be pruned from the network to create a compressed network. We train agents with a reward that incentivizes dropping channels at each layer l_i without affecting network performance significantly. This creates a multi-agent RL framework. As compared to a traditional MDP, the number of time steps (t) for each agent in this framework is just 1, as at each forward pass they take a decision to keep or drop channel only once and receive a reward based on that decision.

3.2. State Representation

We model each layer of a neural network with a channel vector $s_i \in \mathbb{R}^{C_i}$, where C_i is the number of channels in layer i . Each element w_j of s_i represents the weight (importance) of channel j in layer i . A high value of w_j means that channel j is very important and dropping it would affect the network accuracy significantly. We initialize all channel vectors s_i with the same value and our algorithm learns to

increase the weight values of channels that lead to high network accuracy and reduce the weight of those that do not. After training, low weight channels are removed from the network, resulting in a compressed model.

Our state representation does not take into account any other layer dependent variables, such as hidden values, height, width, number of input/output channels, kernel size, etc. [3, 14], which makes this approach very general. For example, this approach could be used to find important input channels for network training [55] or removing channels only from early layers which have a high memory consumption and DRAM access time [2].

3.3. Policy Representation

The agent makes the decision at each layer i to maximize the reward, based on a simple policy:

$$p_j = \frac{1}{1+e^{-w_j}}$$

$$\pi_j = \begin{cases} 1 & \text{with } p_j \\ 0 & \text{with } 1 - p_j \end{cases}$$

$$a_i = \pi_i = \{\pi_0, \pi_1, \dots, \pi_{C_i}\}$$

We convert the weights into probabilities via a sigmoid function and then choose which channels to keep at layer i by taking independent Bernoulli samples. Sample values are either 0 or 1, where 0 means dropping the channel and 1 keeping it. During training, we multiply each layer’s activation l_i with a_i to drop or keep channels. This is similar to dropout [48] but here the probability of dropping the channel depends on the agent weight. There is a single parameter (weight) for each agent to learn. This is in contrast to policies parameterized by neural networks in deep RL [36, 37], which makes our approach simple and fast as we don’t need to train dedicated neural networks to find good policies.

During training, we initialize the state vectors with a value of 6.9, which corresponds to a value of 0.99 after applying the sigmoid function. We assume initially that all channels in a pre-trained network are important and thus the initial probability of keeping a channel should be high. Each channel has a 99% chance of being selected during sampling and 1% chance not, which encourages little exploration and assists policy optimization. Figure 1 shows our policy in action for a layer l_i .

We insert policies only before convolution or linear layers of the network. We avoid inserting policies in batch normalization or depth wise convolution layers, as they operate on individual channels and the decision to drop channels could be affected by the next layer’s policy. Our policy insertion also depends on the network architecture. For example, ResNet or MobileNet architectures consist of residual blocks where computation happens in parallel paths and they are merged with a channel-wise addition. For these blocks, we keep the same policy at the end of parallel

paths so that the same channels are dropped before merging. Figure 1 shows the policy insertion for VGG [47] and ResNet [13] architectures. We use a similar approach to insert policies in computation layers of diverse architectures like GoogleLeNet [50] and DenseNet [19].

3.4. Reward

The objective of our modeling is to compress the network while maintaining predictive accuracy. We combine two different rewards to balance these considerations.

At each layer i , dropped channels represent the compression achieved at that layer. We compute the compression reward $R_{i,C}$ at layer i by summing all the dropped channels in that layer. We compute the number of dropped channels using the layer’s action vector a_i :

$$R_{i,C} = \sum_{j=1}^{C_i} 1 - a_{i,j}$$

After taking a series of actions at each layer, the network makes a prediction, \hat{y} , using only information from preserved channels. Each correct prediction suggests that the agent has taken right actions and thus should be rewarded. An incorrect prediction incurs a penalty, λ , which can be tuned to adjust the preference for compression or accuracy:

$$R_{acc} = \begin{cases} 1 & \text{if } \hat{y} = y \\ -\lambda & \text{else} \end{cases}$$

The penalty is a high negative value which forces the agent not to drop the channels if predictions are wrong. Our final reward at each layer combines these rewards via multiplication. The value is high only when both compression and accuracy rewards are high,

$$R_i = R_{i,C} * R_{acc}$$

3.5. Optimization with Reinforcement Learning

The agent must learn to take decisions to maximize the expected sum of rewards across all layers. We can estimate the expected reward using trajectories τ sampled from the policy. In our framework, the agent takes the decision only once so the trajectory has a single time step. For example, the trajectory for layer i is $\tau_i = (s_i, a_i)$. The optimization objective in this case is:

$$w^* = \arg \max_w J(w) = \sum_{i=1}^L \arg \max_w E_{\tau_i \sim \pi_w(\tau_i)} [R_i]$$

We use the REINFORCE policy gradient algorithm [36, 54] to learn policy parameters w ,

$$\nabla_w J(w) = \sum_{i=1}^L \nabla_w E_{\tau_i \sim \pi_w(\tau_i)} [R_i]$$

$$\nabla_w J(w) = \sum_{i=1}^L E_{\tau_i \sim \pi_w(\tau_i)} [\nabla_w \log \pi_w(a_i | s_i) R_i]$$

$$\nabla_w J(w) = \frac{1}{N} \sum_{b=1}^N \sum_{i=1}^L [\nabla_w \log \pi_w(a_{b,i} | s_{b,i}) R_{b,i}]$$

where N is the batch size used to estimate the gradient and L is a depth of the network. The gradients are used to update the policy or weights at each layer. We train policies and neural networks jointly using their own loss functions.

In traditional RL approaches the state representation is stochastic and the agent takes an action by looking at each state to learn policies. This is contrary to our approach as we try to learn state representation without requiring expensive policy optimization methods like Actor-Critic [25] and PPO [45]. Given that each agent has only one parameter, the policy gradient approach learns fast and efficiently.

4. Experiments

4.1. Datasets and Model Architectures

We evaluate our approach by measuring the accuracy, compression, and acceleration (FLOPs reduction) rates of standard deep networks on common image classification benchmark datasets, such as CIFAR10 [26] and ImageNet [42]. To demonstrate effectiveness of our approach we experimented with multiple widely-used model architectures, including VGG [47] with sequential layers, ResNet [13] with residual blocks, GoogleLeNet [50] with inception module, and DenseNet [19] with dense block.

4.2. Training Details

We used PyTorch [40] to implement the proposed method but our approach is very flexible and can be integrated with any deep learning library like TensorFlow [1]. To train policies, we used the ADAM optimizer [24] with a batch size of 256 and a learning rate of 0.01. Training begins with fixed policy vectors of an initial value 6.9 (Section 3.3), yet different subnetworks are generated in each forward pass due to Bernoulli sampling. For each correct prediction, the agent gets a positive reward, and for each wrong prediction it is penalized. We found the value of the penalty hyperparameter λ through a grid search, running each experiment for 10 epochs. The high penalty leads to better accuracy as the agent is more penalized for every wrong prediction. As we lower the penalty it tries to compress the network more without largely affecting accuracy.

Additionally, we also fine-tune the network at each optimization step so that it may recover from drops in performance after pruning. Joint training of policies and the network is more computationally efficient and sample efficient than approaches that separate learning from search [27, 28, 52]. On the CIFAR dataset, all models are trained for 300 epochs. We stopped policy training after 260 epochs and dropped all the channels with a probability of less than 50%. We continued fine-tuning the remaining subnetwork in the next 40 epochs to improve accuracy. Finally, we removed the dropped channels from the network to get the compressed model. On the ImageNet dataset, models were

trained for just 30 epochs. We stopped policy training after 20 epochs and fine-tune in the remaining epochs.

4.3. Results and Analysis

4.3.1 Results on CIFAR-10

We evaluated the performance of the proposed approach on CIFAR-10 with four diverse networks, VGGNet-16/19, DenseNet-40, GoogleLeNet, and ResNet-56/110. We used VGGNet with the variation of the original VGG-16 for CIFAR-10 similar to [27, 28, 32]. For GoogLeNet, the final output layer was changed to have 10 classes. DenseNet-40 has 40 layers with a growth rate of 12. For a fair comparison, we used the same trained networks used by [28, 32]. For all our experiments, we tried to adjust the penalty parameter (λ) to achieve similar accuracy as reported state-of-the-art results and compared the parameters/FLOPs reduction rates.

VGGNet: Table 1 shows that with a higher penalty ($\lambda = 500$), our approach **DECORE-500** was able to achieve 0.15% higher accuracy than the baseline model with a 63% parameter and 35% FLOPs reduction rate, which shows that this network was over parameterized for this problem and DECORE can be used as a regularizer. As compared to approaches that prune networks using network statistics [27, 28], our approach provides higher compression and acceleration (FLOPs reduction) rates at similar or higher accuracy. As compared to L1 pruning [27], our approach DECORE-200 was able to get a 25% higher compression rate (89% vs 64%) and a 30.5% higher FLOPs reduction rate (64.8% vs 34.3%) at 0.16% higher accuracy (93.56% vs 93.40%). Similarly, as compared to HRANK [28], our approach provides better results at various compression rates. For example, with a high penalty, DECORE-200 was able to achieve 6.1% higher compression and 11.3% higher FLOPs reduction rates at 0.13% higher accuracy (93.56% vs 93.43%) compared to HRANK-1. Similarly, with a low penalty, DECORE-50 was able to achieve 6.3% higher compression and 11.8% higher FLOPs reduction rates at 0.45% higher accuracy (91.68% vs 91.23%) as compared to HRANK-3. Moreover, HRANK [28] manually sets different compression rates at each layer for higher compression or accuracy, and our method automatically (just by changing penalty λ) finds an optimal setting and generates better results overall. For the training process, HRANK trains each layer for 30 epochs (480 epochs for 16 layers), and the training time increases as the number of layers increases, while our approach only needs to train for 300 epochs for all settings.

Our approach also finds better results than existing methods which learn channel importance [20, 60]. For instance, as compared to SSS [20], DECORE-200 achieved 15.2% higher compression and 23.2% higher FLOPs reduction rates at 0.54% higher accuracy (93.56% vs 93.02%). As compared to Zhao et al. [60], DECORE-200 achieved

Model	Top-1(%)	Params(PR)	FLOPs(PR)
VGG16	93.96	14.98M(0.0%)	313.73M(0.0%)
DECORE-500	94.02	5.54M(63.0%)	203.08M(35.3%)
L1 [27]	93.40	5.40M(64.0%)	206.00M(34.3%)
SSS* [20]	93.02	3.93M(73.8%)	183.13M(41.6%)
Zhao et al. [60]	93.18	3.92M(73.3%)	190.00M(39.1%)
HRank-1 [28]	93.43	2.51M(82.9%)	145.61M(53.5%)
DECORE-200	93.56	1.66M(89.0%)	110.51M(64.8%)
GAL-0.05 [32]	92.03	3.36M(77.6%)	189.49M(39.6%)
HRank-2 [28]	92.34	2.64M(82.1%)	108.61M(65.3%)
DECORE-100	92.44	0.51M(96.6%)	51.20M(81.5%)
GAL-0.1 [32]	90.73	2.67M(82.2%)	171.89M(45.2%)
HRank-3 [28]	91.23	1.78M(92.0%)	73.70M(76.5%)
DECORE-50	91.68	0.26M(98.3%)	36.85M(88.3%)
VGG19	93.76	20.0M(0.0%)	398.44M(0.0%)
N2N [3]	91.64	1.0M(95.0%)	-
DECORE-40	91.65	0.3M(98.5%)	43.9M(89.0%)

Table 1. Pruning results of VGGNet on CIFAR10. In all tables and figures, PR means pruned rate, GAL- X refers to GAL with sparsity factor X , DECORE- λ is DECORE with penalty λ , HRANK- N refers to HRANK with a compression rate setting N , and M/B means millions/billions. For a fair comparison, we provide multiple results of HRANK, GAL and others at various compression rates. SSS* shows the results standardized by HRANK and GAL.

15.7% higher compression and 25.7% higher FLOPs reduction rates at 0.38% higher accuracy (93.56% vs 93.18%).

Compared to architecture search approaches such as GAL [32], DECORE-100 provided 19% higher compression and 41.9% FLOPs reduction rates at 0.41% higher accuracy (GAL-0.05). For high compression (low penalty), DECORE-50 was able to achieve 16.1% more compression and 43.1% FLOPs reduction rates at 0.95% higher accuracy than GAL-0.1. As compared to RL based architecture search method N2N [3], DECORE achieves 3.5% higher compression on the VGG-19 network at similar accuracy. N2N [3] training can take more than 2,500 epochs to get the best architecture. DECORE achieves similar accuracy and higher compression with just 300 epochs ($< 9\%$).

GoogleLeNet: Table 2 shows at higher penalty DECORE-500 was able to find 0.15% higher accuracy than the baseline model with 23% compression and 19.8% FLOPs reduction rates. DECORE-200 was able to find more compression (38% vs L1, 25.5% vs HRANK-1) and lower FLOPs (45.6% vs L1, 23.6% vs HRANK-1) at similar accuracy, compared to approaches using network statistics, such as L1 [27] and HRANK [28]. In a low penalty setting, DECORE-175 provided 16.3% higher compression and 14.3% FLOPs reduction as compared to HRANK-2.

Compared to architecture search methods, DECORE-175 provided more compression (36.8% than GAL-0.05 and 32.4% than GAL-ApoZ), lower FLOPs (46.5% than GAL-0.05 and 34.7% than GAL-ApoZ), and higher accuracy (0.4% than GAL-0.05 and 2.22% than GAL-ApoZ). It shows that DECORE can be applied to inception modules

Model	Top-1(%)	Params(PR)	FLOPs(PR)
GoogLeNet	95.05	6.15M(0.0%)	1.52B(0.0%)
DECORE-500	95.20	4.73M(23.0%)	1.22B(19.8%)
L1* [27]	94.54	3.51M(42.9%)	1.02B(32.9%)
HRank-1 [28]	94.53	2.74M(55.4%)	0.69B(54.9%)
DECORE-200	94.51	1.17M(80.9%)	0.33B(78.5%)
GAL-ApoZ* [18]	92.11	2.85M(53.7%)	0.76B(50.0%)
GAL-0.05 [32]	93.93	3.12M(49.3%)	0.94B(38.2%)
HRank-2 [28]	94.07	1.86M(69.8%)	0.45B(70.4%)
DECORE-175	94.33	0.86M(86.1%)	0.23B(84.7%)

Table 2. Pruning results of GoogleLeNet on CIFAR10. L1* and GAL-ApoZ* are results standardized by HRANK and GAL

Model	Top-1(%)	Params(PR)	FLOPs(PR)
DenseNet-40	94.81	1.04M(0.0%)	282.92M(0.0%)
DECORE-175	94.85	0.83M(20.7%)	228.96M(19.1%)
Liu et al.-40% [34]	94.81	0.66M(36.5%)	190.00M(32.8%)
GAL-0.01 [32]	94.29	0.67M(35.6%)	182.92M(35.3%)
HRank-1 [28]	94.24	0.66M(36.5%)	167.41M(40.8%)
DECORE-115	94.59	0.56M(46.0%)	171.36M(39.4%)
Zhao et al. [60]	93.16	0.42M(59.7%)	156.00M(44.8%)
GAL-0.05 [32]	93.53	0.45M(56.7%)	128.11M(54.7%)
HRank-2 [28]	93.68	0.48M(53.8%)	110.15M(61.0%)
DECORE-70	94.04	0.37M(65.0%)	128.13M(54.7%)

Table 3. Pruning results of DenseNet on CIFAR10. Liu et al.- $\alpha\%$ means α percentage of parameters are pruned.

to achieve same performance as state-of-the-art approaches.

DenseNet-40: DenseNet is a sophisticated architecture with channel concat (a growth rate of 12) at each layer, and channels in one layer are used in the following layers. Dropping even a single channel in this architecture means removing this channel from all the following layers. Table 3 shows DECORE results for the DenseNet architecture. With a higher penalty, DECORE-175 was able to achieve 20.7% compression and 19.1% FLOPs reduction rates at similar accuracy as the baseline model. With a slightly lower penalty, DECORE-115 was able to compress baseline by 46% and reduce FLOPs by 39.4% with just a 0.22% accuracy drop, better results compared to other techniques. In a high compression setting, DECORE-70 achieved 65.0% compression and 54.7% FLOPs reduction with just a 0.77% drop in accuracy. These results demonstrate that our approach can be applied to complex architectures too.

ResNet-56/110: Table 4 summarizes results on the ResNet-56/110 architecture. For ResNet-56, DECORE-200 provided 49% compression and 49.9% FLOPs reduction rates without any accuracy loss. Upon a 50% FLOPs constraint, the RL-based architecture search method AMC [14] finds the best network with 91.9% accuracy which is 1.4% less than DECORE-200. In a high compression setting (low penalty), DECORE-55 compressed the network by 85.3% and reduced FLOPs by 81.5% with a 2.4% accuracy drop. As shown in the table, DECORE provided more compres-

Table 4. Pruning results of ResNet-56/110 on CIFAR10.

Model	Top-1(%)	Params(PR)	FLOPs(PR)
ResNet-56	93.26	0.85M(0.0%)	125.49M(0.0%)
DECORE-450	93.34	0.64M(24.2%)	92.48M(26.3%)
L1 [27]	93.06	0.73M(14.1%)	90.90M(27.6%)
NISP [56]	93.01	0.49M(42.4%)	81.00M(35.5%)
GAL-0.6 [32]	92.98	0.75M(11.8%)	78.30M(37.6%)
HRank-1 [28]	93.17	0.49M(42.4%)	62.72M(50.0%)
AMC [14]	91.90	-	62.72M(50.0%)
DECORE-200	93.26	0.43M(49.0%)	62.93M(49.9%)
He et al. [16]	90.80	-	62.00M(50.6%)
GAL-0.8 [32]	90.36	0.29M(65.9%)	49.99M(60.2%)
HRank-2 [28]	90.72	0.27M(68.1%)	32.52M(74.1%)
DECORE-55	90.85	0.13M(85.3%)	23.22M(81.5%)
ResNet-110	93.50	1.72M(0.0%)	252.89M(0.0%)
DECORE-500	93.88	1.11M(35.7%)	163.30M(35.4%)
L1 [27]	93.30	1.16M(32.6%)	155.00M(38.7%)
GAL-0.5 [32]	92.55	0.95M(44.8%)	130.20M(48.5%)
HRank-1 [28]	93.36	0.70M(59.2%)	105.70M(58.2%)
DECORE-300	93.50	0.61M(64.8%)	96.66M(61.8%)
HRank-2 [28]	92.65	0.53M(68.7%)	79.30M(68.6%)
DECORE-175	92.71	0.35M(79.6%)	58.37M(76.9%)

sion and lower FLOPs rates than other approaches.

Similarly for ResNet-110, DECORE-300 provided 64.8% compression and 61.8% FLOPs reduction without accuracy loss, which demonstrates more compression and lower FLOPs rates than other approaches while achieving higher accuracy. DECORE-175 compressed the network even further to 79.6% with 76.9% FLOPs reduction and just a 0.79% drop in accuracy.

4.3.2 Results on ImageNet

Table 5 shows our performance with different penalty values on the more challenging dataset, ImageNet. DECORE-8 shows at higher penalty ($\lambda = 8$), our proposed method was able to achieve 11.0% compression (25.50M to 22.69M) and 13.5% (4.09B to 3.54B) FLOPs reduction at a 0.16%/0.15% Top-1/Top-5 accuracy improvement.

With a slightly lower penalty ($\lambda = 6$), DECORE-6 was able to achieve 74.58%/92.18% Top-1/Top-5 accuracy (1.57%/0.69% drop as compared to the baseline model) with the number of parameters reduced to 14.10M (44.71% of the baseline) and FLOPs reduced to 2.36B (42.30% of the baseline). Compared to network statistics based approaches such as ThiNet-70 [35], DECORE-6 achieved 2.54% higher Top-1 accuracy and a 11.14% higher compression rate with similar FLOPs. As compared to other methods that learn channel importance, such as SSS-32 [20] and He et al. [16], DECORE-6 provided higher accuracy, more compression, and lower FLOPs too.

Furthermore, DECORE-6 achieved 27.85% higher compression and 2.63% (74.58% vs 71.95%) higher Top-1 accuracy than architecture search methods such as GAL-

Table 5. Pruning results of ResNet-50 on ImageNet. X-joint means jointly pruning channels and blocks.

Model	Top-1(%)	Top-5(%)	Params	FLOPs
ResNet-50	76.15	92.87	25.50M	4.09B
DECORE-8	76.31	93.02	22.69M	3.54B
SSS-32 [20]	74.18	91.91	18.60M	2.82B
He et al. [16]	72.30	90.80	-	2.73B
ThiNet-70 [35]	72.04	90.67	16.94M	2.44B
GAL-0.5 [32]	71.95	90.94	21.20M	2.33B
DECORE-6	74.58	92.18	14.10M	2.36B
GDP-0.6 [31]	71.19	90.71	-	1.88B
GDP-0.5 [31]	69.58	90.14	-	1.57B
SSS-26 [20]	71.82	90.79	15.60M	2.33B
GAL-1 [32]	69.88	89.75	14.67M	1.58B
GAL-0.5-joint [32]	71.80	90.82	19.31M	1.84B
HRank-1 [28]	71.98	91.01	13.77M	1.55B
ThiNet-50 [35]	70.01	90.02	12.38M	1.71B
DECORE-5	72.06	90.82	8.87M	1.60B
ThiNet-30 [35]	68.42	88.30	8.66M	1.10B
GAL-1-joint [32]	69.31	89.12	10.21M	1.11B
HRank-2 [28]	69.10	89.58	8.27M	0.98B
DECORE-4	69.71	89.37	6.12M	1.19B

0.5 [32]. It is worth noting that DECORE was able to find better performance in just 30 epochs which is much faster than architecture search approaches such as GAL (trains generator and discriminator in each batch) and pruning-based methods such as HRANK (90 epochs).

Another interesting observation is that on a small dataset like CIFAR10, pre-trained models often have high accuracy and it could become challenging to prune these models without losing the accuracy. To meet this end the λ parameter is set high so that we can achieve compression without any noticeable accuracy drop. On the flip side, for more complex datasets like ImageNet, λ can have a lower value.

5. Analysis: Does DECORE find important channels?

To investigate the effect of learned policies on network compression, we reviewed policies of the VGG16 network trained on CIFAR10 for 200 epochs with 4,736 agents (just for convolution and linear layers). We didn't train the network jointly in this experiment so that the analysis of learned policies is not mixed with network update.

After training the agents, we remove the channels with high weight values and observe how much it changes the accuracy. If the model accuracy drops significantly without high weight channels, it suggests that the RL algorithm has correctly identified the most important channels in the network. Figure 2 shows the effect on network accuracy after low and high weight channels are removed at each layer. The pre-trained network (without any channel drop) has an accuracy of 93.96%. By removing the highest weight channel only (1 channel out of 4736), there is a noticeable accuracy drop of 1.42% (92.54% vs 93.96%). This decrease is




Figure 2. **DECORE finds important weights** After training VGG16 with DECORE, we remove high weight channels and observe significant drops in accuracy. On the flip side, removing low weight channels has a little impact on accuracy, and thus they can be dropped to compress the model.

relatively large, given that removing this channel only compresses the network by 0.004%. It suggests that DECORE is able to learn the most important channels, as dropping even one of them will greatly impact accuracy. In contrast, if we remove the lowest weight channel, the accuracy drop is negligible. For instance, if we remove 20 highest weight channels the network loses half of its accuracy (47.63%), and after dropping top 50 channels (just 1% of all channels) the accuracy remains only 13% which can almost be considered as a random network. However, dropping 50 of the lowest weight channels does not impact accuracy much. This illustrates that DECORE is able to identify channels with the least and most importance.

6. DECORE for compressed Network Architecture Search

In real-world scenarios, there are often stringent limits on resources available. Although we can not directly optimize the network with hard constraints, the penalty reward in our approach can be tuned to ensure trained models falling within those ranges. As shown in section 5, DECORE learns the weight value for each channel in the network which reflects the importance of that channel. We also noticed that dropping low weight channels doesn't really impact the accuracy much, which means it's possible to sort all agents' weight values and drop the lowest weight channels till the constraint is met. This simple yet effective approach can help meet various types of requirements, eliminating the need to train networks with hard constraints. In this experiment, we explore the effect of hard constraints on performance, including our findings on the MobileNetV2 [43] architecture which is hand designed for real-world applications. MobileNetV2 has 2.28M parameters which require around 8.7MB of on-device memory and 86.7M FLOPs during the inference on the CIFAR-10

Table 6. Top-1 accuracy with Size/FLOPs constraint on CIFAR

Model(Top-1%)	Size (MB)			FLOPs (M)		
	0.5MB	1MB	5MB	25M	50M	100M
VGG16	89.74	91.30	93.05	88.32	91.54	93.09
ResNet-56	91.52	92.56	93.26	91.59	92.60	93.31
MobileNetV2	90.97	91.95	92.58	92.01	92.31	92.58

dataset. As a reference, the baseline MobileNetV2 model we are using has 92.58% Top-1 accuracy for CIFAR-10.

In all experiments, we learned the policies for 200 epochs and then fine-tuned the compressed network for 60 epochs. Similar to section 5 we didn't train the network jointly while learning policies. Table 6 shows the performance drop of different models on the CIFAR10 dataset with different size constraints (in terms of memory storage and FLOPs). As explained above, we dropped the low weight channels till the constraint is satisfied. With a storage constraint of just 1MB, there is a maximum of 2.66% performance drop across all models. When we decrease the size limit to 0.5MB, there is a negligible performance difference than the 1MB results. This illustrates that our technique is robust to varying memory constraints.

In addition to size constraints, we explore constraints on the number of FLOPs in a compressed network. FLOPs are highly correlated with latency on CPU-only devices and thus the results in Table 6 reflect performance with potential latency constraints. With a very low 50M FLOP limit, we observed a maximum performance drop of 2.42% across all models. These results suggest that our approach can flexibly handle various hard constraints on computation, and by proxy latency, without a significant drop in performance. More stringent constraints (e.g., size $< 0.01MB$ or FLOPs $< 10M$) lead to degenerate architectures and joint training (Section 4.2) is recommended for these scenarios.

Last but not least, DECORE identifies the importance of each channel in the network once and then remove low priority channels to obtain a target architecture based on constraints. Compared to other methods [3, 14], it doesn't need to retrain the network when the target criteria is changed.

7. Conclusion

We present DECORE as an efficient reinforcement learning approach to network compression and compressed architecture search. DECORE can find a trade-off between accuracy and compression by tuning the penalty (just 1 hyperparameter) for misclassifications, allowing for more control over the compressed model structure. Once trained the proposed method can also be used to search network architectures with different resource constraints without retraining for those constraints. Exploring other reward functions that target different resource constraints is an interesting direction to explore in future work, along with network architectures for other visual reasoning tasks such as object detection, semantic segmentation, etc.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. 2016. [1](#), [5](#)
- [2] Manoj Alwani, Han Chen, Michael Ferdman, and Peter Milder. Fused-layer CNN accelerators. In *49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2016. [4](#)
- [3] Anubhav Ashok, Nicholas Rhinehart, Fares Beainy, and Kris M. Kitani. N2N learning: Network to network compression via policy gradient reinforcement learning. In *International Conference on Learning Representations*, 2018. [2](#), [3](#), [4](#), [6](#), [8](#)
- [4] Lei Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? *arXiv:1312.6184*, 2013. [2](#)
- [5] Emily L. Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in Neural Information Processing Systems*, pages 1269–1277, 2014. [1](#)
- [6] Edgar Galván and Peter Mooney. Neuroevolution in deep neural networks: Current trends and future challenges. *IEEE Transactions on Artificial Intelligence*, 2021. [2](#)
- [7] Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, and Edward Choi. Morphnet: Fast & simple resource-constrained structure learning of deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [1](#), [2](#), [3](#)
- [8] 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. *ACM SIGARCH Computer Architecture News*, 44(3):243–254, 2016. [1](#)
- [9] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv:1510.00149*, 2015. [1](#)
- [10] Song Han, Jeff Pool, John Tran, and William J Dally. Learning both weights and connections for efficient neural networks. *arXiv:1506.02626*, 2015. [1](#)
- [11] Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *NIPS*, 1993. [1](#)
- [12] Babak Hassibi, David G Stork, and Gregory J Wolff. Optimal brain surgeon and general network pruning. In *IEEE international conference on neural networks*, pages 293–299, 1993. [1](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv:1512.03385*, 2015. [1](#), [4](#), [5](#)
- [14] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. AMC: AutoML for model compression and acceleration on mobile devices. In *European Conference on Computer Vision (ECCV)*, September 2018. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [15] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#)
- [16] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. [2](#), [7](#)
- [17] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv:1503.02531*, 2015. [2](#)
- [18] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. *arXiv:1607.03250*, 2016. [1](#), [2](#), [6](#)
- [19] G. Huang, Z. Liu, K. Q. Weinberger, and L. van der Maaten. Densely connected convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [4](#), [5](#)
- [20] Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 304–320, 2018. [2](#), [3](#), [5](#), [6](#), [7](#)
- [21] 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:1602.07360*, 2016. [1](#)
- [22] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *CVPR*, page 2704–2713, 2018. [1](#)
- [23] J. Jin, Z. Yan, K. Fu, N. Jiang, and C. Zhang. Neural network architecture optimization through submodularity and supermodularity. *arXiv:1609.00074*, 2016. [2](#)
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. [5](#)
- [25] Vijay R. Konda and John N. Tsitsiklis. Actor-critic algorithms. In *NIPS*, pages 1008–1014, 2000. [5](#)
- [26] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009. [5](#)
- [27] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [28] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. HRank: Filter pruning using high-rank feature map. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1529–1538, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [29] Shaohui Lin, Rongrong Ji, Chao Chen, Dacheng Tao, and Jiebo Luo. Holistic CNN compression via low-rank decomposition with knowledge transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. [1](#)

- [30] Shaohui Lin, Rongrong Ji, Xiaowei Guo, and Xuelong Li. Towards convolutional neural networks compression via global error reconstruction. *International Joint Conference on Artificial Intelligence*, 2016. 1
- [31] Shaohui Lin, Rongrong Ji, Yuchao Li, Yongjian Wu, Feiyue Huang, and Baochang Zhang. Accelerating convolutional networks via global and dynamic filter pruning. In *International Joint Conference on Artificial Intelligence*. 1, 7
- [32] Shaohui Lin, Rongrong Ji, Chenqian Yan, Baochang Zhang, Liujuan Cao, Qixiang Ye, Feiyue Huang, and David Doermann. Towards optimal structured CNN pruning via generative adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 5, 6, 7
- [33] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv:1806.09055*, 2018. 2, 3
- [34] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2736–2744, 2017. 2, 3, 6
- [35] Jianhao Luo, Jianxin Wu, and Weiyao Lin. ThiNet: A filter level pruning method for deep neural network compression. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 2, 7
- [36] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016. 4
- [37] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013. 4
- [38] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *arXiv:1611.06440*, 2016. 2
- [39] 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. *International Conference on Learning Representations*, 2017. 1
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019. 1, 5
- [41] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. XNOR-Net: ImageNet classification using binary convolutional neural networks. In *ECCV*, 2016. 1
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, (3), 2015. 1, 5
- [43] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Inverted residuals and linear bottlenecks mobile networks for classification, detection and segmentation. *arXiv:1801.04381*, 2018. 1, 8
- [44] Simone Scardapane, Danilo Comminiello, Amir Hussain, and Aurelio Uncini. Group sparse regularization for deep neural networks. *Neurocomputing*, 241:81–89, 2017. 3
- [45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. volume abs/1707.06347, 2017. 5
- [46] Noah Simon, Jerome Friedman, Trevor Hastie, and Robert Tibshirani. A sparse-group lasso. *Journal of computational and graphical statistics*, 22(2):231–245, 2013. 3
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *Proceedings of the International Conference on Learning Representations*, 2014. 1, 4, 5
- [48] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. 4
- [49] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018. 2, 3
- [50] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, and D. Erhan. Going deeper with convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 4, 5
- [51] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019. 2, 3
- [52] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019. 2, 3, 5
- [53] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. *arXiv:1608.03665*, 2016. 3
- [54] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992. 2, 4
- [55] Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. Learning in the frequency domain. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4
- [56] Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S Davis. NISP: Pruning networks using neuron importance score propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, page 9194–9203, 2018. 1, 2, 7

- [57] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006. [3](#)
- [58] Huixin Zhan and Yongcan Cao. Deep model compression via deep reinforcement learning, 2019. [3](#)
- [59] Xiangyu Zhang, Jianhua Zou, Xiang Ming, Kaiming He, and Jian Sun. Efficient and accurate approximations of nonlinear convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. [1](#)
- [60] Chenglong Zhao, Bingbing Ni, Jian Zhang, Qiwei Zhao, Wenjun Zhang, and Qi Tian. Variational convolutional neural network pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [5](#), [6](#)
- [61] Bohan Zhuang, Chunhua Shen, Mingkui Tan, Lingqiao Liu, and Ian Reid. Towards effective low-bitwidth convolutional neural networks. In *CVPR*, page 7920–7928, 2018. [1](#)
- [62] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv:1611.01578*, 2016. [2](#), [3](#)